Bálint Mucsányi, Michael Kirchhof, Elisa Nguyen,
Alexander Rubinstein, Seong Joon Oh

Trustworthy Machine Learning

Theory
Applications
Intuitions

Trustworthy Machine Learning

Bálint Mucsányi[1,*], Michael Kirchhof[1], Elisa Nguyen[1,2],
Alexander Rubinstein[1,2], and Seong Joon Oh[1,2]

[1]University of Tübingen
[2]Tübingen AI Center
[*]`balint.mucsanyi@student.uni-tuebingen.de`

Contents

Contents

Contents

Preface

Preface

As machine learning technology gets applied to actual products and solutions, new challenges have emerged. Models unexpectedly fail to generalize to small changes in the distribution; some models are found to utilize sensitive features that could treat certain demographic user groups unfairly; models tend to be confident on novel data they have never seen, or models cannot communicate the rationale behind their decisions effectively with the end users like medical staff to maximize the human-machine synergies. Collectively, we face a trustworthiness issue with the current machine learning technology. A large fraction of machine learning research nowadays is dedicated to expanding the frontier of Trustworthy Machine Learning (TML). TML has been an explicit topic in the call for papers of the ICML conference since 2020, and other relatively young conferences dealing with TML topics emerged like FAccT, or AIES.

This textbook on TML is an end product of the homonymous course at the University of Tübingen, first offered in the Winter Semester of 2022/23. The book covers a theoretical and technical background for key topics in TML as well as underlying intuitions. We conduct a critical review of important classical and contemporary research papers on related topics. The book is meant to be a stand-alone product accompanied by code snippets and various pointers to further sources on topics of TML.

The goal of this book is to prepare readers to critically read, assess, and discuss research work in TML. Through the provided code snippets, readers will gain the technical background to implement basic TML techniques and, eventually, conduct their own research in TML.

The book has the following prerequisites:

- Familiarity with Python and PyTorch coding.
- Basic knowledge of machine learning concepts and deep learning.
- Basic maths: multivariate calculus, linear algebra, probability, statistics, and optimization.

Throughout the book, definitions will be provided in blue boxes in the following form:

> **Definition 0.1: Mitochondrion**
>
> Mitochondria are the powerhouse of the cell.

These will be displayed right before encountering the notion in the text.

Similarly, yellow boxes will contain additional information that is not crucial for understanding the main concepts introduced in the book. An example is provided below.

> **Additional Information: Nests of Scarlet Tanagers**
>
> Nests of scarlet tanagers are typically built on horizontal tree branches. [207]

We hope that this book will pique readers' interest in TML and encourage them to contribute to this beautiful field of research.

Introduction to Trustworthy Machine Learning

1 Introduction to Trustworthy Machine Learning

1.1 Scale is all we need?

> **Definition 1.1: Generalization**
>
> An ML model generalizes well if the rules found on the training set can be applied to new test situations we are interested in.

The story of Machine Learning (ML) seems to be that a bigger model with more data implies better test loss, as shown in Figure 1.1. Such models generalize well. Of course, more computing resources are needed, but more prominent tech companies possess them.
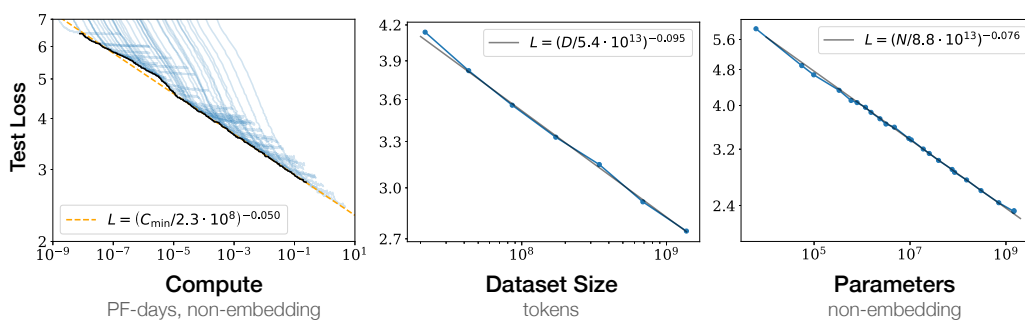
Figure 1.1: "Language modeling performance improves smoothly as we increase the model size, dataset [...] size, and amount of compute [with sufficiently small batch size] used for training. For optimal performance all three factors must be scaled up in tandem. Empirical performance has a power-law [i.e., $y = a \cdot x^b$] relationship with each individual factor when not bottlenecked by the two." [92]. 1 PF-day $= 10^{15} \cdot 24 \cdot 3600$ floating point operations. Figure taken from [92].

Between 2013 and 2020, there was a steady increase in ImageNet [43] top-1 accuracy (Figure 1.2). This increase slowed over time, and between 2020 and 2023, we see a plateau in the top-1 accuracy – seemingly, we "solved ImageNet."

1.1.1 Are we done with ML?

So, are we done with ML? If the reader's answer is 'yes', then the following questions naturally follow:

- Why do we not see ML used in every business?
- Why is ML not changing our lives yet?
- Why have we not gone through a quantum leap in productivity (results, profits, products) owing to ML?
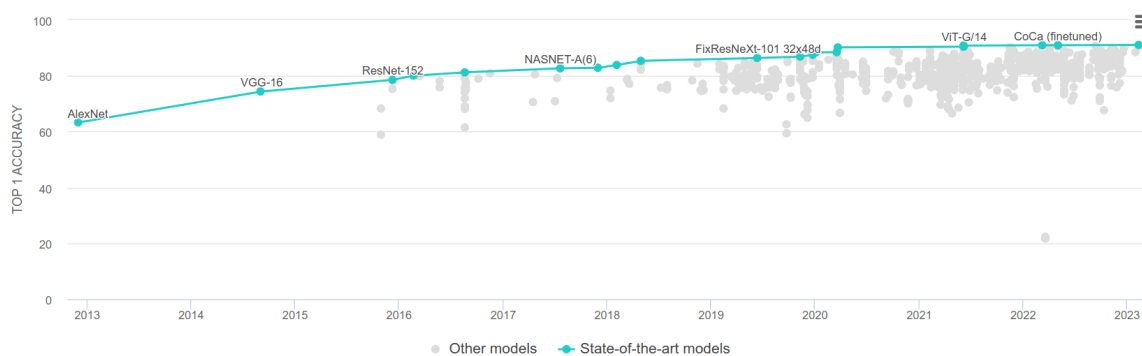
Figure 1.2: ImageNet top 1 accuracy leaderboard on 05.03.2023 [130]. The performance of state-of-the-art methods plateaued over time.

If the reader's answer is 'no', then we ask:

- What are the remaining challenges in ML?
- How can we capture and measure those challenges?

This book aims to answer these questions while showcasing current state-of-the-art approaches in the field of TML.

1.2 Key Limitations of ML

Our answer is 'no': Not all businesses use ML, and we have not yet gone through a quantum leap in productivity because of ML. Let us review the *fundamental limitations of ML*.

1.2.1 ML often does not work.

ML models *do* generalize, but not in the way one would expect. They tend to generalize well, given

1. sufficient amount of data,
2. appropriate inductive biases, and
3. if we stay in the *same distribution* as the training set (in-distribution (ID) generalization).

Our models, however, need to cope with *new situations* in practice. Whenever there are changes in the deployment conditions, our model will usually work *much worse*.

1.2.2 ML has high operating costs.

So, we usually need to constantly adjust our model to the new settings. This requires

1. an ML engineer ($\mathcal{O}$(100k USD/year)),
2. collecting fresh data (on dedicated pipelines) or buying specialized proprietary data, and
3. computing resources or credits for an ML cloud to adjust the models on the new data.

From a business perspective, these points boil down to a money issue. ML has high operating costs if our model constantly has to be adapted to new scenarios. If we had a model that generalized well, we would have less or even none of these costs.

1.2.3   ML is currently not trustworthy.

Even if we address the previous concerns, broad use of ML is not just a matter of whether our model works well or not – *it is difficult to trust ML models*. Extreme cases are when our *life*, *health*, or *money* is at stake.

Example 1: Ten AI doctors say we have stomach cancer and recommend chemo- and radio-therapy. Could we trust this diagnosis and start these treatments? The majority of people would want to have the cancer pointed out in the MRT images. This is an example of *explanability*.

Example 2: We are in a self-driving car driving through a curvy road along a cliff. Should we lift our hands off the wheel? Probably not. We likely *could* not even do that because these cars would insist on human intervention (e.g., by giving warning signs). The automatic detection of an unusual environment is an example of *uncertainty quantification*.

Example 3: It is also hard to trust images generated by DALL-E to be sensible: We often see absurd artifacts in otherwise great ML-generated art. This is a problem of *OOD generalization*, as our model only gives high-quality images for a restricted set of prompts.

1.3   Topics of the Book

The topics this book covers are as follows:

1. **Out-Of-Distribution (OOD) Generalization.** Can we train a model to work well beyond the training distribution?
2. **Explainability.** Can we make a model explain itself?
3. **Uncertainty.** Can we make a model know when it does not know?
4. **Evaluation.** How to quantify trustworthiness? How to measure progress?

The topics we do not cover but are also core parts of Trustworthy Machine Learning:

1. **Fairness.** Demographic disparity is a core concern of fairness, which is the difference between the proportion of rejects and accepts for each population subgroup. The use of sensitive attributes (often implicitly) is also a significant problem regarding trustworthiness.
2. **Privacy and Security.** Data are often proprietary and private. How to keep the data safe? Often we can reverse-engineer the original samples of the training set, e.g., in language models. This way, one can obtain sensitive, private information as well, e.g., medical records of patients.
3. **Abuse of AI tools.** One can use ML to create deepfakes, e.g., to swap faces of people. Disseminating falsehood, e.g., via Large Language Models (LLMs), is also an alarming problem.
4. **Environmental concern.** Accelerated computing consumes much energy.
5. **Governance.** It is important to regulate the use of AI and formalize boundaries of AI usage.

1.4 Trustworthiness: Transition from "What" to "How"

To give an introduction to trustworthiness in ML, let us first define the "What" and "How" parts of an ML problem.

> **Definition 1.2: The "What" Part of a Problem**
>
> The "What" part of a problem is learning the task we want to solve, i.e., the relationship between X and Y. For example, the "What" part might be categorizing images into classes. The "What" point of view is that predicting Y given X is sufficient.

> **Definition 1.3: The "How" Part of a Problem**
>
> The "How" part of a problem specifies how a system comes to its prediction, what cues it is basing its decision-making on, and how it reasons about the prediction.
> For robust AI systems, whether we solve a problem is not enough. How we solve it matters more.

We currently have a "What" to "How" *paradigm shift* in ML. Solving the "What" part is often *not enough*, as detailed in the following section.

1.4.1 Why Solving "What" is Not Enough

A model can use multiple *recognition cues* Z to make its prediction. These cues determine what the model bases its prediction on and what it exploits. There are *two categories* of cues:

1. **Causal, robust cue.** Such cues are robust to environmental changes, as the prediction is not based on that. Indeed, the label is *caused* by this cue. We need to rely on causal, robust cues because otherwise, we will not generalize well to new domains. As an example, consider a car classification task. Then Z could be the car body region of the image, which is a robust cue.
2. **Non-causal, spurious cue.** Such cues are hurtful for generalization. The label is not causally related to this cue, but they are *highly correlated* in the dataset. In the car classification example, a highway in the background would be a spurious cue.

When using vanilla training, nothing stops the model from using only non-causal, spurious cues, e.g., for recognition. The model can achieve high training accuracy (and even high in-distribution test accuracy) if the spurious cues are highly correlated with the label. Whenever the model faces an OOD dataset, however, it can perform arbitrarily poorly based on how predictive the learned bias cue is in the new setting.

Shifted Focus in ML: The "How" parts of problems in Computer Vision

In Computer Vision (CV), we might often be interested in whether an ML system is robust to perturbations. Examples include Gaussian noise, motion blur, zoom blur, brightness, and contrast changes. However, there are even more creative perturbations. For example, we

might measure whether the ML system can still classify objects accurately in quite improbable positions.

Spurious cues that are highly correlated with the task cue but are otherwise semantically irrelevant can greatly harm a model's performance when not acted against. We often want to test whether our classifier exploits spurious cues. This can lead to it breaking down on OOD samples. For example, we can observe the behavior of the classifier in cases where the background is changed, the foreground object is deleted/changed, or the backgrounds and foregrounds are mixed across categories. If our model uses the image background as a spurious cue to make its predictions, it will showcase poor performance in these tests.

Shifted Focus in ML: The "How" parts of problems in Natural Language Processing

We would like to briefly mention Chain-of-Thought (CoT) Prompting. An example is given in Figure 1.3. If we want to teach our Natural Language Processing (NLP) model a new task, we can provide it with some examples of the task and the correct answer and then ask a follow-up question. We supply no explanation of the answer in this case. What happens often is that the LLMs give incorrect answers to the next question. However, when prompting the model with exemplary detailed explanations of each correct answer, called CoT Prompting, the model also explains its prediction and even gets the answer right. It learns to rely on the right cues to provide the answer (and also provides an explanation).
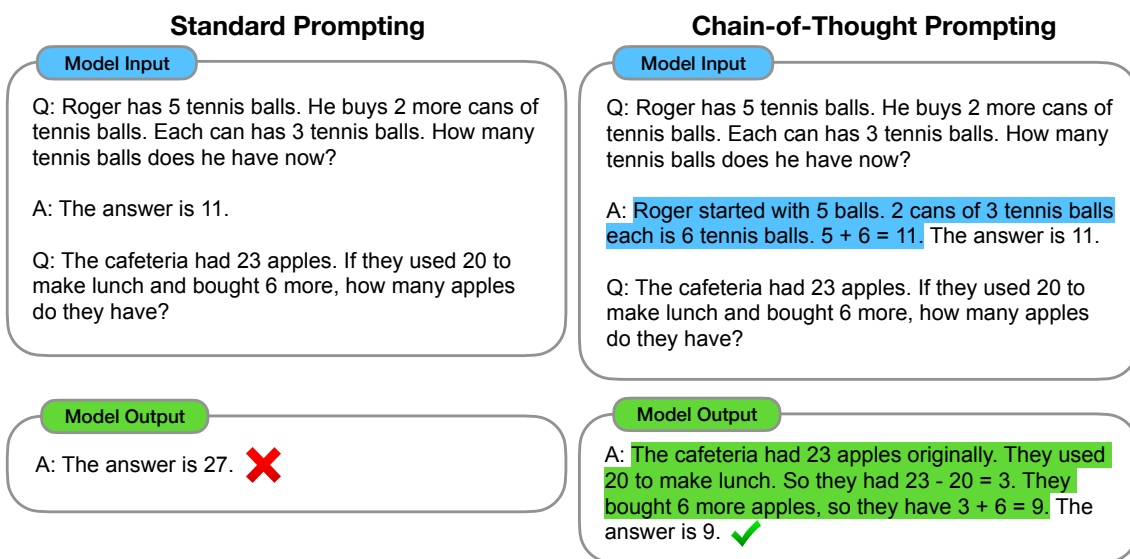
Figure 1.3: CoT Prompting can lead to better model answers. Figure taken from [203].

1.4.2 Machine Learning 2.0

We distinguish two ML paradigms regarding what question they seek answers for: ML 1.0 and 2.0.

Definition 1.4: Machine Learning 1.0

In ML 1.0, we learn the distribution $P(X, Y)$ (or derivative distributions, such as $P(Y \mid X)$), either implicitly or explicitly, from (X, Y) ("What") data. ML 1.0 only considers the "What" task: It does not include the used cues, explanations, or reasoning, i.e., the "How" aspect Z.

Definition 1.5: Machine Learning 2.0

In ML 2.0, we learn the distribution $P(X, Z, Y)$ (or derivative distributions), either implicitly or explicitly, from (X, Y) ("What") data:

$$\text{Input } X \xrightarrow{\hspace{1.5cm}} \underset{\text{The "How" aspect } Z}{\text{Selection of cue, exact mechanism, reasoning.}} \xrightarrow{\hspace{1.5cm}} \text{Output } Y$$
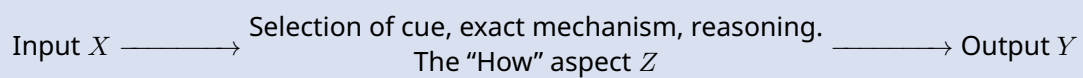
The motivation of ML 2.0 is clear: we want to use the same kind of data to get more knowledge. However, the Z-problem is *not guaranteed to be solvable* from (X, Y) data. Learning $P(X, Z, Y)$ contains all kinds of derivative tasks (a new set of tasks compared to what we had in ML 1.0): Now, we are trying to learn some distribution of X, Z, and Y. For example, we may wish to be able to predict the Ground Truth (GT) Z from input X correctly (learn $P(Z \mid X)$), i.e., to make sure that given an input, the model is choosing the right cue for input X.

In the following chapters, we aim to introduce the reader to various scalable trustworthy ML solutions with a focus on both theory and applications.

Chapter 2
OOD Generalization

2 OOD Generalization

2.1 Introduction to OOD Generalization

OOD generalization stands as a pivotal challenge in modern ML research. It seeks to construct robust models that perform accurately even on data not represented in the training set. This branch of research not only elevates the trustworthiness and reliability of ML systems but also broadens their applicability in real-world scenarios.

Before we get our hands dirty, we have to discuss some terms that are often used in OOD generalization. Let us start with the most basic one: the *task* we want to solve.
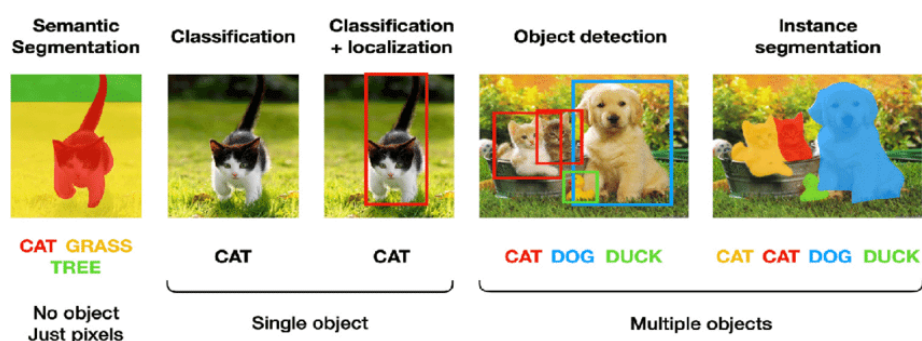
Figure 2.1: Illustrations of various computer vision tasks, taken from [91]. The field of computer vision is vast.

> **Definition 2.1: Task**
>
> Task refers to the ground truth (GT), possibly non-deterministic (see aleatoric uncertainty in Section 4.2) function that maps from the input space $\mathcal{X}$ to the output space $\mathcal{Y}$ that a model is learning, or is a description thereof. Equivalently, the task is the GT distribution $P(Y \mid X = x)$ we wish to model.
>
> **Alternative definition**: Task is the factor of variation (cue) that matters for us, i.e., the factor we want to recognize at deployment. Tasks are not inherent to the data; they are always defined by humans. This slightly differs from the previous definition, but both explain the same concept.

2.1.1   Examples of Tasks

Definition 2.2: ImageNet

ImageNet [44] is a large-scale, diverse dataset initially created for object recognition re-
search. Nowadays, it is popular to use ImageNet for classification, omitting the predic-
tion of a bounding box. It contains millions of annotated images collected from the web
and spans thousands of object categories that are organized according to the WordNet
hierarchy for nouns. The dataset contains hundreds to thousands of samples per node
in the hierarchy.

Ambiguity with "the" ImageNet Dataset: The term "ImageNet dataset" has been used
to refer to mainly two variants of the dataset which has caused a great deal of confusion:
  - **Full ImageNet Dataset/ImageNet-21K/ImageNet-22K**: The full ImageNet
    dataset contains 14,197,122 images associated with 21,841 WordNet cate-
    gories [206]. However, not all of these images are used in typical computer vision
    benchmarks. ImageNet-21K is equivalent to ImageNet-22K, the difference is that
    some researchers round up the number of classes to 22,000 in the name.
  - **ImageNet Large Scale Visual Recognition Challenge (ILSVRC)
    Dataset/ImageNet-1K/ILSVRC2017**: This is the most widely used subset of
    the ImageNet-21K dataset, involving 1,000 object categories. It contains 1,281,167
    training data points, 50,000 validation samples, and 100,000 test images. [206].
    However, the labels for the test set are not released. Therefore, one can only
    use the validation performance for evaluation when writing a paper, making the
    evaluation process less trustworthy. The annual ILSVRC competition, especially
    the 2012 challenge, which was won by the deep learning model AlexNet [111],
    played a pivotal role in the rise of deep learning.

Even within "classification", there exist various tasks: different sets of classes correspond to
different tasks.

  - The Pascal VOC datasets [50] consider 20 classes. These are datasets for object detec-
    tion, instance segmentation, semantic segmentation, action classification, and image
    classification.
  - The COCO datasets [123] contain 80 object categories and 91 stuff categories. Object
    categories strictly contain the Pascal VOC classes. These are datasets for object de-
    tection, instance segmentation, panoptic segmentation, semantic segmentation, and
    captioning. Crowd labels are added when there are too many (more than ten) instances
    of a class in an image. These aggregate multiple objects.[1]

Examples of Tasks in Computer Vision (CV)

An overview of CV tasks is given in Figure 2.1.

  - *Semantic segmentation* aims to predict a semantic label for each pixel in an image.
  - *Classification* is the problem of categorizing a single object in the image.

---

[1]COCO is collected from Flickr. ImageNet is partly also from Flickr and other databases.

- *Classification + localization* aims classify *and* localize a single object in the image.
- *Object detection* classifies and localizes *all* objects in the image. Now we have no restrictions on the number of objects the image might contain.
- *Instance segmentation* assigns a semantic label and an instance label to every pixel in the image. The instance label differentiates between unique instances with the same semantic label.

Examples of Tasks in NLP

> **Definition 2.3: Semantic Analysis**
>
> Semantic analysis in natural language processing (NLP) analyzes the conceptual meaning of morphemes, words, phrases, sentences, grammar, and vocabulary.

> **Definition 2.4: Pragmatic Analysis**
>
> Pragmatic analysis in NLP analyzes semantic meaning but also analyzes context. Instead of examining what an expression means, it studies what the speaker means in a specific context.

Analysis tasks aim to uncover syntactic, semantic, and pragmatic relationships between words/phrases/sentences in a document.

- Tokenization is an essential syntactic analysis technique.
- The semantic analysis of a document might involve sentence classification (like sentiment analysis) or named-entity recognition.
- Word sense disambiguation is a particular example of pragmatic analysis. It aims to unfold which sense of a word is meant in a certain context.
- Part of speech tagging is can be both deemed a semantic and a pragmatic analysis technique. It marks up words in a document with the corresponding part of speech (e.g., noun or verb).

Generation tasks involve generating text.

- Machine translation is an example of conditional text generation where a translation in language B is generated given the original document in language A.
- Question answering is also a conditional text generation problem where the model generates a coherent answer given a natural language question.
- Language modeling is the task of predicting the next word/character in a document or, equivalently, the task of assigning a probability to any text. Here, we condition on the partial sequence we have generated so far.

2.1.2   Generalization Types

Now, we are ready to consider a general overview of generalization types. First, let us introduce some terms that will play a crucial role in our discussion of OOD generalization.

Definition 2.5: Environment (Domain)

The environment is the distribution from which our data are sampled.

Definition 2.6: Cue (Feature, Attribute)

Cues, features, and attributes all refer to the factors of variation in the data sample. Examples include color, shape, and size.

Note: A cue is not necessarily a feature in a vector representation. Cues are also entirely independent of the model. They are characteristics of the dataset.

Definition 2.7: In-Distribution (ID) and Out-of-Distribution (OOD) Samples

In-distribution (ID) samples come from a test dataset which is used to gauge the model's performance on familiar data (in-distribution generalization). Out-of-distribution (OOD) samples, on the other hand, are drawn from a different test dataset to assess the model's performance on unfamiliar or unexpected data (out-of-distribution generalization).

Definition 2.8: Generalization Types

Generalization type		How is training $\approx$ test?	How is training $\neq$ test?
ID		Training and test sets come from the same distribution.	We have different samples.
OOD	Cross-Domain	Training and test sets are for the same task.	They are from different domains.
	Cross-Bias		They have different cue-correlations.
	Adversarial		Test samples are worst-case scenarios.

Note: This is not a comprehensive list of OOD generalization variants.

Let us give examples for each scenario and consider some remarks.

Example of ID Generalization

We consider the task of recognizing a set of people from an office. They might be in different poses or situations, but always the same people, both in dev and deployment. The office theme will be common in the subsequent examples for different generalization types to highlight and emphasize the main differences between these.

Example of Cross-Domain OOD Generalization

Here, we might consider the task of recognizing person A from the office, but for the first time in a party costume during deployment. We have the same (or even different) people from dev in new, unseen clothes. One of the features is changing from training to test, meaning the training and test sets are from different domains. This generalization scenario mixes many factors; we will focus on cross-bias generalization more.

Example of Cross-Bias OOD Generalization

Persons A and B work in the office of the previous examples. We want to recognize person A for the first time in person B's jacket. We have the same people but in exchanged clothes. The biased cue for person A has changed from their jacket to person B's jacket. More formally, the cue that was highly correlated with person B in the training set now co-occurs with person A in the test dataset. The ML system will likely predict person B if we do not counteract the bias. This is because of the well-known shortcut bias of ML systems, which we will discuss later.

In practice, we are usually interested in changing the cue from training to test the model is likely to look at when making a prediction (because of shortcut bias), e.g., clothing. Such benchmarks test whether the model is focusing on a cue that is irrelevant to the task (e.g., a person's clothing is irrelevant to their identity).

Example of Adversarial OOD Generalization

Consider the problem of recognizing person A even when they hide their identity with a face mask (with someone else's face on it or using other tricks). Now person A is the adversary against our face recognition system, but this does not necessarily mean that person A has malicious intentions. Person A might wish to hide their true identity by making the model fail to recognize his face. There are also adversarial patterns to avoid facial recognition systems, e.g., to avoid surveillance. Adversarial generalization is a tough task, and it is even more challenging to obtain guarantees for this generalization type.

2.2   Why do we even care about OOD generalization?

In the YouTube video "Self Driving Collision (Analysis)" [33], we see perfect weather and visibility, with low traffic. Nevertheless, as the Tesla turns onto the road, it does not detect a row of plastic bollards and hits them. This accident is not a one-off occurrence, as later in the video, it tries to hit other bollards too. Why does this happen? Because this is a new street arrangement that the model has not seen before, and it fails to generalize to this situation. To be sure that the model is robust in many situations, we need some kind of OOD generalization.

Many things constantly change in the world. New, unseen events happen all the time, like the Covid pandemic. If we trained a model before the pandemic to predict loungewear sales for a particular date, it might have extrapolated well until national lockdowns were announced. These lockdowns caused a substantial domain shift, in which loungewear sales increased

considerably. The model we trained before the lockdowns failed to reflect reality after this environmental change.

The typical solution to domain shifts is model retraining. Things inevitably change over time, and the model accuracy drop over time is unavoidable if the model is kept fixed. People thus often recollect data, annotate new samples, and retrain the model on new data. We can use this procedure to keep the model's accuracy above a certain threshold, illustrated in Figure 2.2.
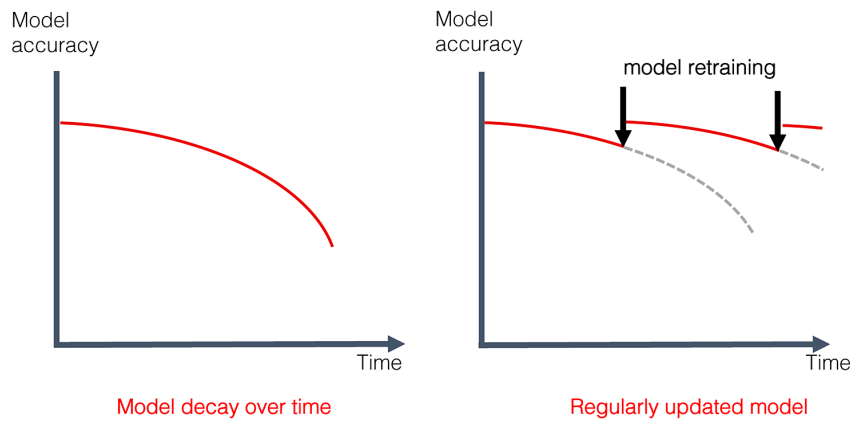
Figure 2.2: Illustration of the use of regular model updates to preserve deployment accuracy, taken from [4]. In many cases, model accuracy would plummet over time if we did not update it regularly.

Definition 2.9: Model Selection

Model selection is the process of selecting the best model after the individual models are evaluated based on the required criteria. One usually has a pool of models specialized for various domains. The expert chooses the best model for the current deployment scenario. For example, Amazon often performs model selection in its cloud services.

Definition 2.10: MLOps

MLOps is an engineering discipline that aims to unify ML systems development and deployment to standardize and streamline the *continuous delivery* of high-performing models in production [198]. An overview is given in Figure 2.3.

However, the constant retraining of models and the model selection expertise (MLOps) is costly.

- **Manpower**: 100k EUR/person/year at least.
- **GPUs, electricity**: 25k EUR/year + 8000 kg CO_2/year (considering a single NVIDIA Tesla A100 unit and Google Cloud).
- **Data management** (schema, maintenance) is also expensive.
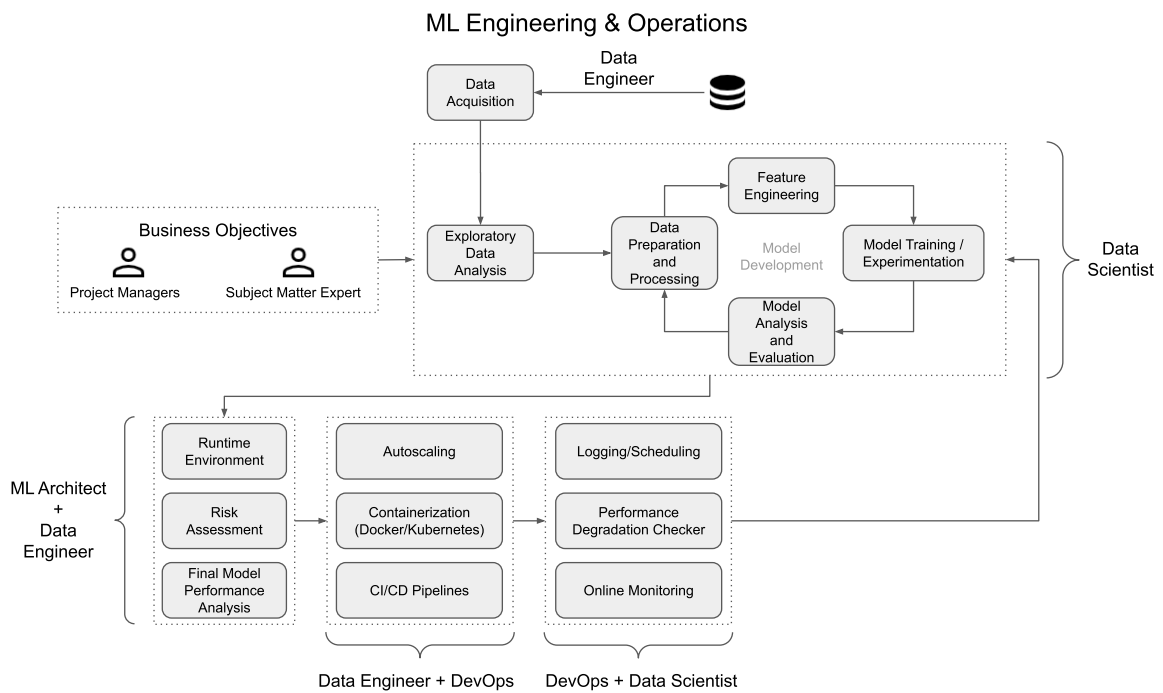
Figure 2.3: MLOps is a complex discipline with multiple participants. **Note**: Data Acquisition is not just a DB query. It also includes the collection of data. The data curation procedure can take a long time. One must keep track of shifting data (data versions), keep annotators in the loop, and update models accordingly. This procedure can be very costly. Figure adapted from [198].

Additional Information: NVIDIA Tesla A100

The NVIDIA Tesla A100 is a tensor core GPU often used for training ML models. It can be partitioned into 7 GPU instances so multiple networks can efficiently operate simultaneously (training or inference) on a single A100. In early 2023, it has one of the world's fastest memory bandwidths, with over 2 TB/s. Training BERT is possible *under a minute* using a cluster of 2048 A100 GPUs [149].

Definition 2.11: DevOps

A set of practices intended to reduce the time between committing a change to a system and the change being placed into production while ensuring high quality. [85]

ML problems arise from business goals. If there is no distribution shift and no need for model selection, there is no need for MLOps, and we only need DevOps. We need MLOps (continuous updates of models) because the data, user, and environment shift continuously. Ideally, we only have to perform continuous updates semi-automatically: We only need a few people to maintain the system. Eventually, however, we wish to get over MLOps as well. We need models that are very robust to domain shifts to achieve this.

2.2.1   Greater Levels of Automation

First, we define *diagonal datasets* that will help us understand the levels of automation in ML and the ill-defined behavior in OOD generalization (Section 2.7.1).

> **Definition 2.12: Diagonal Dataset**
>
> A dataset in which all (or multiple in general) cues vary at the same time (i.e., they are perfectly correlated) that can be used to achieve 100% accuracy. Thus, it is impossible to infer what the deployment task is from the label variation. A model using either of the perfectly correlated cues could achieve 100% training accuracy.

Next, we need to describe the Amazon Mechanical Turk service to reason about annotation costs and crowdsourcing.

> **Definition 2.13: Amazon Mechanical Turk**
>
> The Amazon Mechanical Turk (AMT) is an online labor market for dataset annotation, where one can crowdsource their annotation task.

We consider five levels of automation (1: lowest, 5: highest) in problem-solving.

Level 1: No ML

In this case, we have no ML model to use for our particular problem. The human effort is gigantic: A center with hundreds of personnel is constantly required (which was a common case 40-50 years ago). They take care of input streams on the fly, i.e., they are processing a continuous data stream with *human intelligence*. This procedure is *very costly* and *inefficient*.

Level 2: MLOps with Periodic Annotation

In this setup, we have an ML model available to help with our problem. However, this model can only generalize to the same distribution based on the annotated samples. The human effort is reduced but still considerable: A group of people annotates thousands (possibly millions across projects) of samples every month, as the world is changing quickly. Options for annotations include in-house annotators, outsourcing to annotation companies, or crowdsourcing through AMT. Annotation costs 10-30 USD per hour per person on AMT. (Slightly above minimum wage for US workers.) Harder tasks, e.g., instance segmentation, cost more. For the browser-based annotation of 1 million images, we estimate up to 1 million USD for AMT crowdsourcing. An ML engineer's market price is 100-300k USD per year per person. These costs are prohibitively expensive for small businesses.

Level 3: MLOps with Reduced Annotation

Now, we have an ML model that is minimally resilient to distribution shifts. The human effort is reduced even more: Annotation is required only every year. This resilience reduces the cost of MLOps quite a bit.

Level 4: MLOps with No Annotation

In this hypothetical scenario, our ML model – once trained – is so robust against distribution shifts that it only requires minimal human engineering (e.g., hyperparameter adaptation and model selection). Regarding the human effort, annotation is not needed anymore. Only ML engineers are needed to select the right model suitable for the task at that particular time (based on the needs of business executives). They are also constantly looking for the best models.

Level 5: ML without MLOps

Here, even the ML engineer functionality is (partly) automatized. The model can alter its hyperparameters to adapt to changing distributions. Adapting hyperparameters usually requires fine-tuning; however, the way we choose hyperparameters can be made very efficient, e.g., requiring only very few observations of training sessions and data samples. (In ML, we always need observations.) It can even be automatized with, e.g., Bayesian optimization. Importantly, this does not refer to a meta-model that can automatically choose between the set of candidate models. We cannot even be sure that is possible, as certain factors cannot be inferred from the data. As an example, let us consider a diagonal dataset in which the shape and color cues co-occur perfectly. At one point, users might want a shape-based classifier. Later they might change their mind and want a color-based classifier. This requirement is not reflected in the data stream for a diagonal dataset: it is part of the human specification. This is precisely why model selection always involves human feedback.

Why is an expert still needed for model selection? One might wonder why an expert decision-maker cannot be replaced in this very idealistic hypothetical scenario. This is because some metrics are often unreliable (that look good on paper, but the model that performs well on them might not be what we want), and there are requirements from a model that are often hard to quantify. An ML engineer might also be needed to keep the pool of models up to date, including the latest innovations in ML. There are also always new model architectures and general technologies that appear. This pool needs to be constantly curated and updated to the general needs of the users. These new models might also not be better than previous ones on *all* criteria, just some of them (e.g., better accuracy at the cost of less interpretability).

There might also be many criteria to adhere to. For example, we might be interested in the performance, computational resources, fairness, calibration, or explainability. Accuracy is not the only criterion, and there is no *single* criterion. The single *best* model does not exist in general, no matter how robust our pool of models is; and even if our pool of models is robust, some models might perform (slightly) better in exact deployment scenarios on certain metrics – we want to squeeze out performance. Model selection is not just an argmax. With multiple criteria, it is often too difficult to put some weights on these metrics and use thresholding. Automating model selection is, therefore, a challenging problem with fundamental limitations.

Finally, an expert is always needed to *give the final word*. They must make an executive decision and choose the best model based on the business needs. When there are problems with a new model (e.g., fairness), a human must intervene and roll it back to a previous state. **Note**: The expert discussed here does not have to be an ML expert. The main decisions usu-

ally come from business executives.

2.2.2 Once we "solve" OOD generalization...

What happens if we "solve" OOD generalization (i.e., our models become robust to distribution shifts)?

- Our model will work well even under new situations.
- MLOps will not be needed at the current scale. (However, model selection and ML expertise will probably be needed for a long time.)
- Small businesses will be able to adopt ML more easily.
- ML can be extended to more risky applications because we can be sure that it will work in novel situations, too.
- ML will drive the risky applications, e.g., the industry of healthcare, finance, or transportation. Robust models gain trust. However, we will see later that *explainability* is just as important.

To summarize our introduction to OOD generalization and drive the key points across:

- ML is still costly because it requires periodic annotation and maintenance. There are huge human costs involved.
- When ML models generalize well to novel situations, costs will be reduced.

2.3 Formal Setup of OOD Generalization

2.3.1 Stages of ML Systems

To discuss a more formal setup of OOD generalization, let us first consider two stages of ML systems: *development* and *deployment*.

Definition 2.14: Development (dev)

Development is the stage where we train our model and make design choices (for hyperparameters) within some resource constraints.

Definition 2.15: Deployment

Deployment is the stage where our final model is facing the real-world environment. This environment is called a *deployment environment* and can change over time.

Definition 2.16: Training

Training is the particular action of fitting the model's parameters within the dev stage to the training set, with a fixed hyperparameter setting.

> We do not separate the training phase from the rest of the dev phase, but we *do* separate dev from deployment.

Definition 2.17: Testing

Testing is a lab setup designed to mimic the deployment scenario closely – scientists evaluate their final inventions on test benchmarks and report their results.

Practice point of view:
- This is different from deployment and still a part of development.
  - If we want to be precise: As soon as we have labeled samples from deployment (and we make any design choices based on these or just test our model), we are using information from the deployment setup in dev. We cannot talk about true (domain or task) generalization anymore. The deployment scenario should stay fictitious and unobserved in such settings.

Academia point of view:
- The test set (2.3.2) and the action of testing is treated as a part of the deployment.

The specification of these stages can be bundled into one *setting*.

Definition 2.18: Setting/Setup

A setting specifies the available resources (during development) and an ML system's surrounding (deployment) environment.

Essential components of a setting:
- **Development resources**: What types of datasets, samples, labels, supervisions, guidance, explanation, tools, knowledge, or inductive bias are available?
  - ML engineers are also resources. They have their own knowledge to optimize an ML model the right way. If we have better engineers with better intuition of what to do in a scenario, we can train the model quicker and better.
- **Deployment environment**: What kind of distribution will our ML model be deployed on?
- **Time**: Resource availability changes over time. The deployment environment changes over time. We can only deploy after development, but sometimes we keep developing after deployment.

Example of a Setting

Consider an ID supervised learning setup. This is an ideal scenario ML research has started its exploration in. Various strong results about consistency, convergence rates, and error bounds can be given in this setup [88, 87] that break in OOD settings.

Our *development resources* are labeled (X, Y) samples from distribution P. Our *deployment environment* contains unlabeled samples X from distribution P presented one by one. **Note**:

This is an incomplete description of the development resources and the deployment environment that aims to drive the main points across. In scientific papers, a much more thorough description is required.

We usually specify settings when we have an actual task we want to work on, i.e., we have a *real-world scenario* at hand.

> **Definition 2.19: Real-World Scenario**
>
> A real-world scenario is a projection of a setting onto a hypothetical or actual convincing real-world example. This is a particular situation that fits the setting.

Example of a Real-World Scenario

Consider an ID supervised learning setup again (the simplest case). Our *task* is to build a system for detecting defects (e.g., dents) in wafers (semiconductors, pieces of silicon used to create integrated circuits) through image analysis. Our *development resources* contain a dataset of wafer images with corresponding labels – defective or normal. In our *deployment environment*, the images are of the same distribution, as the wafer products and camera sensors are identical between the dev dataset and the data stream from deployment.

> **Additional Information: How to compare methods with different resources?**
>
> We always want to compare methods fairly. If one method uses fewer resources in development, we cannot compare the two methods fairly.

2.3.2 Dataset Splits in ML

Next, we discuss different general dataset splits used in ML.

> **Definition 2.20: Training Set**
>
> The training set is a (usually large) collection of samples whose purpose is to train the model.
>
> **What is optimized?** Model parameters.
>
> **What is the objective?** The training loss, possibly with regularization.
>
> **What is the optimization algorithm?** A gradient descent variant using Tensor Processing Units (TPUs), or GPUs.
>
> **How frequent are updates of the model?** $\mathcal{O}$(milliseconds-seconds).

Definition 2.21: Validation/Dev Set

The purpose of the validation set is to roughly simulate the deployment scenario by using samples the model has not seen yet and measure ID generalization.

What is optimized? Model hyperparameters and design choices.

What is the objective? Generalization metrics.
- If we consider true OOD generalization without having access to the target domain (i.e., not domain adaptation (2.4.4) or test-time training (2.4.6)), we cannot measure OOD generalizability on the validation set. Therefore, the validation set usually comes from the same domain(s) as the training set. Otherwise, we would already tune our hyperparameters on the domain we wish to generalize to; thus, whether we measure OOD generalization on the test set later is questionable. Such scenarios exhibit 'leakage', which we will cover in Section 2.5.2.

What is the optimization algorithm? For example, Bayesian optimization, "Grad student descent", random search.

How frequent are updates of the model? $\mathcal{O}$(minutes-days).

Definition 2.22: Test Set

The test set is used to simulate the deployment scenario more accurately than during validation by using samples from the distribution we believe the model will face during deployment. The test set can, therefore, also measure OOD generalization.

What is optimized? The methodology and overall approach through the shift of the field.
- For example, the shift from CNNs towards ViTs.
- The line is unclear between the change of methodology and design choices; this is more like a spectrum.

What is the objective? Generalization metrics.
- The test set can be any type of OOD dataset.

What is the optimization algorithm? Paradigm shifts, updating the evaluation or the evaluation standard.
- As the field changes, the set test also changes. For example, for ImageNet, many test sets are available (e.g., for generalization to different OOD scenarios), and more have been added over time.
- We are setting a new goal for the field that many researchers will follow.
- Standard refers to the benchmark, metric, or protocol according to which we evaluate our models. (It has a close connection to the test sets we use.)

How frequent are updates of the model? $\mathcal{O}$(months-years).
- In the scale of months and years, methods are *meant to be optimized* to the test set. The problems this optimization entails are crucial to understand and are discussed in detail in Section 2.3.3.
- The test set must be updated to the user and societal needs over time. Naturally,

27

> the training set and validation set also change over time.

2.3.3   Why Idealists Cannot Evaluate on the Test Set

We measure accuracy on the test set because we wish to *compare* our method to previous methods. This is an implicit way of choosing a model over other methods, which is part of the methodology. Therefore, the test set is still a part of development in practice in the most precise sense.[2]

Whenever we make any decisions based on test results (be it ours or others'), we cannot measure generalizability on the test set anymore. This is almost always violated in practice. However, there is no clear workaround, as benchmarks are essential to progress in any field of ML research. We can only "spoil the test set less," but we can never *not* spoil it if we want to advance the field.

2.4   Common Settings for OOD Generalization

There is no such thing as *the* OOD generalization setting. There are many different scenarios for it. Let us first explain why differentiating between these learning settings is important.

2.4.1   Why are the learning settings important?

Figure 2.4: Collage of different domain labels and corresponding images, taken from [224]. Images of the same kind of objects can be surprisingly different when considering different domains.

Let us first define the notion of *domain labels*.

---

[2] For domain generalization (Section 2.4.5), we never get any annotations from deployment in reality. We consider the deployment scenario as a fictitious entity.

> **Definition 2.23: Domain Label**
>
> The domain label is an indicator of the source distribution of each data point in the form of a categorical variable (e.g., dataset name).

Example: "MNIST" can be a domain label for an image from the MNIST dataset [117]. Samples in different datasets are (almost always) coming from different distributions. Other valid domain labels include "Art Painting", "Cartoon", "Sketch", or "Real World", as illustrated in Figure 2.4.

Distinguishing various learning settings is of crucial importance for the following reasons.

1. To figure out which techniques can be used for the given learning scenario. We want to understand the given ingredients precisely and know the relevant search keywords for googling the papers.
2. To compare against previous methods in the same learning setting. It is key to enumerate the exact (and sometimes hidden) ingredients used by a method and compare it only with methods that use the same ingredients. Some authors may give misleading information about the setting their method operates in. For example, if one claims to have not used domain labels but has used some equivalent form of them, we must be able to notice that and voice our concerns. Comparing methods based on their ingredients is much more justified than comparing based on the name of the settings the authors *claim* to adhere to.

2.4.2 ID generalization

For the sake of comparison, let us start with ID generalization. An illustration of this setting is depicted in Figure 2.5. We have the same domain all the way through development and deployment. During deployment, we get unlabeled samples to which we wish our model to generalize.
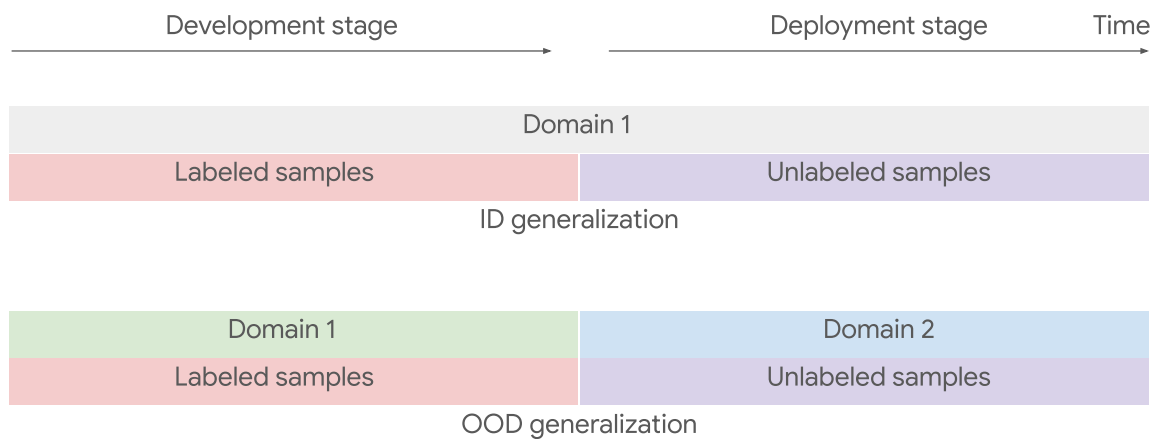

Figure 2.5: Illustration of the ID generalization setting (top) and the general OOD generalization setting (bottom). OOD generalization showcases a change of domain.

2.4.3   Domain-Dependent OOD Generalization

A general view of this setting is shown in Figure 2.5. There are different domains for development and deployment. One needs to generalize to the deployment domain. This is the most general setting for OOD generalization. There are many names of settings. Exact definitions of settings for OOD generalization fluctuate. Therefore, understanding the exact ingredients for each setting matters much more than "which name to put". For the purpose of the book, we will still go over the settings and try to put definite boundaries. We discuss different categorizations of OOD generalization settings below.

Categorizing according to the nature of the difference between dev and deployment

- In *cross-domain* generalization, the deployment environment contains completely unseen cues in dev.
- In *cross-bias* generalization, deployment contains unseen compositions of seen cues in dev.
- *Adversarial* generalization considers a (real/hypothetical) adversary in deployment who tries to choose the worst-case domain.

Categorizing according to the extra information provided to address the ill-posedness

- In *domain generalization*, domain labels are provided.
- In *domain adaptation*, some (un-)labeled target domain samples are available in dev.
- *Test-time training*'s dev continues even after deployment. We get access to deployment (target domain) samples. We may or may not label them.
- *Domain-incremental continual learning* considers a single domain during dev. Domains are added over time during deployment.

These settings all come with different ingredients, and one should not compare methods across different settings. The two axes of variation above are independent.

2.4.4   Domain Adaptation

Development stage	Deployment stage	Time

Domain 1	Domain 2
Labeled samples	Labeled/Unlabeled samples

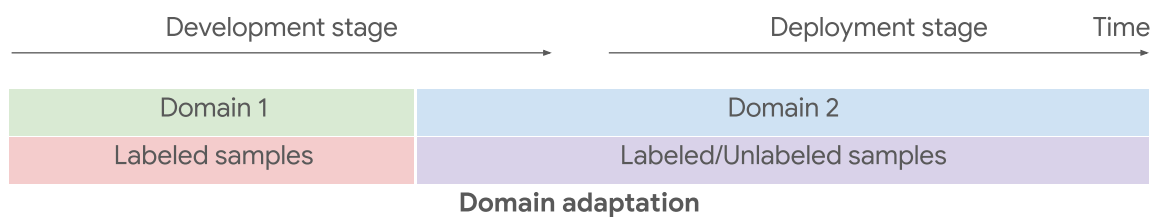**Domain adaptation**

Figure 2.6: Domain adaptation setting. The development stage also comprises samples from domain 2.

Domain adaptation is illustrated in Figure 2.6. During the dev stage, we have access to some labeled or unlabeled samples (depending on the exact situation) from the deployment environment. We can, e.g., align our features with the target domain statistics using moment matching.
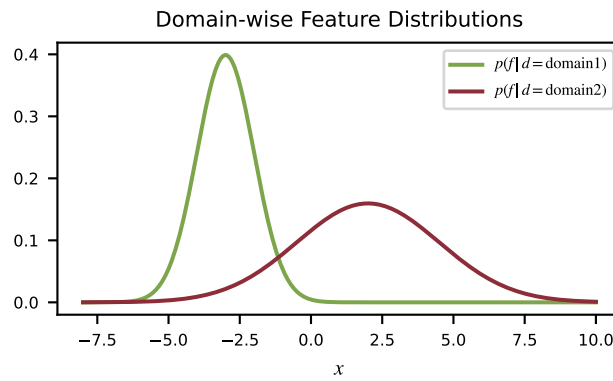
Figure 2.7: Feature embedding distributions can be notably different between the domains available in the development stage of domain adaptation. The figure shows Gaussians fit to domain-wise empirical feature distributions of samples from different domains for feature alignment in domain adaptation using moment matching: we aim to match these Gaussians among domains. f: feature, d: domain.

Moment Matching in Domain Adaptation

An example of domain-wise feature distributions is illustrated in Figure 2.7. These can, e.g., correspond to the penultimate layer's sample-wise activations in a ResNet-50. We represent the empirical distribution of the domain-wise feature values by their expectation and variance (approximated by the sample mean and variance). For domain 2, we have a few unlabeled samples that can be used for this computation. We place, e.g., an L_2 penalty on the differences between domain 1 and 2 statistics (sample mean and variance in our example) of intermediate features, or we can also consider the Wasserstein distance between the Gaussians as a penalty. During training, for samples from domain 1, we compute the task loss (e.g., cross-entropy) and the penalty term. For samples from domain 2, we only compute the penalty, as there are no labels for these samples. We only backpropagate gradients of the penalty through the labeled domain 1 samples and use the unlabeled samples only to calculate the penalty. This way, we only directly train on domain 1 but adapt our model based on domain 2 samples to generalize to domain 2. This tends to give us a small amount of improvement in robustness.

2.4.5 Domain Generalization
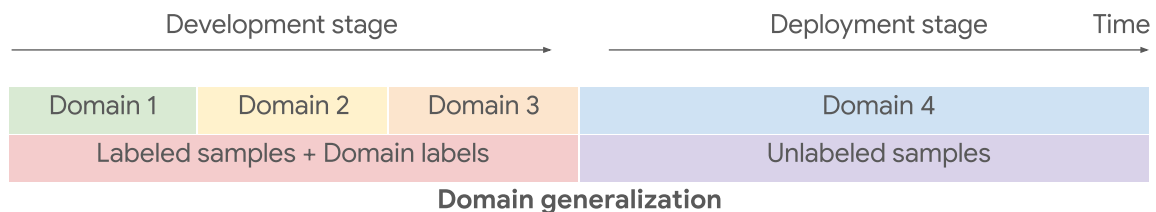
Figure 2.8: Domain generalization setting. We have access to multiple domains during the development stage, but we have to generalize to a novel, unseen domain in the deployment stage.

An overview of domain generalization is given in Figure 2.8. During the dev stage, we have access to labeled samples from multiple domains. We also know the domain label for every sample. Knowing domain labels is usually a hidden assumption; not many papers talk about

this. If we do not know the domain labels, there are techniques for detecting the domains without them, but these are never perfect and come with additional assumptions.

Moment Matching in Domain Generalization

We can "unlearn" domain-related characteristics in our representation by performing *moment matching* similarly to domain adaptation, but now between all domains available in the development stage. Similarly to the domain adaptation case, we compute the sample mean and variance separately for each domain as we have domain labels.[3] We fit Gaussians to the features of samples from different domains. We align the Gaussians for different domains by, e.g., placing an L_2 penalty on the pairwise differences between their corresponding means and covariance matrices. We backpropagate gradients through all samples, using both the task labels and domain labels. If we succeed, we ignore differences among domains in the training set based on moments. We hope that the model becomes independent of domain information (of any kind), so it will probably work well on the next (unknown) domain.

2.4.6 Test-Time Training

Figure 2.9: Test-time training setting. The development stage continues in the deployment stage.

Test-time training is shown in Figure 2.9. After training our model, we keep updating it according to the labeled or unlabeled samples (depending on the exact setting) from the deployment environment. Thus, dev continues even into the deployment because our model keeps being updated. We might not do labeling in domain 2, but it helps to have access to incoming domain 2 samples and correct the feature model on the fly (e.g., by performing moment matching).

This paradigm is becoming more popular these days. A key figure of the approach is Alexei Efros.

2.4.7 Domain-Incremental Continual Learning

An overview of the domain-incremental continual learning setting is given in Figure 2.10. We train on a single domain before deployment. Domains are added over time during deployment. We label a few samples over time and update our model on the way. Only the labeled samples are used for improving our model. We hope that the model also does not forget previous domains. Performance should remain as high as possible for previous domains as

---

[3]The task labels are not used for moment matching, only to compute the task loss.
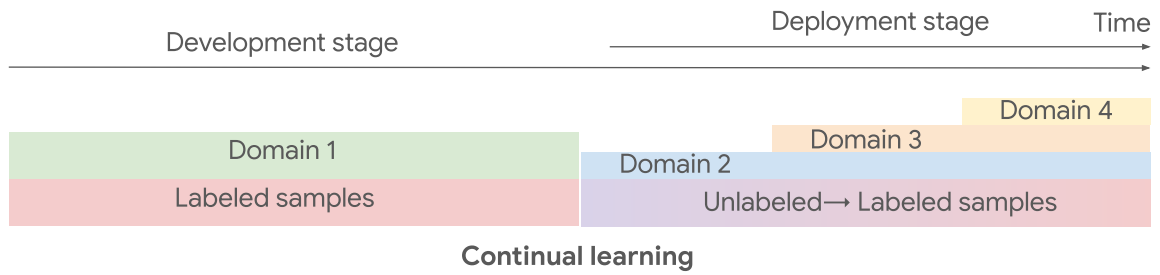
Figure 2.10: Domain-incremental continual learning. New domains are added over time in the deployment stage.

well. Data keeps coming from all deployment domains, but we must adapt quickly to the new domain.

2.4.8   Task-Dependent OOD Generalization

In general, the task being different is a lot harder than the domain being different. Usually, a different task also means a different domain.[4]
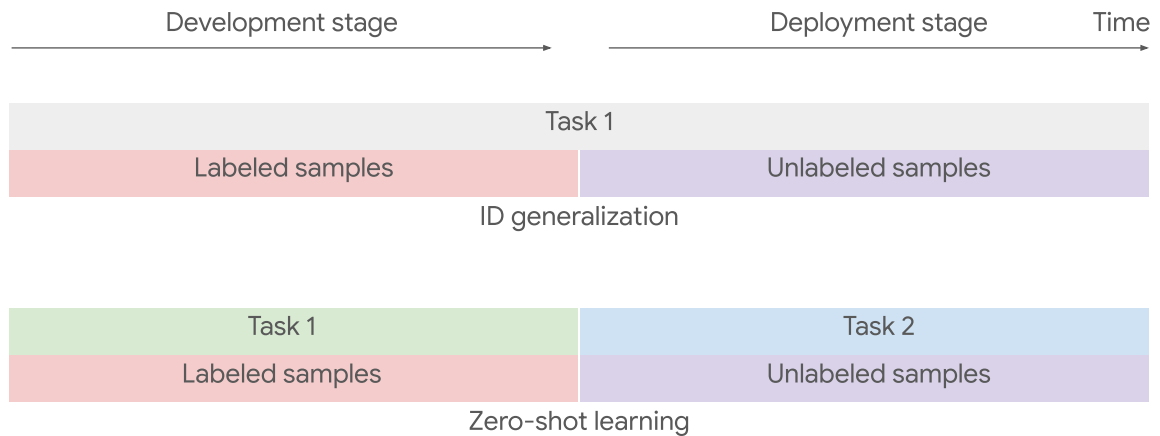
Figure 2.11: Comparison of ID generalization and zero-shot learning. Zero-shot learning aims to tackle a novel, unseen task in the deployment stage.

So far, the task stayed the same across development and deployment. However, the task can also change over time. The best-known scenario of this is *zero-shot learning*, which is compared to ID generalization in Figure 2.11. In ID generalization, the task stays the same. In zero-shot learning, we have a different task for deployment about which we have no information in dev.

Large Language Models and Zero-Shot Learning

Large Language Models (LLMs) are capable of performing zero-shot learning [106] (called zero-shot-CoT prompting). They can encode the task description in natural language, so there is sufficient information for the model to solve the problem in principle. It is, however, still fascinating how LLMs can figure out how to solve new kinds of tasks not presented before that are an output of human creativity.

---

[4]However, we can also come up with counterexamples. When task 1 is to predict numbers 0-4 on MNIST and task 2 is to predict numbers 5-9, the domain stays the same, but the task changes.

Nevertheless, we almost never have any guarantees about benchmarks truly being zero shot for LLMs – their datasets are *huge*, and we can never be certain that the model did not have the same task in its training dataset. CLIP also has zero-shot learning capabilities.

Categorizing according to which tasks are available at the development and deployment stages

- In *ID generalization*, the task stays the same in dev and deployment.
- In *zero-shot learning*, during deployment, we are faced with a new task not present in dev.
- K-*shot learning* gives a "softened" setting where we have K labeled samples of the deployment task in dev.
- *Meta-learning* has different tasks available during dev. This can also be combined with K-shot learning.

2.4.9 K-Shot Learning
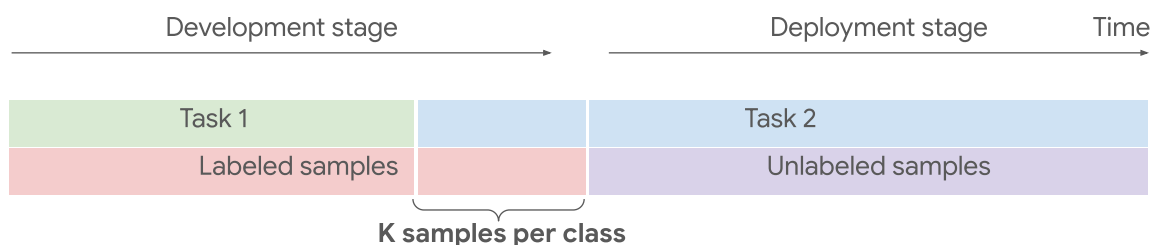
Figure 2.12: K-shot (few-shot) learning setting. K samples are available from task 2 during the development stage to aid the model towards robust generalization.

K-shot learning is illustrated in Figure 2.12. People try to make zero-shot learning easier by introducing some labeled samples for the target task during development. The K samples per class are for the target task. We learn to fit our model to the deployment task using a large number of task 1 samples and a few ($K \times$ #class) task 2 samples.

Example 1: ImageNet pretraining followed by fine-tuning on a downstream task.

Example 2: Linear probing in Self-Supervised Learning (SSL). Here, we do not even need labeled samples for domain 1. We train a strong feature representation in a self-supervised fashion, then we apply a linear classifier on the learned features and fine-tune the model with labeled task 2 data.

2.4.10 Meta-Learning + K-Shot Learning

Meta-learning can be combined with K-shot learning, as shown in Figure 2.13. In this case, we have multiple tasks during development, and we wish to learn features that generalize across tasks, but we still need samples from the target task. In essence, we "learn to learn a new task" with tasks 1-3 (that give rise to a compound task). We then adapt our model to the deployment task 4, using the K samples per class for task 4.[5]

---

[5]For task changes, we also need to change the output head in the parametric case (e.g., linear probing). It is not needed for the non-parametric case (kNN) and CLIP [159]. In CLIP, we need no information about the exact target task (zero-shot learning), but we need an LLM. The information comes from large-scale pretraining.

Figure 2.13: Meta-learning + K-shot learning setting. Multiple (proxy) tasks are available in the development stage. We further have access to K samples per class from the deployment stage task.
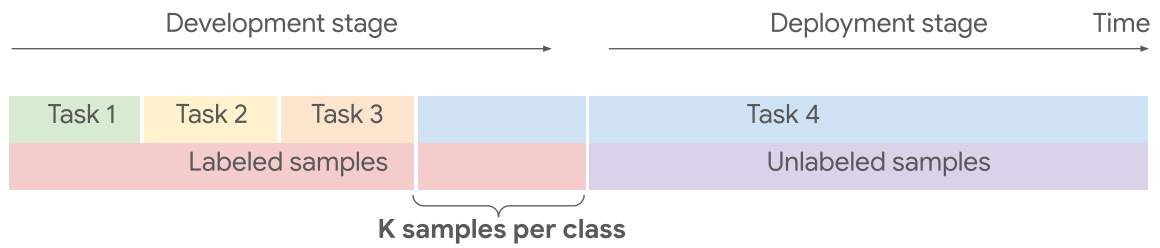
2.4.11 Task-Incremental Continual Learning
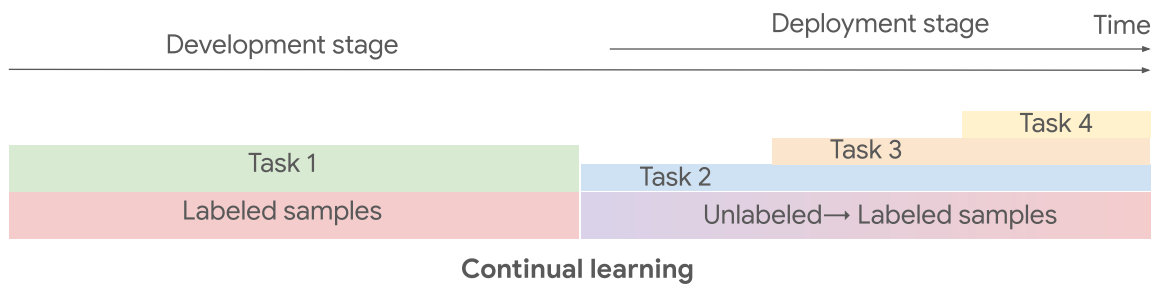
Figure 2.14: Task-incremental continual learning setting. The development stage continues in deployment, adding new tasks over time.

Here, we consider a task-incremental version of continual learning, which is illustrated in Figure 2.14. Tasks are added over time. We label only a few samples over time. (We can only utilize these.)[6] We update our model on the way. Ideally, the model should not forget the previous task.

2.5 ML Dev as a Closed System of Information

To better illustrate the flow of information in the ML development stage, we draw a parallel between it and a closed thermodynamic system. This is illustrated in Figure 2.15.

Here, dev is represented as a closed system that consists of four main parts: *dataset*, *annotation*, *inductive bias* and *knowledge*. Inductive bias can appear in the form of the model architecture or the way we pre-process the data. A good example of knowledge is the expertise of people with much experience in training neural networks (NNs).

In a closed lab environment where no dataset, annotation, or anything else is given to the system, there should be no additional information that suddenly appears. We should not expect new information to be born out of this system. Equivalently: There is no change in the maximal generalization performance we can get out of this system. There are lots of papers that *violate* this principle (see 2.5.2) [67, 83].

Note that it is possible to *kill* information by, e.g., averaging things or replacing measurements with summary statistics.

---

[6]The output layer is always switched for the task accordingly. Sometimes very shallow output heads are enough (e.g., linear probing) if we have a strong backbone feature representation.
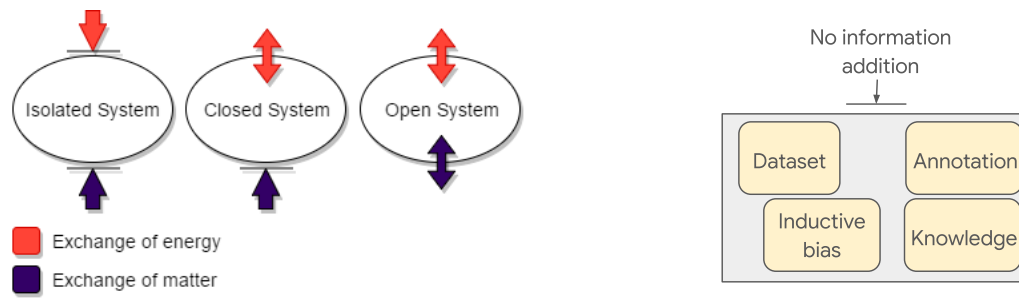
There is no addition of matter to a closed thermodynamic system.

There is no addition of information to an ML development system.

Figure 2.15: Comparison between a closed thermodynamic system [208] (left) and an ML development system (right).

2.5.1 Information Leakage from Deployment

> **Definition 2.24: Information leakage**
>
> *Information leakage* refers to the situation where information intended exclusively for the deployment stage becomes accessible during the development stage. It is an influx of information into a closed system.

Let us consider information leakage in the domain generalization setting, illustrated in Figure 2.16. When one defines domain generalization as in Section 2.4.5, information about domain 4 must not be available during development. That is, we cannot inject new information into this closed ML system that comprises the resources *at dev stage*. If we *do* inject new information, we have to treat it as a new setting: When information about domain 4 is available, we cannot call it a domain generalization setup anymore. This also means we cannot compare against previous domain generalization methods. We need to set up a new setting, build a new benchmark, and compare against methods with the same setting.[7]

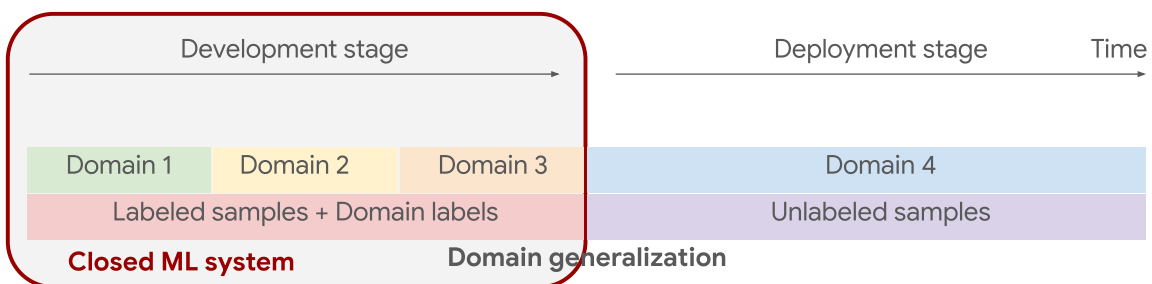Figure 2.16: Closed ML system in the development stage for the domain generalization setting. Information about domain 4 must not be available in this closed system.

---

[7]Reducing the amount of information gained from evaluation helps in not spoiling the test set too much. For example, we might use a hidden server for benchmarking where only the ranking of submissions is shown but not the exact results.

Examples of Information Leakage in the Domain Generalization Setting

Consider the domain generalization setting from 2.4.5. There are several ways how information leakage can surface and spoil our results.

Scenario 1. Some hyperparameters are chosen based on labeled samples from domain 4. In a sense, our dev set is partly taken from domain 4. We cannot talk about true generalization.

Scenario 2. Some hyperparameters are chosen by visually inspecting domain 4. This is still information leakage, just in a less automated way.

Scenario 3. The model is trained on labeled samples from domains 1-3 and unlabeled samples from domain 4. In the particular case of only two domains – domain 1 in the development stage and domain 2 in the deployment stage – and labeled or unlabeled samples being used from domain 2, we are performing domain adaptation, not domain generalization.

Scenario 4. Some hyperparameters are chosen to maximize publicly available scores after evaluation on some benchmarks with domain 4. Strictly speaking, these scores contain information about domain 4. One way to overcome this information leakage is to provide only a ranking of methods but not the scores.

Information Leakage from Domain Generalization Evaluation

Let us consider the particular problem of evaluation domain generalization methods in a bit more detail. It makes perfect sense to have a few labeled samples from deployment because we have to evaluate our system on a new domain anyway. Still, strictly speaking, as soon as we evaluate our model on the new domain, we *use* our target domain (domain 4), so we cannot talk about generalization. We may need to shift the definition of domain generalization into something that allows some validation in the target domain (like domain adaptation). **Evaluating or benchmarking domain generalization is, therefore, contradictory.** Researchers still evaluate their methods on domain generalization benchmarks (observing the test set corresponding to the deployment domain multiple times through their lifetime), as we need to monitor progress somehow.

Information Leakage from Pretraining

Another interesting question arises if we consider pretrained models. As soon as there is a pretrained system in dev resources, we introduce not just a single model but the entire pretraining dataset (which is gigantic in the case of large language models (LLMs) or CLIP [159]). Much expertise is put into the dev scenario, which is also a dev resource. The consequence is that it is hard to say that anything is zero-shot learning in such settings. For example, if we give a new task to an LLM, we can never be sure that the model has never seen that task during training. The use of ImageNet-1K pretraining for zero-shot learning is also criticized: the 1k classes contain much information, and for certain classes we evaluate our model on, we do not have true zero-shot learning at all.

When is information leakage not a problem?

The other side of the argument about the severity of information leakage is that it might not always matter whether something is zero-shot learning. If an LLM contains all information about the world, then there is nothing new in the deployment stage, so we cannot have true zero-shot learning. Nevertheless, the model works very well, so we can make good use of it in many real-world scenarios. Take the example of face recognition. Suppose there is a person our system has never seen before. We want it to recognize (or verify) that this is the same person on subsequent days. This setup is zero-shot verification. However, as soon as our training set contains billions of identities, this does not matter anymore. In the most extreme case of having seen all people in the world, we do not have to generalize to unseen people. Nevertheless, we are still happy with the system if it works well for everyone. Zero-shot learning thus becomes less meaningful at a large scale.

2.5.2   A Case Study on Information Leakage

We consider an example of a paper that is leaking information from deployment, titled "Self-Challenging Improves Cross-Domain Generalization" [83]. It is straightforward to find such papers, even from highly regarded research groups.

> **Definition 2.25: Ablation Study**
>
> We are changing one factor at a time in our method, as we want to see the contribution of each factor towards the final performance. Everything else is kept fixed. Then we can understand the effect of the factor better, and we can also optimize that factor (hyperparameter) separately.

Table 2.1: Benchmark results of various feature drop strategies. Explanation of columns: e.g., for the art painting column, we train the model on {cartoon, sketch, photo} and test it on art painting. Table taken from [83].

Feature Drop Strategies	backbone	artpaint	cartoon	sketch	photo	Avg ↑
Baseline	ResNet18	78.96	73.93	70.59	**96.28**	79.94
Random	ResNet18	79.32	75.27	74.06	95.54	81.05
Top-Activation	ResNet18	80.31	76.05	76.13	**95.72**	82.03
Top-Gradient	ResNet18	**81.23**	**77.23**	**77.56**	95.61	**82.91**

In Tables 1-5 of [83], an ablation study is conducted. We show Table 1 of the paper in Table 2.1 for convenience. We see various hyperparameters chosen based on the performance on the domain they want to generalize to. (For example, the "Feature Drop Strategy" hyperparameter considers different ways to drop features to make the model better regularized.) They are looking at the generalization performance to each of the domains using leave-one-out domain generalization. They finally choose the hyperparameters based on the average accuracy on the left-out domains. If we also validate on the test set, we cannot talk about domain generalization anymore, as we have information leakage. (Even if we consider the academic point of view of the test set belonging to deployment.)

This hyperparameter configuration will be pretty good for the PACS dataset [122] (see below). However, this does not guarantee that this is the best ingredient for non-PACS cases. We might overfit to PACS severely by making such choices. (Of course, this overfitting can also happen even if we do not use the test set as a part of the validation set, but rather as a criterion for method selection across different papers. However, that is a much less severe case of overfitting. Here, the authors make use of the test set many times in a *single* paper.)

Takeaway: Ablation studies are generally great for ID generalization tasks, but one should be very careful with ablation studies for OOD generalization.

2.5.3   Solutions to Information Leakage

There are many (partial) solutions to combat information leakage, discussed below.

Select hyperparameters within dev resources. Section 3 of "In Search of Lost Domain Generalization" [67] discusses information leakage and provides possible solutions for it. Selecting hyperparameters, design choices, checkpoints, and other parts of the system must be a part of the learning problem (i.e., part of the ML dev system). When we propose a new domain generalization algorithm, we must specify a method for selecting the hyperparameters rather than relying on an unclear methodology that invites potential information leakage.

Use the test set once per project. By using a specific test set multiple times, we can always overfit to it. If our goal is to go towards a distribution outside dev, then evaluating on the test set multiple times can be harmful. However, as discussed previously, we *do* have to use it multiple times to compare methods and evaluate our approach. Solution: At least do not use the test set for hyperparameter tuning; tune them on the validation set. For example, if we want to generalize well to the art painting dataset of PACS, tune the hyperparameters on {cartoon, sketch, photo}. Then we measure performance on the art painting test set. We will do the same thing for a new, genuinely unknown domain in deployment: find the hyperparameters on the known domains. Thus, one should use the test set sparingly. A good rule of thumb might be to use it once per paper. This way, we are less likely to overfit to it. (The State-of-the-Art (SotA) architectures are also likely to overfit to standard benchmarks, e.g., to ImageNet-1K. [163])

Update benchmarks. Even if the test set labels are unknown, we can overfit to the test set just based on the reported performance, e.g., on leaderboards. Thus, for many reasons, the test set has to be changed every once in a while. Another helpful idea is to use a non-fixed benchmark, where the data stream changes over time. For comparability, this is an issue: we have a continuously changing target over time. However, it is usually not problematic: In human studies, researchers have been dealing with a changing evaluation set. By using statistical tests, they could always argue about statistical significance. One particular example is the case of *clinical trials*. It is physically impossible to test two related drugs on exactly the same set of people: The test set changes from experiment to experiment. However, statistical tests give a principled way to determine if the observed changes are significant or if they could have happened by chance.

Modify evaluation methods. A different approach is to use a differential-privacy-based evaluation method. This method adds a Laplace noise to the accuracy before reporting it to the practitioner. This is better than overfitting to a single benchmark.

Summary: We are trying to address an impossible problem: to truly generalize to new domains, which requires them to be previously unseen. We can never keep them completely unseen, as then we cannot *measure* generalizability. However, as soon as we measure generalizability, we cannot talk about generalization anymore. This is an unavoidable dilemma for many ML fields, even more so for Trustworthy Machine Learning (TML), as they deal with more challenging cases of generalization where evaluation is very tricky.

2.6 Domain Generalization Benchmarks
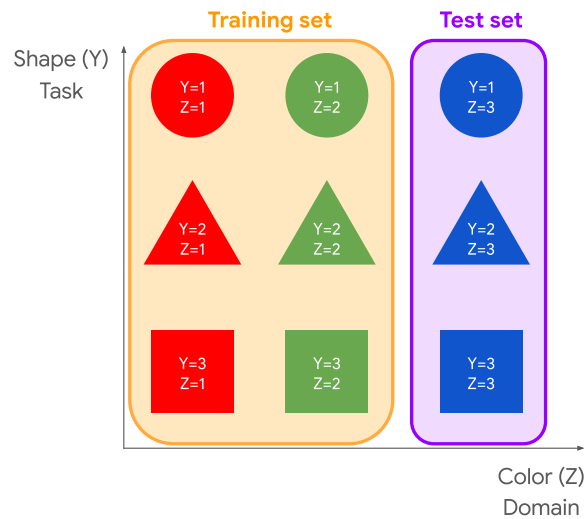
Figure 2.17: Training (yellow) and test (blue) datasets in the domain generalization setting. Shape is the task, color is the domain.

Definition 2.26: Subpopulation Shift Benchmarks

In subpopulation shift benchmarks, we consider test distributions that are subpopulations of the training distribution and seek to perform well even on the *worst-case* subpopulation.

In the previous section, we highlighted the importance of paying special attention to benchmarks in the domain generalization setting. Now, let us take a closer look and discuss some prominent examples in more detail. Some of these benchmarks will be related to the problem of subpopulation shift which is partially connected to domain generalization.

2.6.1 Examples of Domain Generalization Benchmarks

A toy domain generalization problem is shown in Figure 2.17. Our *goal* is to generalize well to blue images: This is a particular instance of cross-domain generalization. The *inputs* are images with mono-colored objects of some shape. The *labels* are $\{0, 1, 2\}$ – we have a three-way classification problem. The set of classes is shared across the domains and is assigned according to the object's *shape* (circle, triangle, or square). The model's *task* is to predict the label of a given input. Here, we consider three domains: red, green, and blue colored

objects. This is not a real problem but we refer to it to illustrate the scheme shared among the subsequent benchmarks.
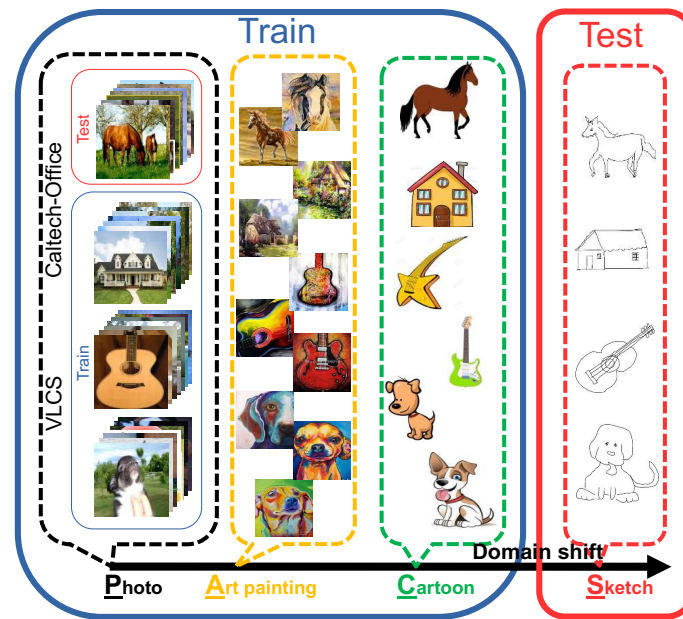
Figure 2.18: The PACS dataset can be used for domain generalization. Figure taken from [121].

The PACS Dataset

The PACS dataset considers four domains: **P**hotos, **A**rt Paintings, **C**artoons, and **S**ketches. Samples from each domain are shown in Figure 2.18. The set of classes is shared across the domains. To benchmark domain generalization, leave-one-out evaluation is used. For example, one might train on **PAC** and test on **S**.

DomainBed

DomainBed is a combination of some popular domain generalization benchmarks into a single suite. It subsumes, e.g., PACS, Colored MNIST, Rotated MNIST, and Office-Home. Office-Home contains four domains. These are (1) *art* that contains artistic images in the form of sketches, paintings, ornamentation, and other styles; (2) *clipart* that is a collection of clipart images; (3) *product* that contains images of objects without a background; and (4) *real-world* that collects images of objects captured with a regular camera. For each domain, the dataset contains images of 65 object categories found typically in Office and Home settings. Samples from each subsumed benchmark are shown in Figure 2.19.

The Wilds Benchmark

The Wilds benchmark [105] comprises several tasks and domains for each task. IT contains domain generalization benchmarks and also subpopulation shift benchmarks. A detailed illustration of the dataset is shown in Figure 2.20.
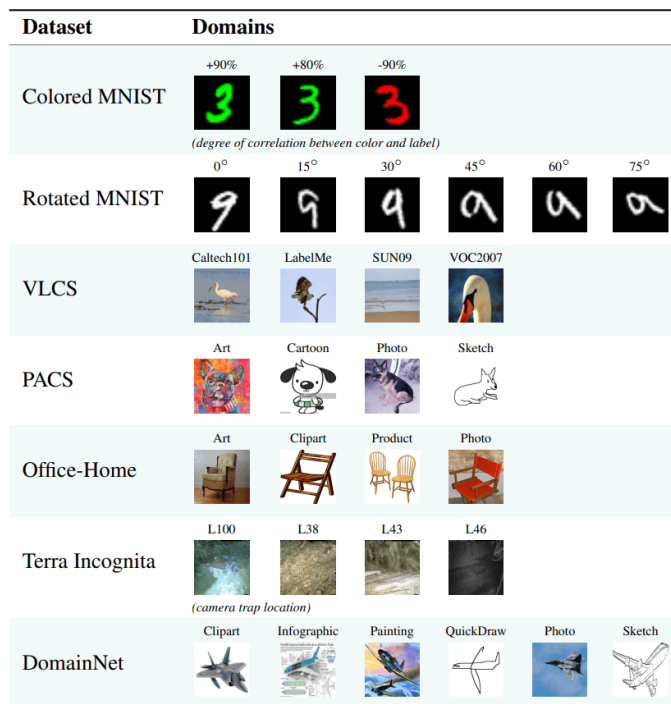
Figure 2.19: The DomainBed suite. Illustration taken from [66].

Figure 2.20: The Wilds suite. Figure taken from [170].

ImageNet-C

ImageNet-C [77] is an extension to the ImageNet dataset [44] with a focus on robustness. For the same image, the dataset contains various corruptions. Corruptions include Gaussian Noise, Defocus Blur, Frosted Glass Blur, Motion Blur, Zoom Blur, JPEG Encoding-Decoding, Brightness Change, and Contrast Change. Examples of these corruption types are shown in Figure 2.21. The ImageNet-C dataset consists of 75 corruptions, all applied to the ImageNet test set images. It simulates possible corruptions under the deployment scenario, thereby measuring the robustness of the model to the perturbation of the data generating process.
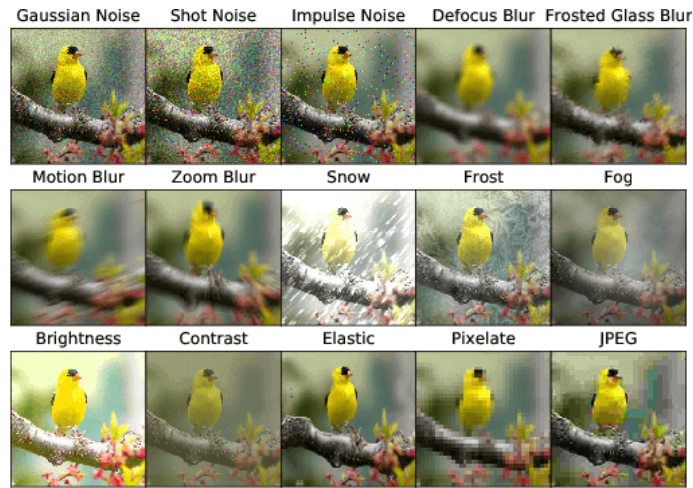
Figure 2.21: Illustration of the various corruptions ImageNet-C employs, taken from [76].

ImageNet-A

ImageNet-A [78] collects common failure cases of the PyTorch ResNet-50 [73] on ImageNet.[8] It contains images that classifiers should be able to predict correctly but cannot. Examples from ImageNet-A are shown in Figure 2.22.
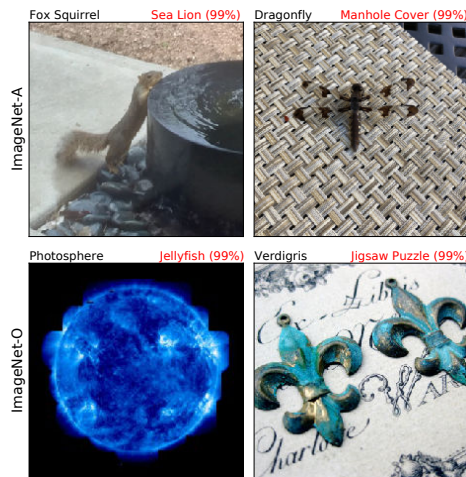
Figure 2.22: Sample from the ImageNet-A and ImageNet-O datasets, taken from [78].

ImageNet-O

ImageNet-O is another extension to ImageNet that contains anomalies of unforeseen classes which should result in low-confidence predictions, as the true class labels are not ImageNet-1K labels. ImageNet-O examples are shown in Figure 2.22.

---

[8]We explicitly mention the used implementation of ResNet-50 because there are subtle differences between versions.

2.7 Domain Generalization Difficulties

We have discussed how easy it is to confuse a setting with domain generalization just by not being careful enough with how one uses information about the target distribution. For those who are ready to accept this difficulty, we would like to point out that there are even more complications with domain generalization. However, we hope that these difficulties will not be an obstacle but rather an invitation to challenge, which is why we gathered the most important ones in this section.

2.7.1 Ill-Defined Behavior and Spurious Correlations
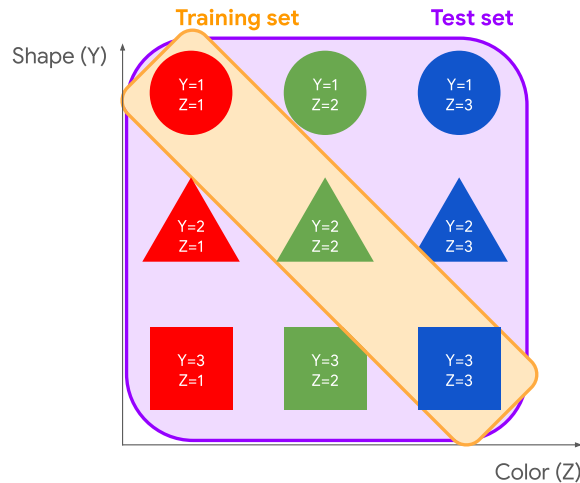
Figure 2.23: Diagonal training dataset and unbiased test set in cross-domain generalization. During training, the model is only exposed to samples where the shape and color labels coincide.

Consider Figure 2.8. For this setting, we outline two main difficulties: the ill-defined behavior on novel domains and the spurious correlations between task labels and domain labels.

Ill-Defined Behavior on Novel Domains

The model does not know what to do in regions without any training data. One could ask how domain generalization is even possible. It works in practice, but there are no rigorous theories as to why. We take it at face value, without any guarantees of the model's behavior on novel domains. This problem can be addressed through calibrated epistemic uncertainty estimation (4.2.3) to make the model "know when it does not know".

Spurious Correlations between Task Labels and Domain Labels

> **Definition 2.27: Spurious correlation**
>
> A spurious correlation is the co-occurrence of some cues, features, or labels, which happens in the development stage but not in the deployment stage.

For example, our prediction of shape may depend a lot on the color. If we have a diagonal dataset, this can become a huge problem, as depicted in Figure 2.23. Here, we have a perfect correlation between the two cues in the training dataset. In other words, there are spurious correlations between *task labels* and *domain labels*. This results in an ill-defined behavior on novel domains.

The problem of spurious correlations is also present in cross-bias generalization. We will consider this setting as it is easier than domain generalization, and ill-definedness is out of the picture.

2.8   Cross-Bias Generalization

We will now discuss cross-bias generalization from Table 2.1.2 that has a particular focus on the problem of spurious correlations. As seen before, we can amplify the spurious correlation between domain (bias) and target label (task) for OOD generalization to arrive at a scenario like Figure 2.23. We also remove the issue with unseen attributes: a model is guaranteed to encounter each attribute (e.g., possible shapes, colors) at least once, but in a heavily correlated fashion.
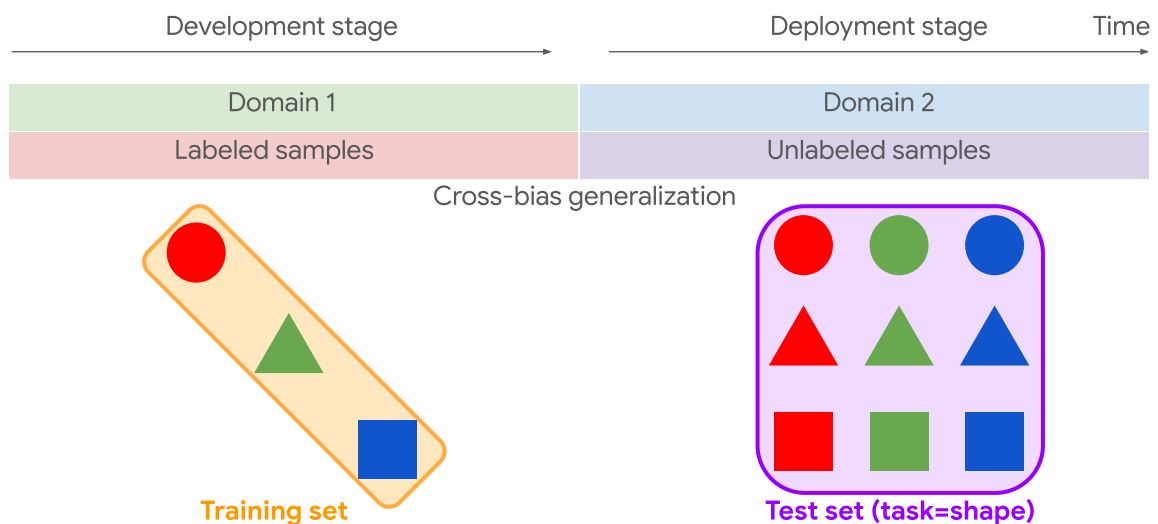
Figure 2.24: Cross-bias generalization setting with an unbiased deployment domain. In the deployment stage, the model has to do well on samples where the correlation between color and shape is broken.

This leads us to textbook cross-bias generalization, a cleaner setup for addressing the spurious correlations, for which an overview is given in Figure 2.24. In the test set, we have to recognize a diverse set of combinations of cues that we have not seen during training.

In general, the situation could be better described as "We still have more dominance along the diagonal, but we have a requirement that every single subgroup (i.e., (color, shape) combination) has to have a similar level of performance.". This formulation is roughly equivalent to having an equal number of samples in each grid cell. It is also possible that the deployment scenario is *still biased*, just has a different bias. We impose no restrictions on the deployment distribution.

> **Additional Information: Compositionality and Cross-Bias Generalization**
>
> Cross-bias generalization has close ties to compositionality [10, 113] that aims to disentangle semantically different parts of the input in the representation of neural networks. If a network leverages compositionality, i.e., treats semantically independent parts of the input independently when making a prediction, spurious correlations cannot arise by definition. This leads to robust cross-bias generalization. Of course, achieving this in practice is much more complicated.

2.8.1 Why is cross-bias generalization still challenging?

ID generalization is already an ill-posed problem. The No Free Lunch Theorem states that without extra inductive bias in the dev scenario, we cannot train a model that generalizes to the same distribution. We need inductive biases to find well-generalizing models ID. Without inductive biases, any model is equally likely to generalize well ID [212, 141].

OOD generalization (in particular, cross-bias generalization) poses another layer of difficulty: the *ambiguity of cues*, discussed next. We need further information in the ML dev system to solve it.

2.8.2 The Feature Selection Problem

We mentioned that the ambiguity of cues brings an additional challenge to cross-bias generalization. We would like to formally define this ambiguity.

> **Definition 2.28: Underspecification**
>
> An ML setting is underspecified when multiple features (e.g., color, shape, scale) let us achieve 100% accuracy on the training set. The training set does not specify what kind of cue the model should be looking at and how to generalize to new samples that do not have a perfect correlation. If the model chooses the incorrect cue, we say a *misspecification* happens.
>
> **Note**: We assume a network with very high capacity that can get 100% accuracy for every cue in the training set. For complex cues, the decision boundary tends to be wiggly, but under our assumption, even this decision boundary can be learned.

Underspecification in the cross-bias generalization setting necessitates the selection of the suitable feature(s) for good generalization to the deployment scenario.

A model under the vanilla OOD (e.g., cross-bias) generalization setting with a diagonal dataset lacks the information to generalize to an arbitrary deployment task well. When predicting in the deployment scenario (considering an uncorrelated dataset), the model cannot simultaneously use all perfectly aligned cues on the training set, as they contradict each other. Any cue the model adopts from training could be correct; the answer depends on the deployment task (chosen by a human, e.g., they can choose the most challenging cue for the model), which is arbitrary out of the perfectly correlated cues.

- If we have an adversarial deployment task selector, it can always fool the system into performing badly by choosing the most difficult cue for the model as the task.

Without any knowledge about the deployment task, cross-bias generalization is not solvable with a diagonal training set. Yet, it happens a lot that someone claims this in ML conference papers. They usually have a hidden ingredient that they implicitly assume. This is a prime example of *information leakage*.

To select the right feature for the task, more information is needed. This also holds for more general OOD settings: an example is shown in Figure 2.25.
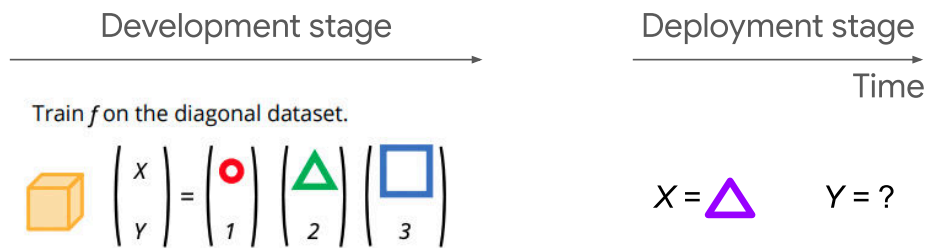
Figure 2.25: Underspecification in a more general toy OOD setting than cross-bias generalization. We are faced with the same problem: We now know that color is not the task, but shape and size can still be tasks. Figure inspired by [173].

The Feature Selection Problem in Fairness

The feature selection problem is also closely connected to the problem of fairness. What is fairness? From the viewpoint of the equality of opportunity as a notion of individual fairness: people who are similar w.r.t. a task should be *treated* similarly. There can be attributes for individuals that are relevant to the task and attributes that are supposed to be irrelevant, e.g., demographic details, such as race or gender. We want the model to only look at relevant features (task cue), not sensitive/prohibited attributes (bias cue). This notion of fairness is comparative: We are determining whether there are differences in how similar people (according to the task cues) are treated.

Decision-makers should automatically avoid differential treatment according to people's race, gender, or other possibly discriminatory factors if we accept in advance that none of these characteristics can be relevant to the task at hand.

2.8.3   Extra Information to Make Cross-Bias Generalization Possible

As we have seen, without extra information, cross-bias generalization is not solvable. We suggest considering a simplified generalization setting where such information is available in the development stage. This is much less exciting than true generalization, but we need this simplification to make the problem feasible. We will consider two ways to add extra information to the setting that makes the problem well-posed.

99% biased samples   +   1% unbiased samples

Development   **Deployment (task=shape)**
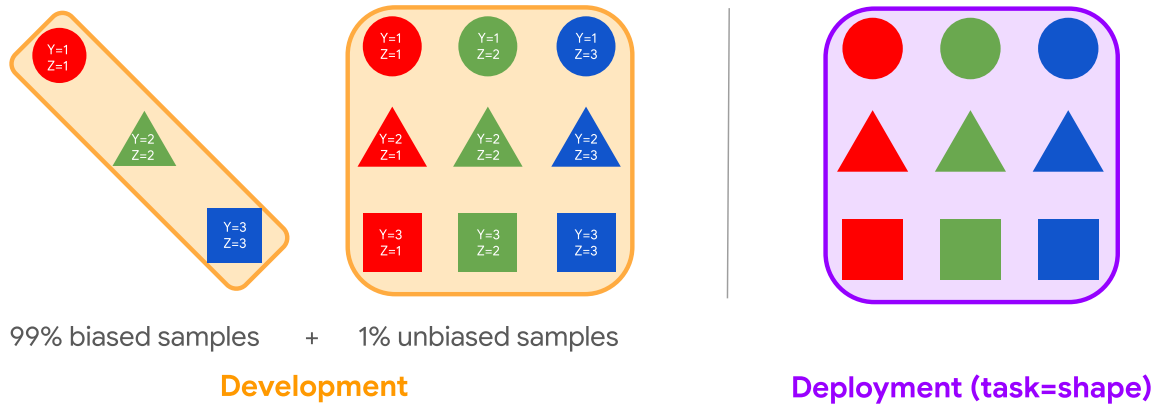
Figure 2.26: New setting that makes cross-bias generalization possible, referred to as the "First way" in the text. We have a few unbiased samples in the development resources and bias labels are also available.

First way to make cross-bias generalization feasible: adding unbiased samples

This approach is illustrated in Figure 2.26. A small number of non-correlated samples are added to the dev resources (these samples are not necessarily deployment samples). We have attribute (Zs – here bias, but it could also be domain) labels for each sample that specify which bias category a sample corresponds to. For example, Zs can correspond to different jackets. It is useful to explicitly tell the model what *not to* use as cues (see DANN in Section 2.12.2) in the form of bias labels.

People control the percentage of unbiased samples using $\rho \in [0, 1]$ in papers. We have to know what ρ they are using; it is a part of the setting. The lower the percentage of unbiased samples, the harder the task becomes. The task can be made arbitrarily hard, up to the point that it is impossible again ($\rho = 0$). The test set is unbiased in this example. However, in the deployment domain, we might just as well have biased samples that are biased in a different way than the dev samples.

A word about domain generalization

As we discussed in 2.4.5, *domain generalization* is supplying additional information by provid-ing domain labels. However, if we simply treat the bias labels (color) as our domain labels for domain generalization, we are sadly still not able to solve the problem: for such a diagonal dataset, the task labels are the same as the bias labels. Under this interpretation, domain generalization does not directly make the problem solvable. That is why we still need access to unbiased samples. In this case, we treat the domain labels as 'unbiased' and 'biased', and

the problem is solvable again. (This is very similar to the first way, only the interpretation is different.)

Second way to make cross-bias generalization feasible: converting the problem to domain adaptation/test-time training

We can also consider *domain adaptation*. Here, the source of extra information is access to the target distribution. This is different from before when we only had unbiased samples that did not necessarily come from the target domain. By performing domain adaptation, we make the target distribution more accessible to the dev stage. Here, we assume labeled samples from the target domain, and knowledge about the domain of each sample.

The same logic applies if we convert the problem to *test-time training*. The only difference is that in test-time training, the target distribution changes continuously during deployment, therefore, we constantly adapt our model to new situations.

2.8.4   How to determine what cue our model learns to recognize?

To understand how well we solved the problem of cross-bias generalization or to gain insights into the model's inner workings, it is often helpful to understand which cue our model uses for predictions. However, answering this question is not straightforward in general.

To diagnose our model, we require labels for different cues (e.g., labels Y and Z from Figure 2.26). In that case, after training on our close-to-diagonal dataset, we label unbiased (off-diagonal) samples from a test set according to different cues and calculate the model's accuracy w.r.t. each labeling scheme on this unbiased test set. The model should achieve high accuracy for the cue it learned on the training dataset and perform close to random guessing for all other cues.

2.9   Shortcut (Simplicity) Bias

We have seen that due to underspecification (Definition 2.28), models can learn different equally plausible cues. But do models prioritize learning one cue over others? It turns out the answer is yes, simpler cues are learned first. This property is usually called *shortcut/simplicity bias*, defined below.

> **Definition 2.29: Shortcut Bias/Simplicity Bias**
>
> The shortcut bias is the ML models' inborn preference for "simpler" cues (features) over "complex" ones.
> When there are multiple candidates of cues for the model to choose from for achieving 100% accuracy (i.e., the setting is underspecified), the model chooses the *easier* cue.

2.9.1 Examples of Shortcut Bias

Let us first define the *Kolmogorov complexity*, which is needed for the details of the first example.

> **Additional Information: Kolmogorov Complexity**
>
> The Kolmogorov complexity measures the complexity of strings (or objects in general) based on the minimal length among programs that generate that string.
>
> Kolmogorov Complexity of a cue $p_{Y|X}$ (KCC) [173]:
>
> $$K(p_{Y|X}) = \min_{f:\mathcal{L}(f;X,Y)<\delta} K(f) \qquad \delta > 0, f : X \to Y.$$
>
> Intuitively, $K(p_{Y|X})$ measures the *minimal* complexity of the function f required to memorize the labeling $p_{Y|X}$ on the training set (i.e., $\mathcal{L} < \delta$).

As a toy example, according to the paper "Which Shortcut Cues Will DNNs Choose? A Study from the Parameter-Space Perspective" [173], Color > Scale > Shape > Orientation in the order of models' preference, regardless of the network architecture and the training algorithm. Why could this be? The reason, according to the authors, is that color is a simpler cue than the others, as measured by the Kolmogorov complexity of the cues. The authors approximate $K(f)$ by the minimal number of parameters of model f to memorize the training set with labels w.r.t. the cue in question.

To better illustrate what simplicity bias is, we provide several examples below. An overview is shown in Table 2.2, which is further detailed in the individual sections.

Table 2.2: Overview of bias types and corresponding cues.

Problem	Task	Bias Cue	Task Cue
Context bias	Classify object	Background context	Foreground object(s)
Texture bias	Classify object	Texture of object	Shape of object
Not understanding sentence structure	Natural language inference	Set of words in a sentence, lexical overlap cue, subsequence cue, constituent cue	The entire sentence
Biased action recognition	Recognize action that human is performing	Scene, instrument, static frames	Human movement
Using single modality for multi-modal tasks	Visual question answering	Question only	Question and image
Use of sensitive attributes	Predict possibility of future defaults	Sensitive attributes (disability, gender, ethnicity, religion, etc.)	Size of the loans, history of repayment, income level, age, etc.

Context Bias

Consider the task of object classification. The task cues are the foreground objects, but a classifier focusing on the background context bias cues can achieve high accuracy when the background is highly correlated with the foreground. The examples, shown in Figure 2.27, are from [179].

Example 1: We have a classification problem where one of the classes is 'keyboard'. On nearly all images, keyboards are accompanied by monitors. The model might learn a shortcut bias for detecting monitors (detecting these might be easier than detecting keyboards): Then, the context (monitor pixels) will influence the keyboard score (logit) more than the actual keyboard presence. This process will not generalize to novel scenes where keyboards and monitors appear separately. If we remove the monitors from the image, the score for 'keyboard' will go down. If we remove the keyboard from the image, the score for 'keyboard' will stay quite high because the monitors are still present. Generally, co-occurring cues/features (diagonal samples) often lead to spurious correlations.

Example 2: The task is 'frisbee', and the bias is 'person'. It is easier to detect people because they are usually larger in images. The same phenomenon can be observed here as in **Example 1**.

Note: We humans also often look at the context to predict what is present in an image (or scene).

Figure 2.27: Context bias can arise in various settings. Figure taken from [179].

Texture Bias

Consider the task of object classification again. In this case, the task cue is the shape of the object and the bias cue is the texture of the object.

Example: Training a cat/dog classifier on a diagonal dataset, where the texture and shape are highly correlated. At test time, we want to predict cats when changing their texture (e.g., to greyscale, silhouette, edges, or to a marginally different texture). The accuracy of humans stays consistently high because we like to look at global shapes. Popular CNN models break down in such scenarios. However, when only the true texture of the original object (cat) is presented, models stay perfectly accurate while humans make more mistakes (90% accuracy). The example is inspired by [60].

Note: Networks are prone to be biased towards textures because it is much easier to learn. If the task is 'shape', such networks will generalize poorly to no/different textures.

NLP Models Not Understanding the Exact Structure of the Sentence

Our task of interest is natural language inference: Given premise and hypothesis, determine whether (1) the premise implies the hypothesis, (2) they contradict each other, or (3) they are neutral. The task cue is the whole sentence pair. However, the model might only use the set of words in the sentences, the lexical overlap cue, the subsequence cue, or the constituent cue. These are explained in the examples below, taken from [138]. We consider three bias cues and corresponding premise-implication pairs for each.

Example 1: Lexical overlap cue. Assumes that a premise entails all hypotheses constructed from words in the premise.

$$\text{The doctor was paid by the actor.} \implies \text{The doctor paid the actor.}$$

Example 2: Subsequence cue. Assumes that a premise entails all of its contiguous subsequences.

$$\text{The doctor near the actor danced.} \implies \text{The actor danced.}$$

Example 3: Constituent cue. Assumes that a premise entails all complete subtrees in its parse tree.

$$\text{If the artist slept, the actor ran.} \implies \text{The artist slept.}$$

These can all lead to wrong implications, as seen above.

Biased Action Recognition

The model's task is to recognize the action that a human is performing on a video. The task cue is the human movement, e.g., swinging, jumping, or sliding. The bias cues might be the scene, the instrument (on/with which the action is performed), or the static frames. The quiz below is taken from [38].

Quiz: Can the reader guess what action the blocked person is doing in the videos of Figure 2.28? Even from the scene alone, we as humans can have a good guess about what the person is likely doing. This tells us that humans also use many cues in the context to make

predictions. However, we also know that there are many other possibilities; we are just giving the most likely prediction. When we observe the actual task cue, we can make predictions based on that. Machines fail miserably because they *only rely on bias cues* from the dataset. We want ML models to be aware that they can be tricked in such cases; a notion of uncertainty and well-calibratedness is needed.

Figure 2.28: Example of four frames in videos where it is remarkably easy to predict a human's (very likely) action based on a single, static frame.

Using a Single Domain for Multi-modal Tasks: Visual Question Answering

The task is to answer a question in natural language using both the question and a visual aid (an image). The task cue is, therefore, both the question and the image. The bias cue is *only* the question. When one of the modalities is already sufficient for making good predictions on the training set, the model can choose to only look at that cue because of the simplicity bias. This generalizes poorly to situations where both modalities are needed. The example below is inspired by [28].

Example: The question is "What color are the bananas?". In the image, we see a couple of green bananas. When the model only relies on the question, it will probably get this question wrong. (Correct answer: green, not yellow.)

ML-based Credit Evaluation System using Sensitive Attributes

The model is tasked to predict the possibility of future defaults for each individual. (Will the person go bankrupt, or will they be able to repay the loan?) The task cue is the size of the loans, history of repayment, income level, age, and similar factors. The bias cues are sensitive attributes that are not allowed to be used for the prediction, such as disability, gender, ethnicity, or religion. When an ML system learns to use bias cues to predict credit risks (that might not be explicit features in a vector representation), the model is not fair. The ML system requires further guidance to not use sensitive cues.

2.9.2   Is the simplicity bias a bad thing?

Whether shortcut bias is a good or a bad thing depends on the task.

Simplicity Bias in ID Generalization

Simplicity bias is actually *praised* in ML in general, especially in ID generalization. There are reports saying

> "Deep learning generalizes because the parameter-function map $[\theta \mapsto f_\theta]$
> is biased towards simple functions." (Valle-Perez et al. 2019)

The parameter space is enormous. If there is no inductive bias (from the training algorithm or the architecture), we can find whatever solution in the parameter space, many of which do not generalize well. But because of the simplicity bias, we will find some simpler rules that are very likely to generalize well *to the same distribution*. (Here, we use the assumption that preference for simple cues usually leads to simple functions.)
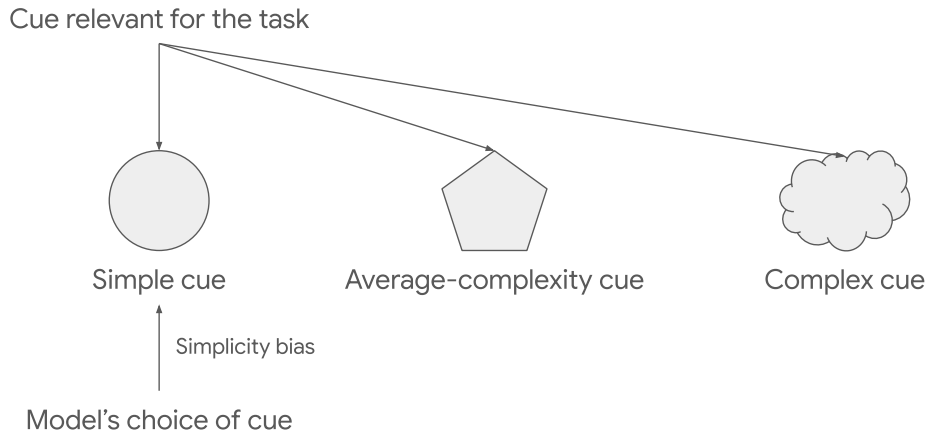
Figure 2.29: Example where the shortcut bias is favorable. For ID generalization tasks, simple cues are often sufficient for generalization.

The usefulness of the shortcut bias in ID generalization is illustrated in Figure 2.29. In the diagonal dataset case, any of the perfectly correlated cues are valid for performing well in deployment, considering ID generalization.

Simplicity Bias in OOD Generalization
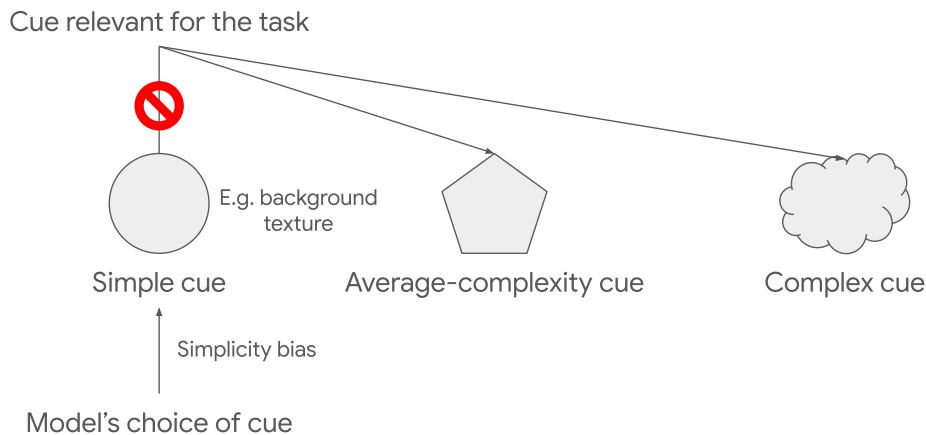
Figure 2.30: Example where the shortcut bias is unfavorable because of misspecification and does not lead to robust generalization. For OOD generalization tasks, simple cues may not work anymore.

For OOD generalization, the picture is a bit different. Simplicity bias is usually not welcome here because there are many OOD cases where the simplest cue is not good for generalization, as it is not relevant to the task. This causes problems during deployment, as the model's natural choice does not necessarily concur with the cue that would let the model generalize. For example, the background texture tends to be simple to recognize because we only have to look at very local parts of the image. The model might be able to use it to fit the training

data well, but it will not usually generalize to different domains. For fairness, simple cues (e.g., parent's income) may also not be *ethical* to use. We wish to prevent the model from using these cues. An example where the shortcut bias is unfavorable is given in Figure 2.30.

2.10   Identifying and Evaluating Misspecification

As discussed in earlier sections, underspecification (as defined in Definition 2.28) poses significant challenges to domain and cross-bias generalization. Therefore, it is crucial to diagnose whether our ML system suffers from misspecification. There are two main strategies to evaluate misspecification [42] (e.g., to determine whether the model uses too much context). Both are *counterfactual* evaluation methods, i.e., they manipulate the input to determine what cue the model is looking at. (Counterfactual evaluation always seeks answers to questions of the form "What would be the prediction if we changed ...?".) We either alter the task cue or the bias cue to observe the behavior of the model.

Altering the task cue. Here, the needed ingredients are the test set with task labels[9] and cue disentanglement (the ability to change cues in the input independently). The evaluation method is as follows.

- **Alter**: For every test sample, alter (or remove) the task-relevant cue.
- **Decide**: If the model performance *does not drop* significantly, our model is biased towards an irrelevant cue, meaning our system is misspecified.

Altering the bias cue. The needed ingredients are the same as when altering the task cue. The evaluation method is detailed below.

- **Alter**: For every test sample, alter (or remove) the bias cue.
- **Decide**: If the model performance *drops* significantly, our model is biased towards the altered cue, which, again, means that our system is misspecified.

Note: This way, we also know *what* our model is biased towards. With the previous method, we could only determine *whether* our model is biased.

These desiderata can be formulated in terms of *differences in accuracy/loss*. As long as there is a straightforward method that ranks biased and unbiased models correctly, it works well. Different papers do it differently.

Examples of changing the bias cue

Example 1 (Figure 2.27): The task cue is 'skateboard', and the bias cue is 'person'. It is improbable to see a skateboard on the road without a person on it: the task cue is highly correlated with the bias cue. We remove the bias cue and see how the score for 'skateboard' changes for a trained model. The needed ingredients are bounding box annotations/segmentation masks for objects and a good inpainting model. If the score for 'skateboard' drops a lot, the model has been relying on the bias cue.

Example 2 (Figure 2.31): The task cue is 'shape', and the bias cue is 'texture'. We consider two ways to change the bias cue: (1) Obtain a segmentation mask of the object and overlay a

---

[9]The task label is needed to calculate the loss.
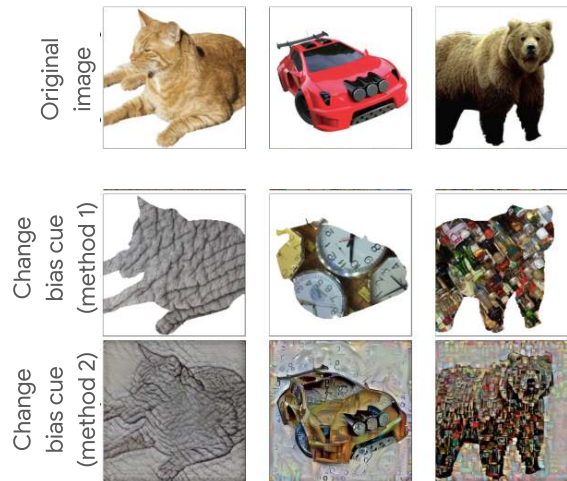
Figure 2.31: Example of two ways to change the (possible) bias cue of texture while preserving the task cue of shape.

texture image of choice. (2) Style-transfer [58] original image with a texture image of choice. The latter causes a less abrupt change: The image stays more reasonable. If the score of the true object drops significantly, the model has been relying a lot on the texture bias.

Example of changing the task cue

The following example is taken from [93]. The task cue is the overall positivity/negativity (sentiment) of the review. The bias cue is "anything but the task cue," e.g., the bag of words representation of a review. We let a human change the task cue (the sentiment analysis labels) by introducing minimal changes (a few words) in the sentences. If the score of the positive label does not change significantly after the update, the model does not rely on the overall meaning of the inputs.

To further illustrate the possible interventions the human annotator can make, we list some examples of changes made to the reviews:

- Recasting fact as "hoped for".
- Suggesting sarcasm.
- Inserting modifiers.
- Replacing modifiers.
- Inserting negative phrases.
- Diminishing via qualifiers.
- Changing the perspective.
- Changing the rating and some words.

Some of these are indeed very subtle and a model that is biased to the bag of words that appear in the review cannot react to such changes.

Easier →

		Easier →	
Settings	**All training samples are biased**	**Some unbiased training samples; no bias label**	**Some unbiased training samples; bias label available**
No test samples	Impossible to solve	Detect unbiased samples & put more weights on them **2**	Put more weights on unbiased samples **1**
A few unlabeled test samples	Train multiple hypotheses & do model selection	Combination of methods	
A few labeled test samples	Train multiple hypotheses & do model selection **3**		

Easier ↓

Figure 2.32: Overview of possible cross-bias generalization scenarios where the problem is made feasible by using different kinds of additional information. Scenario 2 can use prior knowledge about what the bias will be in the dataset. Under these assumptions, it can either detect unbiased samples and put more weight on them or make the final and intentionally biased models different (independent) in other ways. We will not discuss the upper version of scenario 3 (paper: "Test-Time Training with Self-Supervision for Generalization under Distribution Shifts" [190]), as we are not yet convinced that it is a possible case to solve in general deployment scenarios. We have no full trust yet.

2.11 Overview of Scenarios for Selecting the Right Features

So far, we have seen that predictions of models are often based on *bias* cues, while a key to generalization lies in their reliance on the *task* cues. How could we ensure that our model uses the task cue for its predictions? We will see approaches to selecting the right features for many settings. Let us quickly review some possible scenarios with extra information in Figure 2.32. The figure considers several settings that vary in their access to unbiased training samples or test samples as well as corresponding labels. It is important to understand that if we have no information apart from the diagonal dataset, the problem is conceptually unsolvable (top left cell). All the remaining cells describe different scenarios where generalization becomes possible again and we will discuss them in the next sections.

2.12 Scenario 1 for Selecting the Right Features

An example of this scenario is given in Figure 2.26. In this case, we have a small number of unbiased training samples (1% or even less) with bias labels. This is the easiest setting, as we know which samples are unbiased: we simply compare Y with Z. When they are equal, we have an on-diagonal sample. When they are unequal, the sample is off-diagonal (unbiased). We up-weight the off-diagonal samples and perform regular Empirical Risk minimization (ERM). This is the most naive approach, but it can perform well.

Additional Information: How to find unbiased samples?

What we discuss is, of course, a very simplistic setup. It is much more challenging to tell what samples are unbiased for the COCO dataset with 80 categories. However, if we know all target and bias labels, we can compute a matrix of co-occurrences between classes. We can then infer which images are more typical or atypical (e.g. a skateboard without a person is very unlikely), depending on the co-occurrence statistics of the labels. For very unlikely samples, we can, e.g., give a large weight during training. We generally weight samples more where the bias is either missing or different. However, there is an important caveat detailed in the example below.

Example: We have a dataset with many images of cats, dogs, and humans appearing together. The task is to predict whether an image contains a cat. If we see a sample with both a cat and a dog present, can we call it an atypical (unbiased, off-diagonal) sample and give it a large weight? *Only if the model is actually biased towards 'human'.* If the model is biased towards 'dog', this only aggravates the problem. Co-occurrence statistics are useful to give initial weights to samples but are usually only coarse proxies. Many biases are subtle and do not arise in an "interpretable" way. Determining weights post-hoc can directly act upon the problems of our model.

We can only determine weights in such cases using the following routine:
1. Train the network normally.
2. Determine to which combination of cues (such as 'dog' and 'human' jointly, just 'dog', or just 'human') it is biased towards using the unbiased test set.
3. Combat these biases by increasing the weights of samples that contain unlikely combinations of cues w.r.t. the present biases.

The computationally complex part here is annotation. Generally, it is a very strong assumption that we have labels for all possible cues! Once we have the task and bias labels, we create a counting matrix for co-occurrences which is easily computable on the CPU. In the COCO object detection dataset, there are many objects on a single image usually, so co-occurrences are easy to calculate. (Our assumption here is that labeling is complete.)

Additional Information: Model becoming biased again

What happens if we have biased fish images (i.e., fish are always in the hands of fishermen on the images) and we get unbiased images (e.g., fish in water), but the model learns shortcuts again (water background $\implies$ fish)? There are two solutions in general.

Bottom-up, incremental approach. We continuously search for the current model's biases by testing it for different sets of correlations (like testing our model's performance on fish images for a set of potential biases using unit tests). Such sets can be constructed by removing/replacing possible shortcuts (e.g., water background) in the original images. If we find that our model now uses some shortcuts, we incorporate new samples

without the corresponding biases. We continue doing this until the possible ways to learn shortcuts are saturated (i.e., it becomes more complicated than the task itself), and we are happy with the model. This approach does not guarantee that the ultimate model is unbiased, and usually, it is quite complicated to extensively test our model for potential biases.

Top-down approach. Let us assume that some explainability method provided us with a comprehensive and complete list of cues the model is actually looking at. In such a case, we first determine what cues are task cues and what are bias cues by human inspection. Then, we remove the bias cues from our model and include others they should be looking at more. **Disclaimer**: There is no technique for this in general, but it would be very nice to have one. This is very much the frontier of research in explainability.

2.12.1 Group DRO

Let us see how the availability of a small set of unbiased samples can be exploited in practice. In this section, we will discuss a method introduced in the paper "Distributionally Robust Neural Networks for Group Shifts: On the Importance of regularization for Worst-Case Generalization" [171], called Group Distributionally Robust Optimization (Group DRO). The goal of this method is to have the same accuracy for different bias groups (elements of the bias-task matrix depicted in Figure 2.23). This goal is achieved by minimizing the maximum loss across the groups. In the following paragraphs, we will discuss how this minimization is performed.

Optimization problem in Group DRO

In vanilla Empirical Risk Minimization (ERM), we have the following optimization problem:

$$\arg\min_{\theta \in \Theta} \mathbb{E}_{(x,y)\sim\hat{P}} \left[\ell(\theta; (x,y)) \right].$$

To achieve the goal of minimizing the maximum loss across the groups in DRO, the optimization problem is modified to the following one:

$$\arg\min_{\theta \in \Theta} \left\{ \mathcal{R}(\theta) := \sup_{Q \in \mathcal{Q}} \mathbb{E}_{(x,y)\sim Q} \left[\ell(\theta; (x,y)) \right] \right\}.$$

Here, $\mathcal{Q}$ encodes the possible test distributions we want to do well on. It should be chosen such that we are robust to distribution shifts, but we also do not get overly pessimistic models that optimize for implausible worst-case distributions Q.

Let us now choose $\mathcal{Q} := \left\{ \sum_{g=1}^{m} q_g P_g : q \in \Delta_m \right\}$ where Δ_m is the $(m-1)$-dimensional probability simplex and P_g are group distributions. These can correspond to arbitrary groups, but for our use case, the groups are based on spurious correlations. If we go back to our toy example of a (color, shape) dataset, then the individual groups can correspond to all possible (color, shape) combinations). Then

$$\mathcal{R}(\theta) = \max_{g \in \{1,\dots,m\}} \mathbb{E}_{(x,y)\sim P_g} \left[\ell(\theta; (x,y)) \right],$$

as the optimum of a linear program (the way we defined $\mathcal{Q}$) is always attained at a vertex (a particular P_g). Now, if we consider the empirical distributions $\hat{P}_g$, we get **Group DRO**:

$$\arg\min_{\theta \in \Theta} \left\{ \hat{\mathcal{R}}(\theta) := \max_{g \in \{1,\dots,m\}} \mathbb{E}_{(x,y) \sim \hat{P}_g} \left[\ell(\theta; (x, y)) \right] \right\}.$$

The learner aims to make predictions for *the worst-case group* better. Ideally, at the end of training, we have the same loss for each group (considering equal label noise across groups – if one group has huge corresponding label noise, the learner either overfits to the noise severely, which is suboptimal, or we do not have the same loss for each group at the end of training).

Examples for the groups in Group DRO

Toy example. In our previous example (Figure 2.26), all possible combinations of shape and color can be treated as a group. This way, we take into account the underrepresented combinations appropriately. We can also treat the on-diag and off-diag samples as the two groups, which might be a more stable choice if there are very few off-diag samples.

Faces. Assume a dataset of celebrities where the task is to predict gender from the image. The hair color annotation is also available. Further, assume that we have access to many diagonal samples and a small amount of off-diagonal samples where

$$
\begin{array}{lll}
P_1 : & \text{blonde female} & 50\% \\
P_2 : & \text{dark-haired male} & 40\% \\
P_3 : & \text{blonde male} & 3\% \\
P_4 : & \text{dark-haired female} & 7\%.
\end{array}
$$

If we just performed ERM/Regularized Risk Minimization (RRM), the model would usually predict based on a mixture of cues that would still favor the larger groups more and still be able to achieve high accuracy as we explicitly optimize on the average loss. For example, it could predict based on hair color: for dark-haired people, we could predict 'male', and for blonde individuals, we could predict 'female'. However, Group DRO helps us optimize on the worst-case combination, which can help prevent shortcuts.

Humans and skateboards. We consider one group comprising samples that contain a skateboard but not a human and another group comprising samples of skateboards with a human.

The Group DRO algorithm

Roughly speaking, Group DRO minimizes its optimization objective by performing the following steps:

1. Calculate losses for all groups.
2. Select the group with the maximal loss.
3. Set the model's gradient active only on the training samples from the worst-performing group.
4. Repeat.

The actual algorithm (Algorithm 1) is a bit more complicated. It considers an exponential moving average for the weights of different groups and performs gradient steps w.r.t. these weights. This modification allows the method to be trained with SGD. It also has nice convergence guarantees [171].

Algorithm 1: Online optimization algorithm for group DRO

Data: Step sizes $\eta_q, \eta_\theta; P_g$ for each $g \in \{1, \ldots, m\}$
Initialize $\theta^{(0)}$ and $q^{(0)}$
for $t = 1, \ldots, T$ **do**
    $g \sim \text{Uniform}(1, \ldots, m)$ `// Choose a group g at random`
    $x, y \sim P_g$ `// Sample x,y from group g`
    $q' \leftarrow q^{(t-1)}; q'_g \leftarrow q'_g \exp\left(\eta_q \ell\left(\theta^{(t-1)}; (x, y)\right)\right)$ `// Update weights for group g`
    $q^{(t)} \leftarrow q' / \sum_{g' \in \{1, \ldots, m\}} q'_{g'}$ `// Renormalize q`
    $\theta^{(t)} \leftarrow \theta^{(t-1)} - \eta_\theta q_g^{(t)} \nabla \ell\left(\theta^{(t-1)}; (x, y)\right)$ `// Use q to update θ`
end

> **Additional Information: Comments for the Group DRO algorithm**
>
> **Smoothed group-wise updates.** In Algorithm 1, $q_g^{(t)}$ influences the step size for the sample (and the corresponding group in general). This formulation can be considered a smoothed version of the original one, as we do not select the worst-performing group but still base the update on the group-wise performances.
>
> **Looking at the worst-group metric.** In general, the method performs worse than ERM on the average accuracy metric, as ERM directly optimizes on that. However, Group DRO shines on the worst-group accuracy metric, which is directly optimized by the method. ERM usually breaks down completely on the worst-group accuracy metric when there are notable group imbalances in the dataset.

Ingredients for Group DRO

In Group DRO, we have

$$\text{samples for } (X, Y, G) = (\text{input}, \text{output}, \text{group})$$

where the groups come from, e.g., spurious correlations or demographic groups. As we have group labels in addition to the usual setup (the difference is highlighted in red), we expect better worst-case accuracy. By explicitly optimizing on the worst-case spurious correlation/group, our model might generalize better in deployment.

> **Definition 2.30: Attribute Label**
>
> Attribute labels are indicators of all possible factors of variation in our data. Domain labels are a particular case of these.

The group label can not only be a bias or domain label, but even a general attribute label

(Definition 2.30). This additional cue makes cross-domain generalization less ill-posed.[10]

2.12.2 Domain-Adversarial Training of Neural Networks (DANN)
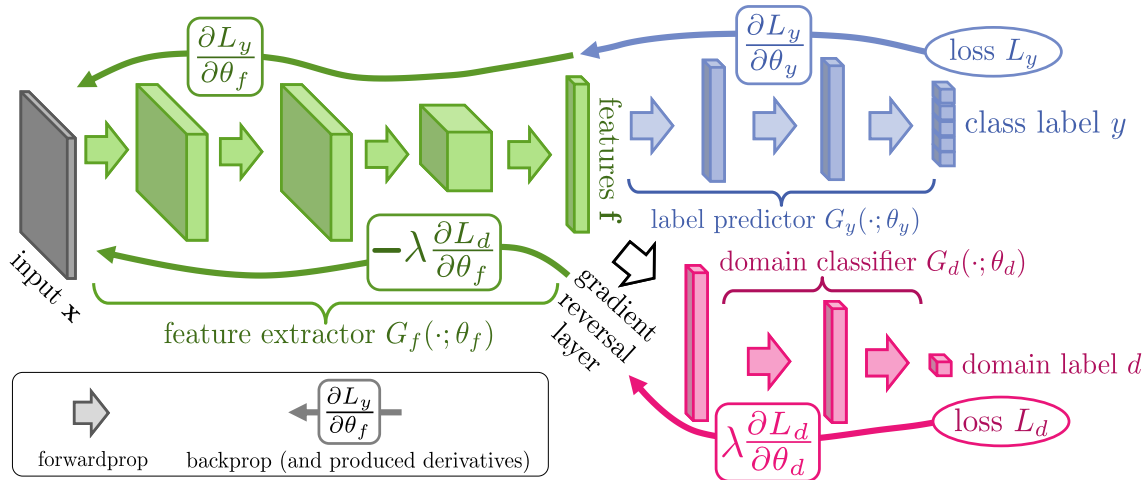
Figure 2.33: Overview of the DANN method. The feature extractor is encouraged to provide strong representations for predicting the class label and not contain any information about the domain label. Figure taken from the paper [56].

Apart from Group DRO, we have one more algorithm for Scenario 1 to discuss, called "Domain-Adversarial Training of Neural Networks" (DANN). The DANN method was introduced in the paper "Domain-Adversarial Training of Neural Networks" [56] and is another method to select good cues given bias labels by removing domain information from the intermediate features. An overview of the method taken from the original paper is shown in Figure 2.33.

The idea of the method is to add an additional head to the model (magenta in the image) that would predict bias labels (named domain labels in the paper) and adversarially train a feature extractor (green in the image) such that the features it extracts are maximally non-informative for the additional head to predict bias labels but still informative for the original head (blue in the image) to solve the main task.

This is achieved by splitting the training process into two parts. In the first part, the original head and feature extractor are jointly trained with gradient descent for the main task. In the second part, the bias-predicting head is trained with gradient descent for domain label prediction. The feature extractor parameters are adversarially trained with gradient *ascent* to maximize the loss of the bias-predicting head. Intuitively, we optimize the bias-predicting head to "squeeze out" any domain information left in the extracted features.

DANN can be assigned to the group of methods that select task cues given bias labels by removing information about the bias from the intermediate features.

---

[10]This is a more general statement than only considering cross-bias generalization – whenever we are presented with a (nearly) diagonal dataset, we need additional information, and this can happen in any cross-domain setting.

DANN Optimization Objective

We denote the prediction loss by

$$\mathcal{L}_y^i(\theta_f, \theta_y) = \mathcal{L}_y(G_y(G_f(x_i; \theta_f); \theta_y), y_i)$$

and the domain loss by

$$\mathcal{L}_d^i(\theta_f, \theta_d) = \mathcal{L}_d(G_d(G_f(x_i; \theta_f); \theta_d), d_i).$$

The training objective of DANN is

$$E(\theta_f, \theta_y, \theta_d) = \frac{1}{n} \sum_{i=1}^{n} \mathcal{L}_y^i(\theta_f, \theta_y) - \lambda \left(\frac{1}{n} \sum_{i=1}^{n} \mathcal{L}_d^i(\theta_f, \theta_d) + \frac{1}{n'} \sum_{i=n+1}^{N} \mathcal{L}_d^i(\theta_f, \theta_d) \right),$$

and the optimization problem is finding the saddle point $\hat{\theta}_f, \hat{\theta}_y, \hat{\theta}_d$ such that

$$\left(\hat{\theta}_f, \hat{\theta}_y \right) = \underset{\theta_f, \theta_y}{\arg\min}\, E \left(\theta_f, \theta_y, \hat{\theta}_d \right),$$

$$\hat{\theta}_d = \underset{\theta_d}{\arg\max}\, E \left(\hat{\theta}_f, \hat{\theta}_y, \theta_d \right).$$

Breaking DANN apart

First, we discuss the above formulation, which is for *domain adaptation*. The DANN method was originally proposed for this task. The first term of the training objective is the loss term for correct task label prediction on domain 1. The second term is the loss term for correct domain label prediction. We have two sums in the second term for domain 1 and domain 2 samples, respectively. For domain 2, we only have *unlabeled samples*, but we *do* have domain labels. The set of domain labels we have is simply {domain 1, domain 2}. Obtaining $\left(\hat{\theta}_f, \hat{\theta}_y \right)$ means minimizing the first term in θ_f, θ_y and maximizing the second term in θ_f. Similarly, we obtain $\hat{\theta}_d$ by minimizing the second term in θ_d.

Using DANN for cross-bias generalization

We can easily adapt the DANN formulation to cross-bias generalization. In particular, we treat y as the task label (e.g., shape: {circle, triangle, square}) and d as the bias label (e.g., color: {red, green, blue}). Here, the first term enforces correct predictions on both the biased and unbiased samples and the second term is used to kill out information about the bias from the representation. On off-diagonal samples, the bias label is not the task label, thus, f will be optimized to "forget" the bias labels while predicting the task labels correctly. We do not need unbiased samples as long as we have access to labeled samples from the target domain. It could happen that, e.g., the target domain is also biased, just in a different way than the training set.

We could also treat the set of biased samples as domain 1, the set of unbiased samples as domain 2, and use the original formulation of DANN for cross-bias generalization. This approach also works with target domain samples instead of unbiased ones.

Table 2.3: Classification accuracy on the Amazon reviews dataset. DANN leads to very slight improvements on most (source, target) pairs. Table taken from [56].

		ORIGINAL DATA			**MSDA REPRESENTATION**		
SOURCE	TARGET	DANN	NN	SVM	DANN	NN	SVM
BOOKS	DVD	.784	.790	**.799**	.829	.824	**.830**
BOOKS	ELECTRONICS	.733	.747	**.748**	**.804**	.770	.766
BOOKS	KITCHEN	**.779**	.778	.769	**.843**	.842	.821
DVD	BOOKS	.723	.720	**.743**	.825	.823	**.826**
DVD	ELECTRONICS	**.754**	.732	.748	**.809**	.768	.739
DVD	KITCHEN	**.783**	.778	.746	.849	**.853**	.842
ELECTRONICS	BOOKS	**.713**	.709	.705	**.774**	.770	.762
ELECTRONICS	DVD	**.738**	.733	.726	**.781**	.759	.770
ELECTRONICS	KITCHEN	**.854**	**.854**	.847	.881	**.863**	.847
KITCHEN	BOOKS	**.709**	.708	.707	.718	.721	**.769**
KITCHEN	DVD	**.740**	.739	.736	**.789**	**.789**	.788
KITCHEN	ELECTRONICS	**.843**	.841	.842	.856	.850	**.861**

Results of DANN for domain adaptation

To obtain good results with DANN, it is crucial to choose the hyperparameters well. All hyperparameters are chosen fairly in the paper, and there is no information leakage (e.g., by using the test set for choosing hyperparameters). Most hyperparameters are chosen using cross-validation and grid search on a log scale. Some are kept fixed or chosen among a set of sensible values. This is a usual practice in machine learning research. For large-scale experiments, the authors give fixed formulas for the LR decay and the scheduler for the domain adaptation parameter λ for the feature extractor (from 0 to 1), and fixed values for the momentum and the domain adaptation parameter for the domain classifier ($\lambda = 1$ to ensure that the domain classifier trains as fast as the label predictor).

The model is evaluated on generalizability between different *Amazon review topics* on the sentiment analysis task. The results are shown in the top table of Table 2.3. There is no significant difference between how NNs, SVMs, and DANN generalize. DANN is very slightly better on most review topic combinations. DANN is also evaluated on generalizability between MNIST and MNIST-M, SVHN and MNIST, and other datasets for the same task. The results of these experiments are shown in Table 2.4. On these benchmarks, DANN performed a lot better than NNs and SVMs.

We should always take a look at how papers choose hyperparameters. For a more complicated model, like DANN, there are many hyperparameters to choose from. Depending on how smartly we choose them, we get dramatically different results. When comparing methods, we also need to make sure that we spend the same resources for tuning the hyperparameters of all methods. The DANN paper provides fair comparisons.

Ingredients for DANN

In DANN, like in Group DRO, we also have access to

$$\text{samples for } (X, Y, G) = (\text{input}, \text{output}, \text{group}).$$

Method	Source Target	MNIST MNIST-M	Syn Numbers SVHN	SVHN MNIST	Syn Signs GTSRB
Source only		.5225	.8674	.5490	.7900
SA (Fernando et al., 2013)		.5690 (4.1%)	.8644 (−5.5%)	.5932 (9.9%)	.8165 (12.7%)
DANN		**.7666** (52.9%)	**.9109** (79.7%)	**.7385** (42.6%)	**.8865** (46.4%)
Train on target		.9596	.9220	.9942	.9980

Table 2.4: Classification accuracy of DANN compared to other methods on images of digits, assuming different source and target domains. "MNIST-M corresponds to difference-blended digits over non-uniform background. The first row corresponds to the lower performance bound (i.e. if no adaptation is performed). The last row corresponds to training on the target domain data with known class labels (upper bound on the DA performance)." [56] Table taken from [56].

The group label can again be a bias or domain label, but even a general attribute label. By using group supervision, we make cross-domain generalization less ill-posed.

2.13 Scenario 2 for Selecting the Right Features

Let us consider another cross-bias generalization setting from Figure 2.32: Scenario 2. Here, we consider an abundance of biased samples, a few available unbiased training samples ($<$ 1%), and no bias labels. As we do not know which samples are biased (we only have task labels), we need additional assumptions/information on the bias to solve the problem.[11] The question becomes how to identify unbiased samples and how to amplify them.

Before answering this question, let us first think about what assumptions we can make about the bias. The usual assumption is that the bias cue is simple and the task cue (what we want to learn) is more complex. For example, when the task is 'shape', and bias is 'color', this assumption holds. When reversing the roles, the assumption is violated.

This assumption on simplicity leads us to the following possible additional assumptions:

1. Bias is the first cue that a generic model learns.
2. Bias is the cue that is learned by a model of a certain limited capacity (i.e., by a short-sighted, myopic model).

Note: Sometimes, the assumption of the bias cue being a simpler cue than the task cue is violated. Practitioners have to understand the complexity of task cues and possible bias cues to successfully leverage methods with the above assumptions.

In the next sections, we will describe a set of methods that identify unbiased samples based on these assumptions. The framework depicted in Figure 2.34 is a clear basis for our discussion. Before diving into it, we would like to explain two important modules from this framework: "Intentionally biased model" and "Be different" supervision.

---

[11]Unless there are a lot of unbiased samples. Then we simply do ERM, and we basically have ID training.
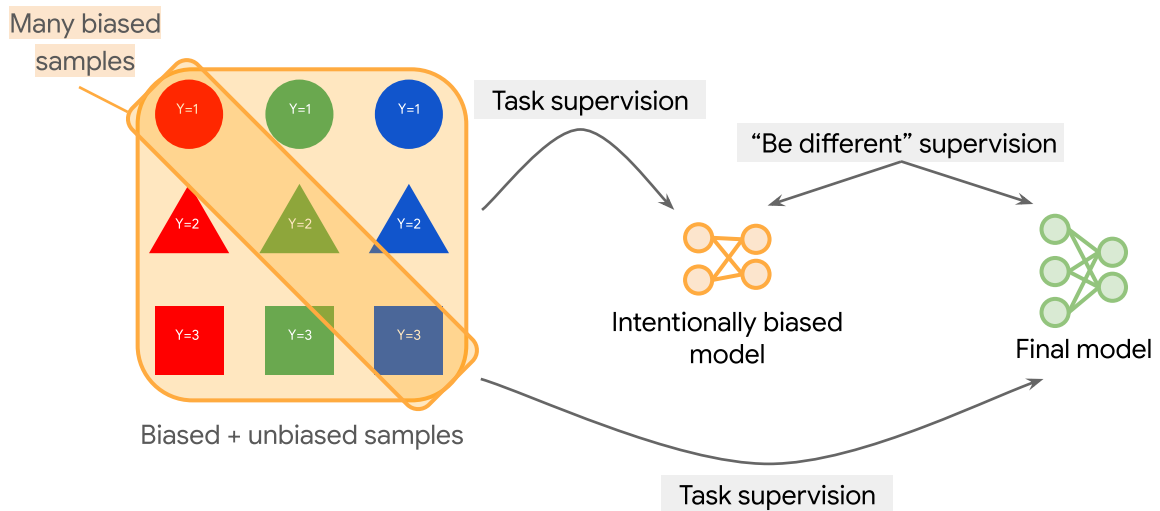
Figure 2.34: A general framework for selecting the right features, referred to as "Scenario 2" in the text. The *intentionally based model* is trained on the entire training set using task supervision.

Definition 2.31: Intentionally Biased Model

An intentionally biased model is designed to learn bias cues quickly, based on the assumptions we made before.

We consider several examples of an intentionally biased model:
- The model is trained for a small number of epochs. Whatever pattern that can already be learned in the first few epochs is considered bias.
- The model is not trained for a few epochs, but its initial correct predictions are amplified during training. This is conceptually very similar to the previous example but is perhaps more performant.
- The model has an architectural constraint: (1) CNN with a smaller receptive field. It can only extract very local information (e.g., texture patterns), not global shape. When the bias is 'texture', this is the way to go. (2) Transformer with shallow depth. It can only learn very simplistic relationships. When our bias is simple, this can work. (3) Single-modality model. This is one way to go when the actual task requires looking at multiple modalities to solve the problem.

Definition 2.32: "Be different" Supervision

"Be different" supervision is a type of regularization that forces the final model to be different from the intentionally biased model. The final model is trained on the original task loss with regularization based on the biased model. The biased model might be trained *before* the final model or *in tandem* (Learning from Failure: Section 2.13.1, ReBias: Section 2.13.2).

Examples of the "be different" supervision:

- Sample weighting based on biased model.
- Achieving representational independence.

2.13.1 Learning from Failure
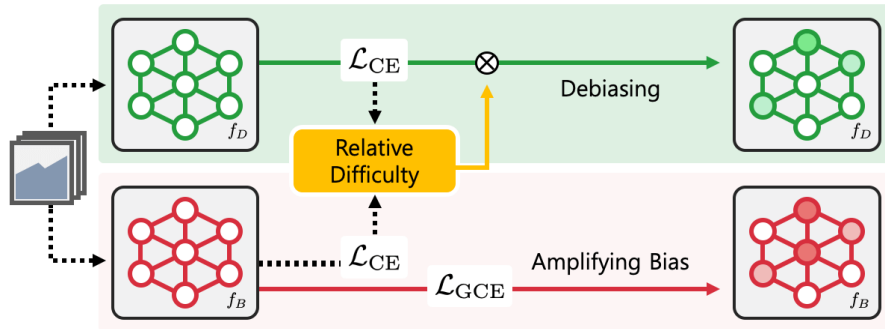
Figure 2.35: Overview of the Learning from Failure method. The intentionally biased model is used for determining the sample weights in the loss of the debiased model based on relative difficulty. Figure taken from the paper [145].

The method we consider now was introduced in the paper "Learning from Failure: Training Debiased Classifier from Biased Classifier" [145]. An overview, taken from the paper, is shown in Figure 2.35.

Here, an *intentionally biased* model is obtained by training with the following special loss that amplifies biases:

$$\mathcal{L}_{\mathrm{GCE}}(p(x;\theta), y) = \frac{1 - p_y(x;\theta)^q}{q}$$

where y is the GT class and $q > 0$.

This loss forces the intentionally biased model to focus on samples for which the predicted ground truth probability is already high. To understand why it happens, it can be shown that

$$\frac{\partial \mathcal{L}_{\mathrm{GCE}}(p(x;\theta), y)}{\partial \theta} = p_y(x;\theta)^q \frac{\partial \mathcal{L}_{\mathrm{CE}}(p(x;\theta), y)}{\partial \theta}$$

and as $q \downarrow 0, \mathcal{L}_{\mathrm{GCE}} \to \mathcal{L}_{\mathrm{CE}}$.

The final model f_D is trained to be *different* from the intentionally biased model by assigning the following sample weights:

$$\mathcal{W}(x) = \frac{\mathcal{L}_{\mathrm{CE}}(f_B(x), y)}{\mathcal{L}_{\mathrm{CE}}(f_B(x), y) + \mathcal{L}_{\mathrm{CE}}(f_D(x), y)}$$

where

$$\mathcal{L}_{\mathrm{CE}}(p(x;\theta), y) = -\log p_y(x;\theta).$$

Such weights force the final model to focus on the samples on which an intentionally biased model makes more mistakes. The final training algorithm, as presented in the paper, is shown in Algorithm 2.

Algorithm 2: Learning from Failure

Data: θ_B, θ_D, training set $\mathcal{D}$, learning rate η, number of iterations T
Initialize two networks $f_B(x; \theta_B)$ and $f_D(x; \theta_D)$
for $t = 1, \ldots, T$ **do**
    Draw a mini-batch $\mathcal{B} = \{(x^{(b)}, y^{(b)})\}_{b=1}^B$ from $\mathcal{D}$
    Update $f_B(x; \theta_B)$ by $\theta_B \leftarrow \theta_B - \eta \nabla_{\theta_B} \sum_{(x,y) \in \mathcal{B}} \mathcal{L}_{\mathrm{GCE}}(f_B(x), y)$
    Update $f_D(x; \theta_D)$ by $\theta_D \leftarrow \theta_D - \eta \nabla_{\theta_D} \sum_{(x,y) \in \mathcal{B}} \mathcal{W}(x) \cdot \mathcal{L}_{\mathrm{CE}}(f_D(x), y)$
end

Breaking LfF Apart

The intentionally biased model is trained with $\mathcal{L}_{\mathrm{GCE}}$. It amplifies whatever is predicted at the first iterations through the rest of the training. For example, if the model first learns 'color', then the loss amplifies color-based predictions and enforces the same predictions throughout training.

The final model is then forced to think of different hypotheses than the first model. If the biased model correctly predicts a sample, it gets less weight in the loss for the final model.

With $q > 0$, $\mathcal{L}_{\mathrm{GCE}}$ assigns more weight on confident samples, which results in larger gradient updates for these. The larger q is, the more the perfect predictions are weighted compared to imperfect ones. We train wrong predictions very slowly and initial predictions are strengthened over time.

Assumption on bias: Biases are the cues that are learned first. The method rewards easy samples to be learned quickly, and harder samples that were not predicted correctly to be given up by the intentionally biased model. Thus, this model is indeed biased towards easy cues.

For hard samples, $\mathcal{L}_{\mathrm{CE}}(f_B(x), y)$ is large throughout the training procedure. Both $\mathcal{L}_{\mathrm{CE}}(f_B(x), y)$ and $\mathcal{L}_{\mathrm{CE}}(f_D(x), y)$ are high for all $x \in \mathcal{D}$ in the beginning. The better f_D becomes on a sample, the more it is weighted (as $\mathcal{L}_{\mathrm{CE}}(f_D(x), y)$ decreases). However, the weight is multiplied by $\mathcal{L}_{\mathrm{CE}}(f_D(x), y)$, which balances this trend out. An illustration of $\mathcal{W}(x) \cdot \mathcal{L}_{\mathrm{CE}}(f_D(x), y)$ is given in Figure 2.36. Samples with high $\mathcal{W}(x)$ are ones that the biased model cannot handle well. Under our assumptions on the bias, samples with high $\mathcal{W}(x)$ are the unbiased ones. Thus, $\mathcal{W}(x)$ replaces the missing bias labels. Sample weights have a similar effect as the "upweighting" of the underrepresented group in Group DRO.

Note: In LfF, depending on the predictions of the first iteration, we choose the samples on which we wish and do not wish to train further. As a simpler baseline, we could also just train the intentionally biased model for 1-2 epochs but with the original cross-entropy loss. However, researchers usually prefer more 'continuous' solutions rather than such thresholds and rules of thumb.

Results of LfF

The paper showcases results on the Colored MNIST [12] dataset where the task is the shape of the digit and the bias is the color of the digit. A sample from this dataset can be seen in Figure 2.37, and the results are shown in Figure 2.38. The results show that if we train a model
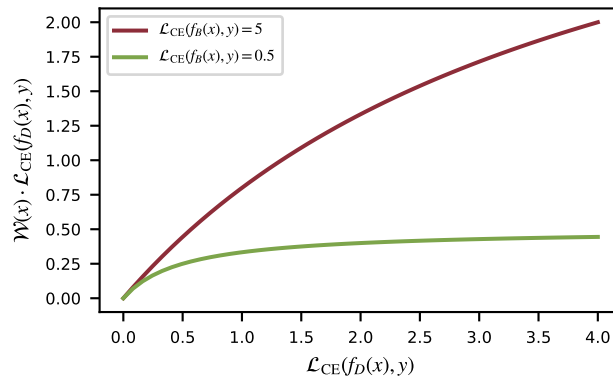
Figure 2.36: Illustration of $\mathcal{W}(x) \cdot \mathcal{L}_{\mathrm{CE}}(f_D(x), y)$ as a function of $\mathcal{L}_{\mathrm{CE}}(f_D(x), y)$ for $\mathcal{L}_{\mathrm{CE}}(f_B(x), y) \in \{0.5, 5\}$. Samples with a higher loss for the biased model are more important for the unbiased model.

for digit classification, it tends to pick up color much more quickly than the actual digit shape. The fact that we are improving performance by using LfF shows that

1. color is indeed learned first; and
2. color was indeed a bias that should be removed from consideration for digit recognition.

The lower the percentage of unbiased samples we include, the larger the relative effect LfF has over vanilla ERM. As expected, if we change the bias cue to digit and the task cue to color, LfF fails.

> **Additional Information: Changing the task on Colored MNIST**
>
> If color were the task and we evaluated LfF on Colored MNIST, we would see a drop in accuracy, as color is learned first, not digit. Thus, compared to the vanilla baseline, the final model generalization performance can verify whether the biased model learned the bias cue and whether what was learned was indeed a bias cue.

Figure 2.37: A representative sample from the Colored MNIST dataset [145].

Ingredients for LfF

In LfF, we use the usual ingredients for supervised learning (samples for $(X, Y) = (\mathrm{input}, \mathrm{output})$) plus an additional assumption:

> Biased samples are the ones that the intentionally biased model learns first.

Dataset	Ratio (%)	Vanilla ○	Ours ○	HEX ◑	REPAIR ●	Group DRO ●
Colored MNIST	95.0	$77.63_{\pm0.44}$	$\mathbf{85.39}_{\pm0.94}$	$70.44_{\pm1.41}$	$82.51_{\pm0.59}$	$84.50_{\pm0.46}$
	98.0	$62.29_{\pm1.47}$	$\mathbf{80.48}_{\pm0.45}$	$62.03_{\pm0.24}$	$72.86_{\pm1.47}$	$76.30_{\pm1.53}$
	99.0	$50.34_{\pm0.16}$	$\mathbf{74.01}_{\pm2.21}$	$51.99_{\pm1.09}$	$67.28_{\pm1.69}$	$71.33_{\pm1.76}$
	99.5	$35.34_{\pm0.13}$	$\mathbf{63.39}_{\pm1.97}$	$41.38_{\pm1.31}$	$56.40_{\pm3.74}$	$59.67_{\pm2.73}$

Figure 2.38: Results of LfF on Colored MNIST [12]. LfF is significantly better than vanilla training but also shows improvements compared to other debiasing methods. There are [Ratio]% biased samples and [1 - Ratio]% unbiased samples. Table taken from the paper [145].

Simply put: the bias is the simplest cue out of the ones with high predictive performance on this biased dataset. This is sometimes true, sometimes not. However, whenever it *is* true, we have a great solution for it. It can still happen, however, that the bias is not the easiest cue to learn. Then, the procedure misses the point.

> **Additional Information: When is something a "bias"?**
>
> What is bias is defined by humans. It is not an algorithmic concept. Only when humans declare something as a bias does it become a bias. It depends on the task (i.e., the setting we wish to generalize to) that humans specify. Whatever is not the task is a potential bias. Once we have a fixed task, we identify biases by, e.g., performing counterfactual evaluation.

> **Additional Information: Possible Extension of LfF**
>
> In the first few epochs, we could already condition the intentionally biased model to look for parameter regions where there are a lot more correct solutions with a bit more complex cues. This is already achieved in a way for regular LfF: when a very simple cue results in very poor training performance, it will not be chosen, no matter how simple it is.

2.13.2 ReBias: Representational regularization

Another method which introduces a similar concept to LfF is "Learning De-biased Representations with Biased Representations" [18]. An intuitive overview is given in Figure 2.39. The paper considers texture bias as the key problem to solve. We build CNNs that are *intentionally biased* towards texture by reducing their receptive fields. By constraining the intentionally biased model to this architecture, it is forced to capture local cues like texture.

The *final model* has a large receptive field. It might be, e.g., a ResNet-50. The *intentionally biased model* has a small receptive field, like the BagNet [24] model. A large receptive field can capture both local and global cues. However, the model might not look at global cues if the dataset is structured so that the net can simply learn very local cues to perform well.
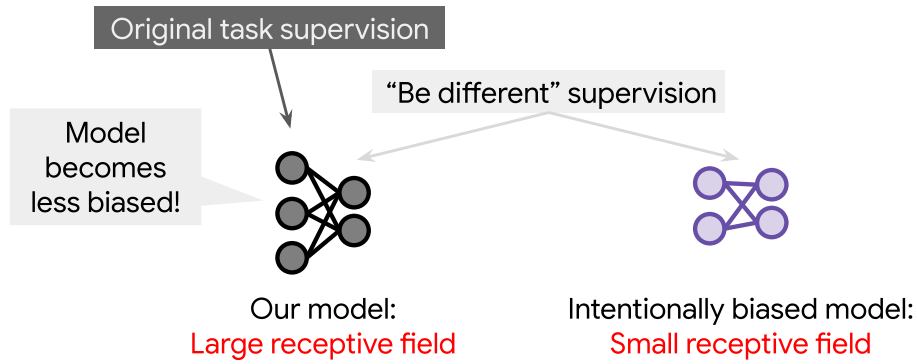
Figure 2.39: High-level and informal overview of the ReBias method. The intentionally biased model has a small receptive field to amplify texture bias. The debiased model is encouraged to be different from the intentionally biased one.

> **Additional Information: Receptive Fields Beyond the Input Image**
>
> We usually use padding to have the kernel centered at every pixel and influence the output dimensionality. If we use padding and regular (e.g., 3×3) convolutions, the receptive field of a deeper layer can be even beyond the image (but there, neurons only output zeros, constants, mirrors, or other redundant values). The field of view is huge in this case.

How can we perform "be different" supervision in this setup? The ReBias method leverages *statistical independence* instead of giving specific weights to samples. We train a debiased representation by encouraging the final model's outputs to be statistically independent from the intentionally biased model's outputs. We measure this independence with the Hilbert-Schmidt Independence Criterion (HSIC) between two random variables U, V:

$$\text{HSIC}^{k,l}(U, V) = \|C_{UV}^{k,l}\|_{\text{HS}}^2$$

where C is the cross-covariance operator in the Reproducing Kernel Hilbert Space (RKHS) corresponding to kernels k and l, and $\| \cdot \|_{\text{HS}}$ is the Hilbert-Schmidt norm which is, intuitively, a "non-linear version of the Frobenius norm of an infinite-dimensional covariance matrix." Kernels k and l correspond to random variables U and V, respectively. Essentially, we embed U and V in the infinite-dimensional RKHS corresponding to the kernels k and l, and compute their covariance there.

We use this criterion to make the invariances learned by these two models different. Our "be different" supervision is to minimize the HSIC between the two models.

Important property: It is well known [18] that for two random variables U, V and RBF kernels k, l,

$$\text{HSIC}^{k,l}(U, V) = 0 \iff U \perp\!\!\!\perp V.$$

> **Additional Information: Why is HSIC needed?**
>
> Why is making U and V uncorrelated not enough? If we have a covariance matrix and try to make it the identity matrix, we can enforce the correlation between the variables to be 0, but they will not necessarily be independent. There can be higher-order, non-linear dependencies. However, the HSIC lifts our random variables to an infinite-dimensional Hilbert space, and we consider the covariance "matrix" there. By doing so, we remove higher-order dependencies too at the same time, making the two variables truly independent.

If we just train a model f on some image classification dataset, it is very likely that the model finds a solution that is also representable by the small receptive field network g, as the model can usually perform well by looking at very small patches for predictions and we have previously discussed the simplicity bias of DNNs. Therefore, for our final model f and the intentionally biased model g, we want to enforce statistical independence $f(X) \perp\!\!\!\perp g(X)$ (that are random variables in $\mathbb{R}^C$) to ensure that the model f we find is not equivalent to some other network g with a small receptive field. The paper uses a finite-sample unbiased estimator $\mathrm{HSIC}_1^k(f(X), g(X))$ and the authors choose k and l to be both RBF kernels. Therefore, we consider the shorthand $\mathrm{HSIC}_1(f, g)$.

We know that

$$\mathrm{HSIC}(f(X), g(X)) = 0 \iff f(X) \text{ and } g(X) \text{ are independent}$$
$$\iff \text{The models } f, g \text{ have "orthogonal invariances".}$$

Let us detail the last equality further. If g discriminates color (i.e., its decision boundary separates objects of different colors), then f should learn invariance for color (i.e., changing of object color does not influence the distance w.r.t. decision boundary of f), and vice versa: if g is treating two samples similarly, then f should consider these far away from each other in the feature representation.[12] We train a de-biased representation by encouraging our model to be statistically independent of the intentionally biased representation.

ReBias Optimization Problem

The optimization problem in ReBias is

$$\operatorname*{arg\,min}_{g \in G} \underbrace{\mathcal{L}(g, x, y)}_{\text{Original task loss}} - \lambda_g \underbrace{\mathrm{HSIC}_1(f(x), g(x))}_{\text{Minimize independence}}$$

for the intentionally biased model and

$$\operatorname*{arg\,min}_{f \in F} \underbrace{\mathcal{L}(f, x, y)}_{\text{Original task loss}} + \lambda_f \underbrace{\mathrm{HSIC}_1(f(x), g(x))}_{\text{Maximize independence}}$$

---

[12] This is somewhat like a metric learning objective for HSIC.

for our model. The minimax game being solved is thus

$$\min_{f \in F} \max_{g \in G} \mathcal{L}(f) - \mathcal{L}(g) + \lambda \operatorname{HSIC}_1(f, g).$$

During training, we update f once, then update g for a fixed f n times ($n = 1$ in the official implementation). There are many other options, e.g., training f and g together on the same loss value.

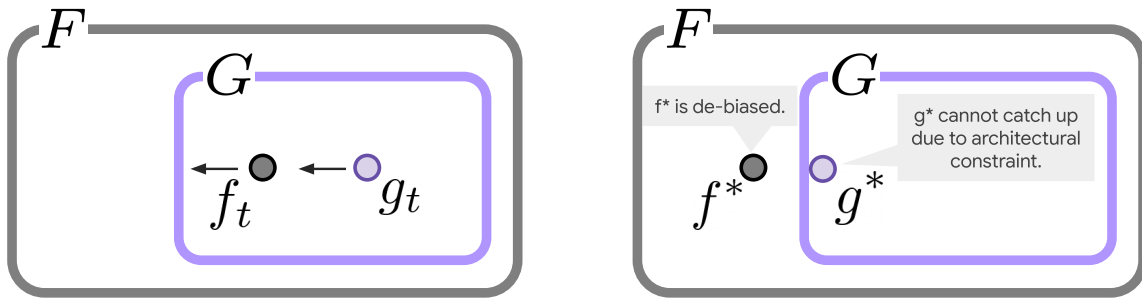### Illustration of Training ReBias

Figure 2.40: *Left.* Illustration of the ReBias minimax optimization problem. The function f is optimized to be highly different from g while still solving the task. The function g is incentivized to stay as similar to f as possible. *Right.* The optimal, de-biased function f^* leaves hypothesis space G. Therefore, no function exists in G that can match f^*.

The training procedure is illustrated in Figure 2.40. Functions f and g are elements of function spaces F and G, respectively. The function g is architecturally constrained and we have $G \subset F$. (We can pad kernels of g by zeros to get a valid model $f \in F$ that simulates a model with a small receptive field.) During the optimization procedure, g tries to catch up to f (solve the task and maximize dependence). In turn, f tries to be different (run away) from g (solve the task and minimize dependence). Eventually, after doing this for a few iterations, f finally escapes the set of models G. Thus, no function in G can represent f anymore (due to the architectural constraint), and f cannot leverage the simple cue that g uses. Now, e.g., f looks at global shapes instead of texture: f becomes debiased.

Figure 2.41: A versatile sample from the Colored MNIST dataset variant used in [18], taken from the paper.

Results of ReBias

Let us first consider the results of the method on the Colored MNIST dataset. In Colored MNIST, the color highly (or perfectly) correlates with the digit shape in the training set. Learning color is a shortcut to achieving high accuracy. Naively trained models will be biased towards color because of simplicity bias. The paper uses a variant of Colored MNIST in which all digits are white, but the background colors are perfectly correlated with the digits. A versatile sample from the dataset can be seen in Figure 2.41. The model we wish to debias is a LeNet architecture that can capture both color and shape. The intentionally biased model is

ρ	Biased						Unbiased					
	Vanilla	Biased	HEX	LearnedMixin	RUBi	ReBias	Vanilla	Biased	HEX	LearnedMixin	RUBi	ReBias
.999	**100.**	**100.**	71.3	2.9	99.9	**100.**	10.4	10.	10.8	12.1	13.7	**22.7**
.997	**100.**	**100.**	77.7	6.7	99.4	**100.**	33.4	10.	16.6	50.2	43.0	**64.2**
.995	**100.**	**100.**	80.8	17.5	99.5	**100.**	72.1	10.	19.7	78.2	**90.4**	76.0
.990	**100.**	**100.**	66.6	33.6	**100.**	**100.**	89.1	10.	24.7	88.3	**93.6**	88.1
avg.	**100.**	**100.**	74.1	15.2	99.7	**100.**	51.2	10.	18.0	57.2	60.2	**62.7**

Table 2.5: Results of ReBias on the Colored MNIST dataset, compared to various previous methods we do not cover in the book. ReBias shows notable improvements in some settings, resulting in the best average performance. ρ controls the fraction of unbiased samples in the training set. The vanilla and biased results show the performance of $f \in F$ and $g \in G$, respectively, trained using ERM. The results are taken from the paper [18].

a BagNet architecture that uses 1×1 convolutions. This is very much liable to overfit to color. The evaluation is performed both on biased and unbiased test sets. When evaluating the trained model on a test set with bias identical to the training set, we measure ID generalization performance. When using a test set with unbiased samples (colors randomly assigned to samples), the model relying on the bias cue would perform poorly. The exact results are shown in Table 2.5. ReBias improves unbiased accuracy while managing to retain biased accuracy.

Let us now turn to the task of action recognition with a strong static bias. The authors use the Kinetics dataset [32] for training the model, which has a strong bias towards static cues. For evaluation, the Mimetics dataset [204] is used that is ripped off the static cues and only contains the pure actions. The model to be debiased is a 3D-ResNet-18 [197] that can capture both temporal and static cues. The intentionally biased model is a 2D-ResNet-18, which can only capture static cues (i.e., cues from individual frames). As the results in Table 2.6 show, ReBias improves unbiased accuracy while also managing to improve biased accuracy.

Table 2.6: Results of ReBias on the Kinetics (biased) and Mimetics (unbiased) datasets, compared to various previous methods we do not cover in the book. Notably, ReBias is the most performant approach on *both* the biased and unbiased datasets. The vanilla and biased results show the performance of $f \in F$ and $g \in G$, respectively, trained using ERM. The results are taken from the paper [18].

Model description	Biased (Kinetics)	Unbiased (Mimetics)
Vanilla (3D-ResNet18)	54.5	18.9
Biased (2D-ResNet18)	50.7	18.4
LearnedMixin (Clark et al., 2019)	12.3	11.4
RUBi (Cadene et al., 2019)	22.4	13.4
ReBias	**55.8**	**22.4**

The Myopic Bias in Machine Learning

Let us first provide a definition for a *myopic model*.

> **Definition 2.33: Myopic Model**
>
> A myopic (short-sighted) model in ML refers to a model that is limited in its scope or focus, and, therefore, may not be able to capture all of the relevant features or information needed for robust prediction and decision-making.
>
> For example, a myopic model that only looks at texture may not be able to capture other important visual cues such as shape, motion, or context, which can be critical for accurate image recognition or object detection. Similarly, a myopic model that only considers static frames in a video may miss important information conveyed by the temporal dynamics of the video, such as motion or changes over time, which can be critical for accurate action recognition or activity detection. A language model may also focus on word-level cues for the overall sentiment of the sentence (e.g., frequency of 'not's).

The intentionally biased models we are considering in ReBias are myopic. The myopic bias appears a lot in ML in general: A very large model that is capable of modeling all kinds of relationships in the data does not learn complex relationships if the data itself is too simple and very conducive to simple cues.

To avoid myopic models, we introduce a second network that is very myopic, and use "be different" supervision, just like in LfF or ReBias. Our model will then be able to leverage complex cues and relationships better.

Example: Considering a language model $f \in F$ biased to word-level cues, we can "subtract" a simple Bag-of-Words (BoW) model (or a simple word embedding) $g \in G$ from the language model by using "be different" supervision to obtain more global reasoning and a more robust model.

Ingredients of ReBias

In ReBias, we use the usual ingredients for supervised learning (samples for $(X, Y) = (\text{input}, \text{output})$), plus additional assumptions:

1. The bias is "myopic".
2. One can intentionally confine a family of functions to be myopic.

Using "be different" supervision by enforcing statistical independence, we aim to obtain unbiased models that leverage robust cues.

2.14   Scenario 3 for Selecting the Right Features

The last cross-bias generalization scenario from Figure 2.32 we would like to discuss is Scenario 3. A more detailed overview of this setting can be seen in Figure 2.42. Here, we assume biased training samples (a labeled diagonal training set) without bias labels and a few labeled test samples. In such a case, we can train multiple models with diverse OOD behaviors, i.e., that have substantially different decision boundaries in the input space.[13] Considering the

---

[13]This is possible since the diagonal problem is highly ill-posed and the problem admits a versatile set of solutions.

shape-color dataset, the decision boundaries do not have to clearly cut any of the human-interpretable cues (predict only based on color vs. predict only based on shape). By having a diverse set of models, we can recognize samples according to many cues that *might be* task cues in deployment. We hope that one of them encodes what we want in the deployment scenario. At deployment time, we choose the right model from this set based on *a few labeled test samples*, then use it during deployment. This corresponds to *domain adaptation* or *test-time training* – different OOD generalization types where we have access to labeled deployment (test) samples. In practice, this is usually done in the context of *test-time training*, as the models are usually updated through the deployment procedure. By labeling samples on the fly (test-time training), one can perform model selection robustly.
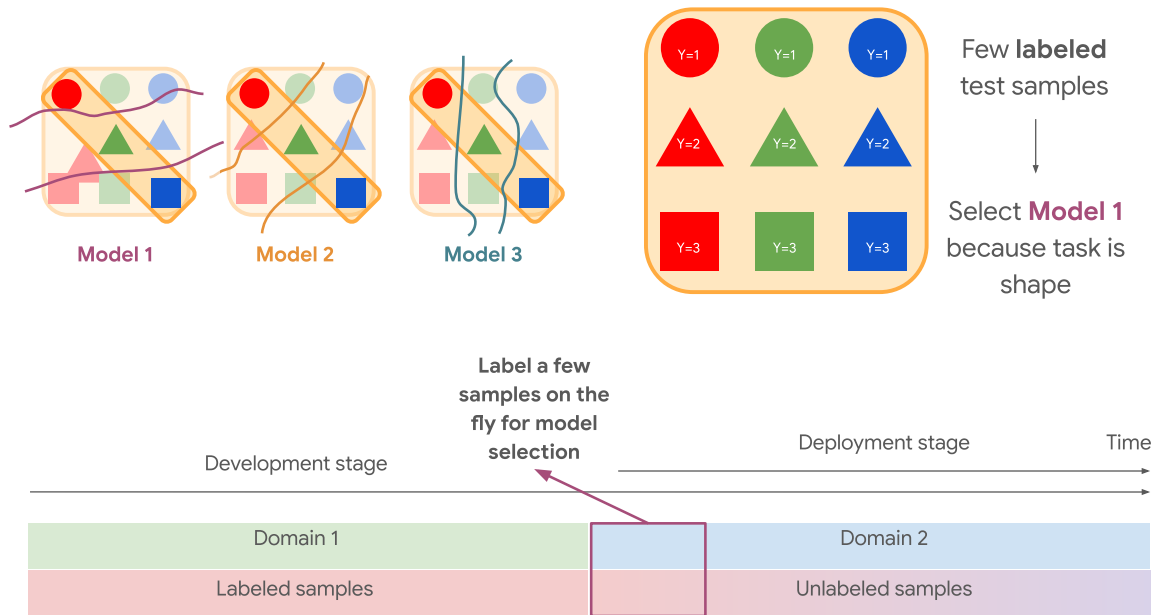

Figure 2.42: Overview of "Scenario 3" for selecting the right features.

Note: If we have deployment samples that are not unbiased and we also have bias labels, we can still use group DRO, sample weighting, and DANN.

In this scenario, the deployment domain is not necessarily unbiased. It can be equally biased, just in other ways. The labeled test samples decide the *task*. We select the best-performing model on the test dataset (which is usually very small in size), e.g., based on accuracy.

> **Additional Information: Difference between having a few test samples and a few unbiased training samples**
>
> In practice, we are unlikely to have unbiased samples at test time. When we do (e.g., as depicted in Figure 2.42), these scenarios *can* be the same, but there can also be other distributional shifts between train/test. The most likely case is that the deployment scenario contains many biased samples but with biases that differ from the training set biases. In this case, we aim to fine-tune/adapt our model to the specific bias at test time rather than aiming to do well on an unbiased set. Scenario 3 ensures that we can adapt to any shift at deployment (test) time, as we have direct access to deployment-time (test-

> time) data. This is a more straightforward setting, providing more information about the deployment scenario.

Here is one of the possible *recipes* to deal with such a setting:

1. Train an ensemble of models with some "diversity" regularization.
2. At test time, use a few labeled samples or human inspection (if it costs less than annotation time or we have special selection criteria) to select the appropriate model that generalizes well.

This recipe gives rise to two questions:

1. How can we know that the samples we base our decision on are representative of the whole test domain as time progresses?
2. How can we make sure that the set of models uses a diverse set of cues?

For the first question, we have two possible answers:

- Adapt the model very frequently (e.g., every batch of data we obtain).
- Trust that the deployment distribution is not going to change, e.g., for the next month, and update only every month.

By choosing either of the above, we also assume that these few labeled samples are enough to determine the most performant model in the deployment scenario.

For the second question, we cannot give a quick answer. If we naively train n models separately, all of them will likely focus on easy cues because of the simplicity bias of DNNs. That is why we need explicit regularization to enforce diversity. In the next section, we will focus on one of the methods that do exactly that.

2.14.1 Predicting is not Understanding

To look at one of the methods for diversifying models, let us discuss the paper "Predicting is not Understanding: Recognizing and Addressing Underspecification in Machine Learning" [193]. The intuition behind this method is that diverse ensemble training can be achieved by enforcing "independence" between models through the orthogonality of input gradients.

One way to achieve this is to add an orthogonality constraint to the loss.[14] Such a constraint can be represented as the squared cosine similarity of the input gradients for the same input:

$$\mathcal{L}_{\text{indep}}\left(\nabla_x f_{\theta_{m_1}}(x), \nabla_x f_{\theta_{m_2}}(x)\right) = \cos^2\left(\nabla_x f_{\theta_{m_1}}(x), \nabla_x f_{\theta_{m_2}}(x)\right).$$

Our goal is to have orthogonal input gradients. As this constraint is differentiable, we optimize it using Deep Learning (DL).

---

[14]HSIC could also be used as an independence criterion.

> **Additional Information: Shape of Gradients**
>
> In the orthogonality constraint, the gradients are of the logits, not of the loss. This results in a 4D tensor for multi-class classification. We simply flatten this tensor and calculate the squared cosine similarity. We only have a 1D output for binary classification, so the gradients will have the same shape as the input image. The paper focuses on binary classification. The independence loss used by the paper requires $\mathcal{O}(M^2)$ network evaluations, where M is the number of models in our diverse set.

Intuition for orthogonal input gradients

Suppose that we have two models, m_1 and m_2, and two different regions of the image: background and foreground. If m_2 is looking at the background, there is a significant focus on the background parts in the input gradient. We want the input gradient of m_1 to be orthogonal to that of m_2, as that will result in m_1 focusing more on the foreground.

Formal reasoning about independence

We define "independence" as the statistical independence of model outputs for a local Gaussian perturbation around every x in the input space. We measure the change in output for model 1 and model 2 using this Gaussian perturbation. The perturbation is small enough to approximate a model via its linear tangent function (input gradient). For infinitesimally small perturbations ($\sigma \downarrow 0$), changes in logits between x and $\tilde{x}$ can be approximated through linearization by the input gradients $\nabla_x f$. In particular, for $\sigma \downarrow 0$, the relative change in the logits from x to $\tilde{x}$ is exactly given by the directional derivative $\left\langle \nabla_x f(x), \frac{\tilde{x}-x}{\|\tilde{x}-x\|} \right\rangle$.

Why can we use the orthogonality of the input gradients for measuring statistical independence? It can be shown that the statistical independence of the model outputs is equivalent to the geometrical orthogonality of the input gradients when $\sigma \downarrow 0$ for the local Gaussian perturbation. The local independence for a particular input x is defined as

$$f_{\theta_1}(\tilde{x}) \perp\!\!\!\perp f_{\theta_2}(\tilde{x}), \tilde{x} \sim \mathcal{N}(x, \sigma I) \in \mathbb{R}^{d_{\text{in}}},$$

and global independence for a particular input x means that in a set of predictors $\{f_{\theta_1}, \ldots, f_{\theta_M}\}$, all pairs are locally independent around x.

We need one more ingredient to ensure that the models are diverse in meaningful ways. The set of orthogonal models increases exponentially with the input dimensionality. For images, we have overwhelmingly many orthogonal models – the input space might be close to being 1M-dimensional. However, the relevant subset of images that make sense inside this space is quite low-dimensional. This low-dimensional subset is the *data manifold*. We want to confine our exploration of decision boundaries to the manifold rather than the entire space. The reason is that diversification regularization without on-manifold constraints may result in models that are only diversified in the vast non-data-manifold dimensions, which means that they behave similarly on on-manifold samples.

We visualize an intuitive example of how models with orthogonal input gradients might still

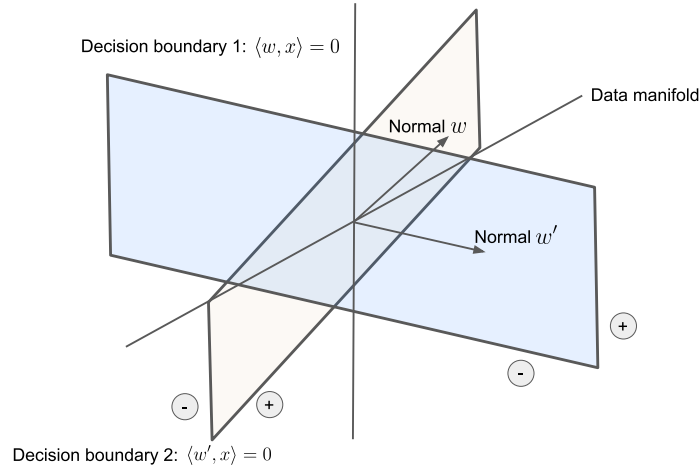behave identically on the data manifold in Figure 2.43.

Figure 2.43: Example that highlights the importance of the on-manifold constraint in Predicting is not Understanding [193]. We consider a 1D line as our data manifold and a binary classification problem. In the case of a linear classifier, the normal of the decision boundary is exactly the input gradient. If we project the decision boundaries onto the data manifold, they become identical. This means that even though the weights of the two models are orthogonal, they make identical decisions on the data manifold.

On-Manifold Constraints

In Predicting is not Understanding, the input gradient is regularized to be "on" the data manifold. We use a Variational Autoencoder (VAE) [100] to learn an approximation of the data distribution from unlabeled samples, i.e., to learn the data manifold $\mathcal{M}$. One can then project any vector $\in \mathbb{R}^{d_{\text{in}}}$ in the input space onto this data manifold by using the VAE $\text{proj}_{\mathcal{M}} \colon \mathbb{R}^{d_{\text{in}}} \times \mathbb{R}^{d_{\text{in}}} \to \mathcal{M}$. This VAE is trained to be capable of projecting a vector v (the gradient in the application) to the tangent plane of the manifold at point x. For OOD samples, this means that we want

$$\text{proj}_{\mathcal{M}}(x, v) \approx v \quad \forall x \sim P_{\text{OOD}}, x + v \sim P_{\text{OOD}},$$

which is achieved by training the VAE to reconstruct the OOD images and applying a similar series of transformations to the vector v as well. Further details can be read in the paper.

The on-manifold constraint is

$$\mathcal{L}_{\text{manifold}}(\nabla f(x)) = \| \text{proj}_{\mathcal{M}}(x, \nabla_x f(x)) - \nabla_x f(x) \|_2^2,$$

where $\text{proj}_{\mathcal{M}}$ is the projection of the gradient onto the tangent space of the manifold at point x. This loss term forces the input gradient to be aligned with the data manifold.

Used together with the independence constraint, the model is constrained to have orthogonal gradients that are roughly inside the data manifold. Intuitively, when the independence constraint influences a model's gradients in dimensions oriented outwards from the manifold, it does not impact its predictions on natural data. Consequently, models that produce identical predictions on every natural input could satisfy the independence constraint because their decision boundaries are identical when projected onto the manifold. This dras-

tically reduces the search space for new models and ensures that the next model in the ensemble will look at a *meaningful* new cue.

> **Additional Information: How to choose the dimensionality of the VAE latent space?**
>
> We do not have to know the dimensionality of the manifold, as it is perfectly fine if we choose the number of models more than that. Adding new models can still lead to more diversity, but it will be impossible to enforce the perfect orthogonality of the input gradients. It is also not a problem if we have fewer dimensions than the actual number of dimensions of the manifold in the VAE latent space, as one can embed higher-dimensional factors of variation into lower dimensions.

Putting it all together
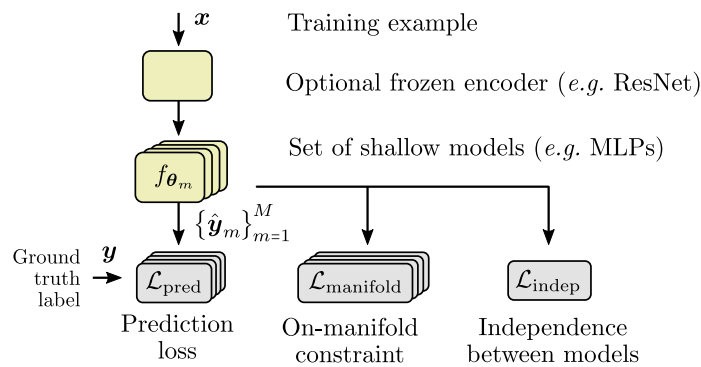
Figure 2.44: Overview of the Predicting is not Understanding method, taken from [193].

Our final loss function is

$$\mathcal{L}(\mathcal{D}_{\mathrm{tr}}, \theta_1, \ldots, \theta_M) = \sum_{(x,y) \in \mathcal{D}_{\mathrm{tr}}} \left[\frac{1}{M} \sum_{m=1}^{M} \mathcal{L}_{\mathrm{pred}}\left(y, \sigma(f_{\theta_m}(x))\right) \right.$$

$$+ \frac{1}{M^2} \sum_{m_1=1}^{M} \sum_{m_2=1}^{M} \lambda_{\mathrm{indep}} \mathcal{L}_{\mathrm{indep}}\left(\nabla_x f_{\theta_{m_1}}(x), \nabla_x f_{\theta_{m_2}}(x)\right)$$

$$\left. + \frac{1}{M} \sum_{m=1}^{M} \lambda_{\mathrm{manifold}} \mathcal{L}_{\mathrm{manifold}}\left(\nabla_x f_{\theta_m}(x)\right) \right]$$

that encapsulates the prediction losses, the independence losses, and the on-manifold losses for the M models.

Figure 2.45: Examples of collages of four tiles in Predicting is not Understanding [193].

Table 2.7: Results of Predicting is not Understanding [193]. The ERM baseline only learns to look at the MNIST tile and performs random prediction for all other cues. The independence loss is also not enough by itself. **Fine-tuning**: After training a set of models, the authors remove the independence and on-manifold constraints and fine-tune the models by applying a binary mask on the pixels/channels such that each model is fine-tuned only on the parts of the image most relevant to themselves (as measured by the magnitude of the gradient $\nabla_{\theta_m} f_{\theta_m}(x)$ among the models for each pixel/channel). **Pairwise combinations**: After training and fine-tuning a set of models, they combine the best of them (as given by our metric of choice on the OOD validation set) into a global one that uses all of the most relevant features. They train this global model from scratch, without regularizers, on masked data, using masks from the selected models combined with a logical OR. They repeat this pairwise combination as long as the accuracy of the global model increases. They always append the new models into the set of models. Independence + on-manifold constraints + VAE + FT + pairwise combinations (3x) performs best, achieving almost the upper bound (training on test-domain data). The upper bound accuracy can be achieved, e.g., if we have four models specializing perfectly in different quadrants. Based on these results, the models are indeed very diverse.

Collages dataset (accuracy in %)	MNIST	SVHN	Fashion	CIFAR-10	Average
			Best model on		
Upper bound (training on test-domain data)	99.9	92.4	80.8	68.6	85.5
ERM Baseline	99.8	50.0	50.0	50.0	62.5
Spectral decoupling [156]	99.9	49.8	50.6	49.9	62.5
With penalty on L1 norm of gradients	98.5	49.6	50.5	50.0	62.1
With penalty on L2 norm of gradients [81]	96.6	52.1	52.3	54.3	63.8
Input dropout (best ratio: 0.9)	97.4	50.7	56.1	52.1	64.1
Independence loss (cosine similarity) [169]	99.7	50.4	51.5	50.2	63.0
Independence loss (dot product) [194]	99.5	53.5	53.3	50.5	64.2
With many more models					
Independence loss (cosine similarity), <u>1024</u> models	99.5	58.1	66.8	63.0	71.9
Independence loss (dot product), <u>128</u> models	98.7	84.9	71.6	61.5	79.2
Proposed method (only 8 models)					
Independence + on-manifold constraints, PCA	97.3	69.8	62.2	60.0	72.3
Independence + on-manifold constraints, VAE (*)	96.5	85.1	61.1	62.1	76.2
(*) + FT (fine-tuning)	99.7	90.9	81.4	67.4	84.8
(*) + FT + pairwise combinations (1×)	99.9	92.2	79.3	66.3	84.4
(*) + FT + pairwise combinations (2×)	99.9	92.5	80.2	67.5	85.0
(*) + FT + pairwise combinations (3×)	**99.9**	**92.3**	**80.8**	**68.5**	**85.4**

Results of Predicting is not Understanding

The method is evaluated on a *collage dataset* with controllable correlation among the four collage images. The four datasets used are MNIST, Fashion MNIST [216], CIFAR [110], and SVHN [147]. We have ten classes for each dataset. We put one sample from each dataset in a window of four elements, as shown in Figure 2.45. During training, a dataset with a perfect correlation between the four labels is used (e.g., the meta-class 0 is the quadruple (zero, pullover, automobile, zero)). This is a biased, diagonal dataset. For evaluation, we use a dataset with no correlation between the four labels. This is an unbiased dataset with off-diagonal samples as well. We label this w.r.t., e.g., MNIST or CIFAR, and – based on the results – we find out what cue (which quadrant) each model learned.

This is *cross-bias generalization*: At test time, we break the correlation among the four quadrants. Our expectation is that by training independent models, we should be able to get four different models that look at different quadrants. The results are shown in Table 2.7. To measure performance, the authors perform test-time oracle model selection: they are using test-time information by choosing the best model on each test-time dataset. A possible justification of test-time oracle model selection is that, in practice, we have a few labeled samples at test time to select the most performant model. These reported numbers show the upper bound on the performance in the scenario above because they are chosen based on the entire test set, not just a few samples.

Ingredients for Predicting is not Understanding

The Predicting is not Understanding method uses the usual ingredients for supervised learning (samples for $(X, Y) = (\text{input}, \text{output})$) to find a set of diverse hypotheses $m_1, \ldots, m_N$. The model selection takes place at *test time*. This work only shows the upper bound of attainable performance using the perfect test-time model selection.

Definition 2.34: Kullback-Leibler Divergence

The Kullback-Leibler divergence from distribution Q to distribution P with densities q, p is given by

$$\mathrm{KL}\left(P \| Q\right) = \int_{\mathcal{X}} p(x) \log \frac{p(x)}{q(x)} \, dx.$$

Additional Information: Independence and Input Gradients

Proposition 2.1. *A pair of predictors $f_{\theta_1}, f_{\theta_2}$ are locally independent at x iff the mutual information* $\mathrm{MI}(f_{\theta_1}(\tilde{x}), f_{\theta_2}(\tilde{x})) = 0$ *with $\tilde{x} \sim \mathcal{N}(x, \sigma^2 I)$ [193].*

Proof. A pair of predictors $f_{\theta_1}, f_{\theta_2}$ are defined to be locally independent at x iff their predictions are statistically independent for Gaussian perturbations around x:

$$f_{\theta_1}(\tilde{x}) \perp\!\!\!\perp f_{\theta_2}(\tilde{x})$$

with $\tilde{x} \sim \mathcal{N}(x, \sigma^2 I)$.
The definition of mutual information is

$$\mathrm{MI}(f_{\theta_1}(\tilde{x}), f_{\theta_2}(\tilde{x})) = D_{\mathrm{KL}}(P_{f_{\theta_1}(\tilde{x}), f_{\theta_2}(\tilde{x})} \| P_{f_{\theta_1}(\tilde{x})} \otimes P_{f_{\theta_2}(\tilde{x})}).$$

It is a well-known fact that $D_{\mathrm{KL}}(P \| Q) = 0 \iff P \equiv Q$. From this, we immediately see that

$$\mathrm{MI}(f_{\theta_1}(\tilde{x}), f_{\theta_2}(\tilde{x})) = 0 \iff f_{\theta_1}(\tilde{x}) \perp\!\!\!\perp f_{\theta_2}(\tilde{x}).$$

$\square$

For infinitesimally small perturbations ($\sigma \downarrow 0$), the variables $f_{\theta_1}(\tilde{x}), f_{\theta_2}(\tilde{x})$ can be approximated through linearization wrt. the input gradients $\nabla_x f$:

$$f(\tilde{x}) \approx f(x) + \nabla_x f(x)^\top (\tilde{x} - x) =: \hat{f}(\tilde{x}).$$

Claim 2.1. *Following the above definition, $\hat{f}_{\theta_1}(\tilde{x})$ and $\hat{f}_{\theta_2}(\tilde{x})$ are 1D Gaussian variables.*

Proof. By definition, $\hat{f}(\tilde{x}) = f(x) + \nabla_x f(x)^\top (\tilde{x} - x) = \underbrace{\nabla_x f(x)^\top}_{A :=} \tilde{x} + \underbrace{f(x) - \nabla_x f(x)^\top x}_{b :=}.$

As $\tilde{x} \sim \mathcal{N}(x, \sigma^2 I)$, we know that

$$\hat{f}(\tilde{x}) = A\tilde{x} + b \sim \mathcal{N}(Ax + b, \sigma^2 AA^\top)$$
$$\sim \mathcal{N}(Ax + f(x) - Ax, \sigma^2 AA^\top)$$
$$\sim \mathcal{N}(f(x), \sigma^2 \nabla_x f(x)^\top \nabla_x f(x)).$$

Substituting f_{θ_1} or f_{θ_2} into f directly gives the statement. $\square$

Claim 2.2. *The correlation of $\hat{f}_{\theta_1}(\tilde{x})$ and $\hat{f}_{\theta_2}(\tilde{x})$ is given by $\cos(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x))$.*

Proof. We know from before that

$$\hat{f}_{\theta_1}(\tilde{x}) \sim \mathcal{N}(f_{\theta_1}(x), \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x))$$
$$\hat{f}_{\theta_2}(\tilde{x}) \sim \mathcal{N}(f_{\theta_2}(x), \sigma^2 \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)).$$

It follows that

$$\rho(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) = \frac{\text{Cov}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x}))}{\sqrt{\sigma^4 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)}}.$$

First, we calculate the covariance:

$$\begin{aligned}
\text{Cov}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) &= \mathbb{E}[\hat{f}_{\theta_1}(\tilde{x})\hat{f}_{\theta_2}(\tilde{x})] - \mathbb{E}[\hat{f}_{\theta_1}(\tilde{x})]\mathbb{E}[\hat{f}_{\theta_2}(\tilde{x})] \\
&= \mathbb{E}[\hat{f}_{\theta_1}(\tilde{x})\hat{f}_{\theta_2}(\tilde{x})] - f_{\theta_1}(x)f_{\theta_2}(x) \\
&= \mathbb{E}[(f_{\theta_1}(x) + \nabla_x f_{\theta_1}(x)^\top (\tilde{x} - x))(f_{\theta_2}(x) + \nabla_x f_{\theta_2}(x)^\top (\tilde{x} - x))] \\
&\qquad\qquad - f_{\theta_1}(x)f_{\theta_2}(x) \\
&= f_{\theta_1}(x)f_{\theta_2}(x) + f_{\theta_1}(x)\mathbb{E}[\nabla_x f_{\theta_2}(x)^\top (\tilde{x} - x)] + f_{\theta_2}(x)\mathbb{E}[\nabla_x f_{\theta_1}(x)^\top (\tilde{x} - x)] \\
&\quad + \mathbb{E}[\nabla_x f_{\theta_1}(x)^\top (\tilde{x} - x)\nabla_x f_{\theta_2}(x)^\top (\tilde{x} - x)] - f_{\theta_1}(x)f_{\theta_2}(x) \\
&= \mathbb{E}[\nabla_x f_{\theta_1}(x)^\top (\tilde{x} - x)\nabla_x f_{\theta_2}(x)^\top (\tilde{x} - x)] \\
&= \nabla_x f_{\theta_1}(x)^\top \mathbb{E}[(\tilde{x} - x)(\tilde{x} - x)^\top]\nabla_x f_{\theta_2}(x) \\
&= \nabla_x f_{\theta_1}(x)^\top \left(\text{Cov}(\tilde{x} - x) + \mathbb{E}[(\tilde{x} - x)]\mathbb{E}[(\tilde{x} - x)]^\top\right)\nabla_x f_{\theta_2}(x) \\
&= \nabla_x f_{\theta_1}(x)^\top (\sigma^2 I)\nabla_x f_{\theta_2}(x) \\
&= \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x).
\end{aligned}$$

Plugging this back into the correlation formula, we obtain

$$\rho(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) = \frac{\mathrm{Cov}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x}))}{\sqrt{\sigma^4 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)}}$$

$$= \frac{\sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x)}{\sigma^2 \sqrt{\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x)} \sqrt{\nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)}}$$

$$= \frac{\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x)}{\sqrt{\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x)} \sqrt{\nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)}}$$

$$= \cos(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x)).$$

$\square$

Claim 2.3. *The mutual information of $\hat{f}_{\theta_1}(x)$ and $\hat{f}_{\theta_2}(x)$ is given by $-\frac{1}{2}\log(1 - \cos^2(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x)))$.*

Proof.

$$\mathrm{MI}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) = D_{\mathrm{KL}}(P_{\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})} \| P_{\hat{f}_{\theta_1}(\tilde{x})} \otimes P_{\hat{f}_{\theta_2}(\tilde{x})})$$

$$= \int_{x_1 \in \mathcal{X}} \int_{x_2 \in \mathcal{X}} p(\hat{f}_{\theta_1}(x_1), \hat{f}_{\theta_2}(x_2)) \log \frac{p(\hat{f}_{\theta_1}(x_1), \hat{f}_{\theta_2}(x_2))}{p(\hat{f}_{\theta_1}(x_1)) p(\hat{f}_{\theta_2}(x_2))} dx_2 dx_1$$

$$= \mathbb{H}(\hat{f}_{\theta_1}(\tilde{x})) + \mathbb{H}(\hat{f}_{\theta_2}(\tilde{x})) - \mathbb{H}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})).$$

Now we notice that $\hat{f}_{\theta_1}(\tilde{x})$ and $\hat{f}_{\theta_2}(\tilde{x})$ are also *jointly* Gaussian:

$$\begin{pmatrix} \hat{f}_{\theta_1}(\tilde{x}) \\ \hat{f}_{\theta_2}(\tilde{x}) \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_{\hat{f}_{\theta_1}(\tilde{x})} \\ \mu_{\hat{f}_{\theta_2}(\tilde{x})} \end{pmatrix}, \begin{bmatrix} \mathrm{Var}(\hat{f}_{\theta_1}(\tilde{x})) & \mathrm{Cov}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) \\ \mathrm{Cov}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) & \mathrm{Var}(\hat{f}_{\theta_2}(\tilde{x})) \end{bmatrix} \right)$$

$$\sim \mathcal{N}\left(\begin{pmatrix} f_{\theta_1}(x) \\ f_{\theta_2}(x) \end{pmatrix}, \begin{bmatrix} \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) & \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x) \\ \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x) & \sigma^2 \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x) \end{bmatrix} \right).$$

Below, we derive the formula for the entropy of a multivariate Gaussian $x \sim \mathcal{N}(\mu, \Sigma) \in \mathbb{R}^n$:

$$\mathbb{H}(x) = -\int p(x) \log p(x) dx$$

$$= -\mathbb{E}_x[\log \mathcal{N}(x \mid \mu, \Sigma)]$$

$$= -\mathbb{E}_x \left[\log \left(\frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp \left(-\frac{1}{2}(x-\mu)^\top \Sigma^{-1}(x-\mu) \right) \right) \right]$$

$$= \mathbb{E}_x \left[\frac{n}{2} \log(2\pi) + \frac{1}{2} \log|\Sigma| + \frac{1}{2}(x-\mu)^\top \Sigma^{-1}(x-\mu) \right]$$

$$= \frac{n}{2} \log(2\pi) + \frac{1}{2} \log|\Sigma| + \frac{1}{2} \mathbb{E}_x[(x-\mu)^\top \Sigma^{-1}(x-\mu)]$$

$$= \frac{n}{2} \log(2\pi) + \frac{1}{2} \log|\Sigma| + \frac{1}{2} \mathbb{E}_x[\text{tr}((x-\mu)^\top \Sigma^{-1}(x-\mu))]$$

$$= \frac{n}{2} \log(2\pi) + \frac{1}{2} \log|\Sigma| + \frac{1}{2} \mathbb{E}_x[\text{tr}(\Sigma^{-1}(x-\mu)(x-\mu)^\top)]$$

$$= \frac{n}{2} \log(2\pi) + \frac{1}{2} \log|\Sigma| + \frac{1}{2} \text{tr}(\Sigma^{-1} \underbrace{\mathbb{E}_x[(x-\mu)(x-\mu)^\top]}_{\Sigma})$$

$$= \frac{n}{2}(1 + \log(2\pi)) + \frac{1}{2} \log|\Sigma|.$$

Finally, we plug this into our formula for the mutual information:

$$\text{MI}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x})) = \mathbb{H}(\hat{f}_{\theta_1}(\tilde{x})) + \mathbb{H}(\hat{f}_{\theta_2}(\tilde{x})) - \mathbb{H}(\hat{f}_{\theta_1}(\tilde{x}), \hat{f}_{\theta_2}(\tilde{x}))$$

$$= \frac{1}{2}(1 + \log(2\pi)) + \frac{1}{2} \log \left(\sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \right)$$

$$+ \frac{1}{2}(1 + \log(2\pi)) + \frac{1}{2} \log \left(\sigma^2 \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x) \right)$$

$$- 1 - \log(2\pi) - \frac{1}{2} \log(\sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \cdot \sigma^2 \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)$$

$$- \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x) \cdot \sigma^2 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x))$$

$$= \frac{1}{2} \log \frac{\sigma^4 \nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \cdot \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)}{\sigma^4 \left(\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \cdot \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x) - (\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x))^2 \right)}$$

$$= -\frac{1}{2} \log \left(1 - \frac{(\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_2}(x))^2}{\nabla_x f_{\theta_1}(x)^\top \nabla_x f_{\theta_1}(x) \cdot \nabla_x f_{\theta_2}(x)^\top \nabla_x f_{\theta_2}(x)} \right)$$

$$= -\frac{1}{2} \log(1 - \cos^2(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x))).$$

$$\square$$

Putting everything together: For an infinitesimal perturbation ($\sigma \downarrow 0$), we know that $\hat{f}_{\theta_1}(\tilde{x}) = f_{\theta_1}(\tilde{x})$ and $\hat{f}_{\theta_2}(\tilde{x}) = f_{\theta_2}(\tilde{x})$, i.e., the linearization is exact. Of course, we have to re-linearize after every gradient step. By driving the mutual information to zero, we enforce statistical independence between $f_{\theta_1}(\tilde{x})$ and $f_{\theta_2}(\tilde{x})$. It is also easy to see that for

$\sigma \downarrow 0$,

$$\min_{\theta_1, \theta_2} \mathrm{MI}(f_{\theta_1}(\tilde{x}), f_{\theta_2}(\tilde{x})) = \min_{\theta_1, \theta_2} -\frac{1}{2} \log(1 - \cos^2(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x)))$$

$$= \max_{\theta_1, \theta_2} \log(1 - \cos^2(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x)))$$

$$= \max_{\theta_1, \theta_2} 1 - \cos^2(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x))$$

$$= \min_{\theta_1, \theta_2} \cos^2(\nabla_x f_{\theta_1}(x), \nabla_x f_{\theta_2}(x)).$$

Therefore, the local independence loss

$$\mathcal{L}_{\mathrm{indep}}(\nabla_x f_{\theta_{m_1}}(x), \nabla_x f_{\theta_{m_2}}(x)) = \cos^2(\nabla_x f_{\theta_{m_1}}(x), \nabla_x f_{\theta_{m_2}}(x))$$

for a pair of models (m_1, m_2) indeed encourages the statistical independence of the models' outputs, considering an infinitesimal Gaussian perturbation around the input x. The obvious minimizer of the term is any constellation where the two input gradients are orthogonal.

Note: It is also easy to see from the correlation and mutual information formulas that for Gaussian variables, zero correlation is equivalent to independence. This is, of course, not true in general.

2.15 Adversarial OOD Generalization

OOD generalization is about dealing with uncertainty. It is easy to make a model generalize well to a single possible environment. As we introduce more environments, this becomes harder and harder until we arrive at an infinite number of environments or "any environment". This tendency is illustrated in Figure 2.46. As we have more and more knowledge about what will happen at deployment time, the space of possible environments shrinks, and thus we become more certain. A parallel can be drawn with the notion of entropy: If we already have much knowledge, additional information has a small entropy.
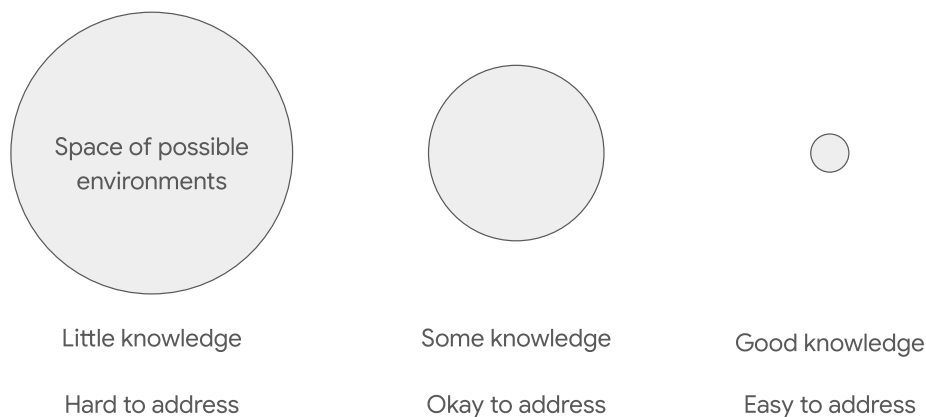

Figure 2.46: The size of the space of possible environments shrinks as we have more and more information about deployment.

The question is: How can we take care of an infinite number of possible environments? There are two general methods for dealing with uncertainty when we do not know the deployment environment perfectly (or the *enemy* who is trying to give us a hard environment):

1. **Make an educated guess.** For a good guess, this is a nice, practical solution that is easy to carry out. However, we obtain no guarantees: We do not know if we made the right guess. In the worst case, we are not making any progress. Many methods seen so far fall into this category, e.g., ReBias, Predicting is not Understanding, and Learning from Failure (by guessing the bias).

2. **Prepare for the worst.** Here, we have a so-called *adversarial environment*. By following this principle, we can obtain theoretical lower bound guarantees: Our model's performance against the worst-case environment provides a lower bound on its performance against the space of possible environments. (We are safe for the worst-case scenario from a set of possible scenarios, so we are also safe for all of them.) An important caveat is that the guarantee is only within the pre-set space of possible environments (the strategy space). Outside of this, we have no guarantees. This approach can also lead to unrealistically pessimistic solutions.

As they both have their pros and cons, there is no single right answer: it is a matter of choice and depends on our application.

So far, we have only considered OOD generalization methods for making an educated guess. Now, let us discuss *adversarial generalization* that comprises methods that prepare for the worst.

> **Definition 2.35: Adversarial Generalization**
>
> Adversarial generalization is an ML technique for "preparing for the worst-case scenario" when we do not know the target scenario/distribution in deployment.

The following subsections will describe this type of OOD generalization in more detail.

2.15.1   Formulation of a General Adversarial Environment

Before discussing adversarial generalization, let us first introduce the notion of a *devil*.

> **Definition 2.36: Devil**
>
> The devil is a (known or unknown) adversary that actively tries to find the worst environment for us from the strategy space according to the adversarial goal and knowledge. The more knowledge it has, the worse environments it can specify for us.

A general adversarial environment is specified by three parts: the *adversarial goal*, the *strategy space*, and the *knowledge*.[15] The exact definitions of these parts are given below.

---

[15] If one of them is left unspecified, we are missing critical ingredients.

> **Definition 2.37: Adversarial Goal**
>
> The adversarial goal is a key component of an adversarial setting that specifies which environment is considered "worst" for our model.

> **Definition 2.38: Strategy Space**
>
> The strategy space in an adversarial setting defines the space of possible environments the devil can choose from.

> **Definition 2.39: Knowledge**
>
> The knowledge of the devil in an adversarial setting specifies the devil's ability to pick the worst environment for our model. In short, it defines what the devil knows about the model.

In the next sections, we will discuss how exactly the devil can achieve their goals.

2.15.2 Fast Gradient Sign Method (FGSM)

First, we start with the definition of a *white-box attack*, as the Fast Gradient Sign Method (FGSM) falls into this category.

> **Definition 2.40: White-Box Attack**
>
> When the adversary knows the model architecture and the weights, we call the attack a white-box attack.

If the devil does not want to think too much, then FGSM can be a popular first choice, as it is one of the simplest ways to achieve adversarial goals. The FGSM attack, introduced in "Explaining and Harnessing Adversarial Examples" [63] is a type of L_∞ adversarial attack. Its three ingredients are listed below.

- **Adversarial Goal**: Reducing classification accuracy while being imperceptible to humans.
- **Strategy Space**: For every sample, the adversary may add a perturbation dx with norm $\|dx\|_\infty \leq \epsilon$ (to make sure it is imperceptible).
- **Knowledge**: Access to the model architecture, weights, and thus gradients. (White-box attack.)

The iconic image from [63], shown in Figure 2.47, depicts the attack, where a small perturbation applied completely destroys the model's prediction performance ("gibbon" with 99.3% confidence). A more general informal illustration of this scenario is given in Figure 2.48.

$$+ .007 \times$$ $$=$$

x

"panda"
57.7% confidence

$\mathrm{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, y))$

"nematode"
8.2% confidence

$\boldsymbol{x} +$
$\epsilon \mathrm{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, y))$
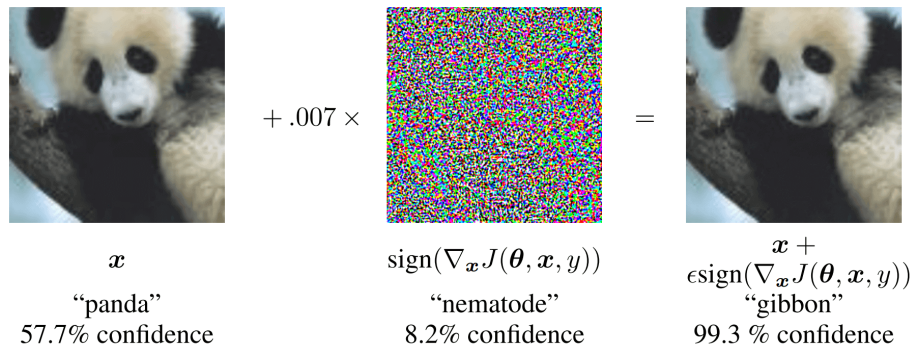
"gibbon"
99.3 % confidence

Figure 2.47: Demonstration of FGSM, taken from [63]. By adding some noise of small magnitude, the network very confidently predicts an incorrect class, destroying the performance of the model. J is the cost function (loss) we wish to *maximize*.
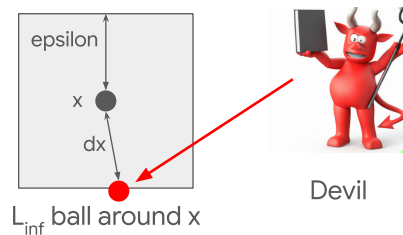

Figure 2.48: Informal illustration of the FGSM method's strategy space. The devil aims to find an adversarial sample in the L_∞ ϵ-ball around the original input x.

FGSM Method

The FGSM attack perturbs the image x as

$$x + \epsilon \ \mathrm{sgn}\left(\nabla_x \mathcal{L}(\theta, x, y)\right)$$

where $\mathcal{L}$ is the loss function used for model θ, y is the ground truth label, and the sign function is applied element-wise. One also has to take care about the image staying in the range $[0, 1]^{H \times W \times 3}$ by clipping or normalizing. ϵ is the size of the perturbation, which is determined by the strategy space. It defines the maximal L_∞ norm of the perturbation.

Let us consider the pros and cons of FGSM below.

- **Pros**: The method is very simple. We take a binary map of the gradient of the loss, i.e., the direction in which the loss increases the most around x. The method is also cheap. It only requires one forward and backward pass per sample to create an adversarial perturbation which makes it swift to obtain.
- **Cons**: The method does not give an optimal result. The perturbed image does not necessarily correspond to the worst-case sample in the L_∞ ball (but it generally gives a good adversarial attack still for unprotected networks).

2.15.3 Projected Gradient Descent (PGD)

If the devil wants to do something more sophisticated to succeed, the Projected Gradient Descent (PGD) might be a favorable choice for them. The PGD attack, introduced in the paper "Towards Deep Learning Models Resistant to Adversarial Attacks" [133], is a type of L_p

adversarial attack ($1 \leq p \leq \infty$), which is the strongest white-box attack to date (also because not many people are looking into strong attacks anymore). The three ingredients of it are detailed below.

- **Adversarial Goal**: Same as for FGSM.
- **Strategy Space**: For every sample, the adversary may add a perturbation dx with norm $\|dx\|_p \leq \epsilon$.
- **Knowledge**: Same as for FGSM. The devil knows everything about the model, both structural details and the weights. It tries to generate a critical perturbation direction based on θ.

An illustration depicting this scenario is shown in Figure 2.49. The devil is *trying to find* the worst-case sample for a fixed x in the L_p ball.
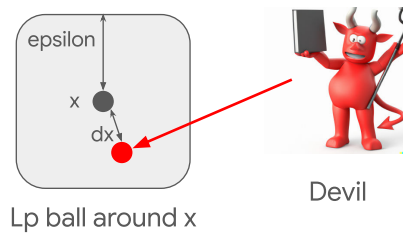

Figure 2.49: Informal illustration of the PGD method's strategy space. The devil aims to find a strong attack in the L_p ball around input x.

PGD Method

The PGD attack solves the optimization problem

$$\max_{dx \in \mathbb{R}^{H \times W \times 3}} \mathcal{L}(f(x + dx), y; \theta)$$

$$\text{s.t. } x + dx \in [0, 1]^{H \times W \times 3}$$

$$\text{and } \|dx\|_p \leq \epsilon.$$

It perturbs the image x iteratively as

$$x^{t+1} = \prod_{x+S} \left(x^t + \alpha \operatorname{sgn} \left(\nabla_x \mathcal{L}(f(x^t), y; \theta) \right) \right)$$

where $\mathcal{L}$ is the loss function used for model θ, y is the ground truth label, t is the iteration index, α is the step size for each iteration, and $\prod_{x+S}$ is the projection on the L_p ϵ-sphere around x.[16] ϵ is the size of the perturbation, which is determined by the strategy space. It defines the maximal L_p norm of the perturbation.

Being an iterative algorithm, PGD usually finds an even worse-case sample than FGSM (which only performs a single step). We iteratively follow the sign of the gradient with step size α and project back onto the L_p ϵ-ball around x. According to the properties of the sign function, in each step, we go in an angle of $\beta \in \{\pm 45°, \pm 90°, \pm 135°, 0°, 180°\}$ from the previous x^t before projecting back onto the ϵ-ball.[17] (Usually, in visualization, this means traveling along

---

[16] $S = \{y \in \mathbb{R}^{H \times W \times 3} \,|\, \|y\|_p \leq \epsilon\}$.
[17] The projection *can* change this angle.

the boundary of the L_p ball.) Now we can go out of the L_p ball of ϵ even in a single step (especially around the 'corners' of the L_p ball), depending on how we choose α. Like in FGSM, we also take care of the image staying in the range $[0, 1]^{H \times W \times 3}$ using clipping or normalization. Convergence happens when, e.g., $\|x^{t+1} - x^t\|_2 \leq 1\mathrm{e}{-5}$ or some similar criterion is satisfied.

> **Additional Information: Using the Gradient's Sign**
>
> Why do we use the sign of the gradient in these methods and not the magnitude? Either case works. However, e.g., Adam [101] is also taking the sign of the gradient for updates (considering the formula without the exponential moving average) and is one of the SotA methods. In high dimensions, the choice of taking the sign does not matter much. This is usually a choice we make based on empirical observations.

2.15.4 FGSM vs. PGD

Let us briefly compare the two attacks we have seen so far, FGSM and PGD. In both cases, the optimization problem for the adversary is non-convex, as the loss surface is non-convex in x. We also have no guarantee for the globally optimal solution, even within a small ϵ-ball (which is very tiny in a high-dimensional space). The strength of the attack depends a lot on the optimization algorithm. We have many design choices, and not all Gradient Descent (GD) variants perform similarly. **PGD is generally much stronger than FGSM; it finds better local optima.** FGSM does not even find local optima in general, as it consists of just a single gradient step. PGD is generally a SotA white-box attack even as of 2023.

> **Additional Information: Size of the ϵ-ball in High-Dimensional Spaces and Distribution of Volume**
>
> Why is the ϵ-ball tiny in a high-dimensional space for a small value of ϵ? The volume of a ball with radius r in $\mathbb{R}^d$ is
> $$V_d = \frac{\pi^{d/2}}{\Gamma\left(1 + \frac{d}{2}\right)} r^d$$
> and $\Gamma(n) = (n-1)!$ for a positive integer n. Therefore, the denominator increases much faster than the numerator, driving the volume to 0 as $d \to \infty$.
>
> The volume is thus concentrated near the surface in high-dimensional spaces: For a fixed dimension d, the fraction of the volume of a smaller ball with radius $r < 1$ inside a unit ball is r^d (as the scalar multiplier cancels). For $r \approx 1$ but d very large, this is around 0.

> **Additional Information: How to choose ϵ in L_p attacks?**
>
> The hyperparameter ϵ is usually chosen to be very small. Even more importantly, one should fix it across studies, as we typically wish to compare against previous attacks/defenses. There are unified values in the community but the exact value does not matter

that much, as below a certain threshold, the perturbations are (mostly) not visible to humans anyway.

2.15.5 Different Strategy Spaces for Adversarial Attacks

Figure 2.50: Example of two image pairs where humans would choose the *left* pair as more similar, but regarding L_2 distances, the *right* pair is much closer. This is because of the translation in the first image pair. Figure-snippet taken from [221].

So far, we have considered perturbations inside an L_p ball determined by ϵ. The problem with this strategy space is that it is not aligned with human perception – it is in the pixel space. It is missing some perturbations that are not visible to humans (such as shifting all pixels up by one), but it also captures some changes apparent to humans (such as additive noise at initially very clear and homogeneous surfaces in images). The L_p strategy space is thus not well-aligned with the adversarial goal. Sometimes it does not satisfy the goal (as the adversary's goal is to produce imperceptible perturbations), and sometimes it technically satisfies the goal but could do it even better (as the adversary's goal is usually also to decrease accuracy as much as possible). The misalignment of additive perturbations and the adversarial goal is further illustrated in Figure 2.50. According to this observation, in the following subsections, we will consider strategy spaces that are different from the L_p-ball-based ones.

Flow-Based Perturbations

In general, images closer in perception space (ones that look more similar to humans) can have a larger L_p difference than obviously different image pairs. Suppose that we have a robust model against any L_p perturbations w.r.t. a ball parameterized by a small ϵ. In this case, the adversary could still be able to find a one-pixel shift of the image that destroys the model's predictions, even though this perturbation is imperceptible. This is a "blind spot" of an adversary that uses an L_p-ball-based strategy space.

> **Definition 2.41: Total Variation**
>
> The total variation of a vector field $f \colon \mathbb{R}^2 \to \mathbb{R}^2$ is defined as
>
> $$\|f\|_{\mathrm{TV}} = \int \|\nabla f_1(x)\|_2 + \|\nabla f_2(x)\|_2 \, dx =: \int \|\nabla f(x)\|_2 \, dx.$$
>
> It is often considered a generalization of the L_2 (or L_1) norm of the gradient to an entire vector space.

Such small image translations (shifts) generally result in huge L_2 distances (as images are not smooth, and along the object boundaries, we have a significant pixel distance), but correspond to perceptually minor differences. Luckily, we can define a metric that assigns small distances to small *per-pixel* image translations. We consider *optical flow* transformations, and we measure small changes by the *total variation (TV)* norm, for which translations of an image have a "size" of zero. Such attacks are discussed in detail in Section 2.15.6.

Physical Attacks

The plausibility of the previously discussed strategy spaces is questionable. L_p attacks and other attacks (e.g., flow-based ones) alter the *digital image*. Do such adversaries exist in the real world? Are the previous attacks plausible at all? Basic security technology can already prevent such adversaries, with access as depicted in Figure 2.51. Therefore, looking into other strategy spaces is well-motivated.
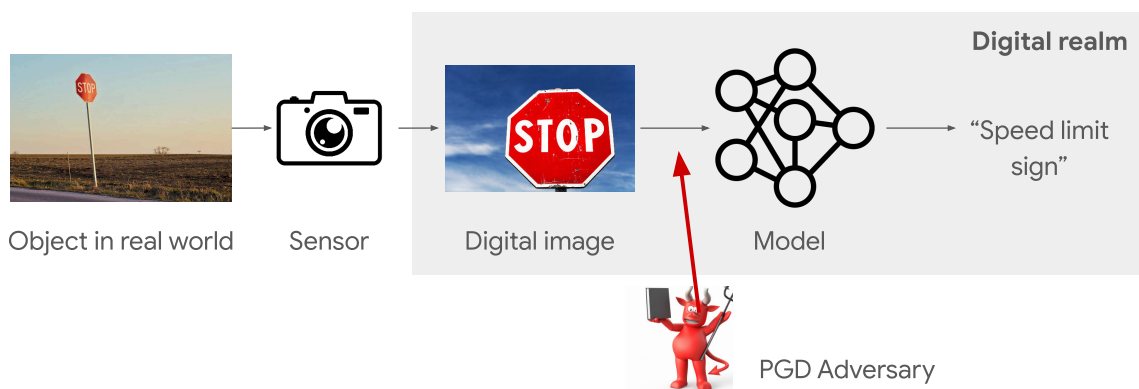

Figure 2.51: The PGD adversary, being a white-box attack, has access to the model in the digital realm.

> **Definition 2.42: Black-Box Attack**
>
> When the adversary only observes the inputs and outputs of a model and does not know the model architecture and the weights, we call the attack a black-box attack.

The PGD adversary perturbs pixels of a digital image after it is captured in the real world. Does this scenario make sense? Should we even defend against such an adversary? We just have to ensure that no one gets to see our compiled code and that no one can change the data stream. This is basic information security. Furthermore, even if one gains access to the data stream, one also needs access to the exact model for white-box attacks. (Once the adversary is that deep in, they might as well just change the prediction directly...) For black-box attacks in the digital realm, this is not needed, but it is still a strange scenario where one has access to the data stream but not the model output. We even go one step further in the discussion about plausibility: When we use an API and have no access to internal data streams, we can indeed construct black-box attacks for the model (as we will see in Sections 2.15.9 and 2.15.10). However, this only ruins the accuracy for us, which seems to be a very poor adversarial goal. **By focusing on attacks in the digital realm, we are probably looking at a non-existent problem.**

Another very similar scenario in the digital realm is when images are uploaded to the cloud, as shown in Figure 2.52. It is very unrealistic for an adversary to come into this pipeline and make changes.
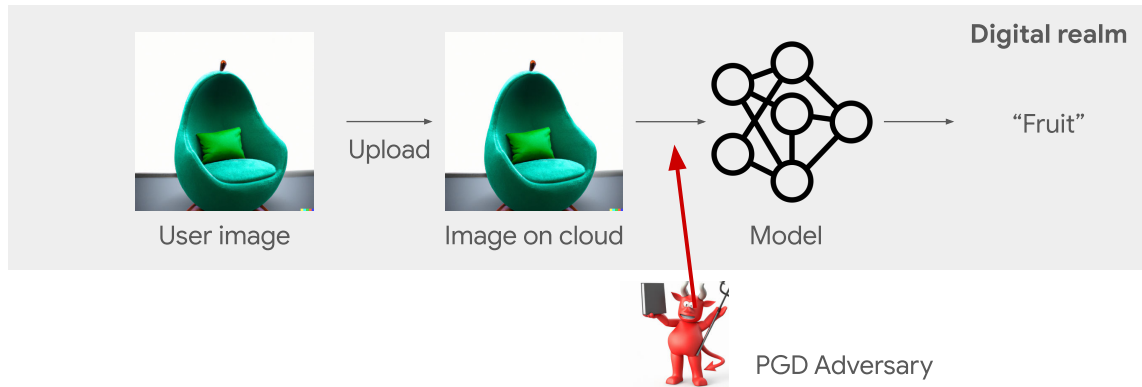

Figure 2.52: In a cloud setting, the PGD adversary still acts in the digital realm.

This lack of realism in attacks in the digital realm inspires the search for a new strategy space in the real world: Let us discuss *physical attacks*. In contrast to previously mentioned attacks, they induce physical changes in objects in the real world. These involve, e.g., putting a carefully constructed sticker (or graffiti) on stop signs to make sure that self-driving cars do not detect it or printing a pattern on cardboard (and e.g. wearing it around the neck) such that the person carrying the sign does not get detected. These options are illustrated in Figure 2.53. This is much more realistic, as the adversary intervenes in the real world, not in a secure stage in a pipeline. The adversaries usually *do* have the necessary access to real-world objects. We argue that we should instead be focusing on defending against such attacks, shown in Figure 2.54. **Note**: Such attacks can be both black-box and white-box attacks.

Figure 2.53: Physical attacks are more realistic than those in the digital realm [195, 134].

Object Poses in the 3D World

We briefly discuss an interesting boundary between adversarial robustness and OOD generalization that also introduces a new strategy space. This is the paper "Strike (with) a Pose: Neural Networks Are Easily Fooled by Strange Poses of Familiar Objects" [6] which focuses on changing poses of objects in 3D space (which is similar to physical attacks but can also be done digitally given a sophisticated image synthesis tool). A collage of synthetic and real images the authors considered is shown in Figure 2.55.

The three ingredients of this "attack" are as follows.

- **Adversarial Goal**: Reducing classification accuracy by changing object poses.
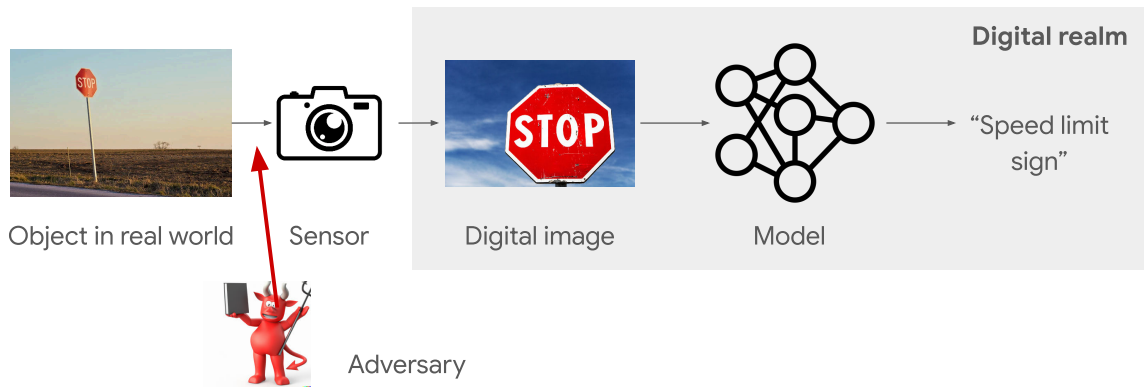
Figure 2.54: In a physical adversarial setting, the adversary has access to the object in the real world. The adversary might also know the internals of the model (considering a white-box setting), but still only intervene in the physical world.

Figure 2.55: Collage of synthetic and real images with the model's corresponding max-probability predictions. According to the human eye, images in (row, column) positions (1, 4), (2, 2), (2, 4), (4, 1), (4, 4) are quite plausible.

- **Strategy Space**: For every sample, the adversary may arbitrarily change the object poses.
- **Knowledge**: Same as for FGSM. (White-box attack.)

Here, the adversary does not necessarily care about small changes in the object pose. Larger changes can still be plausible for the human eye. Once it becomes obvious to humans, they can, of course, intervene. The devil knows everything about the model, both structural details and the weights. It tries to generate a critical pose perturbation based on weights θ.

One may ask how this is a real threat at all. The threatening observation is that the model completely breaks down for the plausible examples, even though these could be observed in the real world. This work is on the boundary of adversarial robustness and OOD generaliza-

tion to real-world domains.[18] If the perturbation grows larger and we do not have a notion of a devil and worst-case samples anymore, we enter the realm of generalization to plausible real-world domains, across biases, or in other OOD generalization schemes.

2.15.6 Optical Flow

Let us now discuss the optical flow approach in more detail. Optical flow is used a lot for visual tracking and videos. It provides the smallest warping of the underlying image mesh to transform image x_1 into x_2.[19] It specifies the *apparent* movement of pixels which is needed to transform image x_1 into x_2. We obtain it by performing (regularized) pixel matching between images/frames.

Optical flow is represented as a vector field over the 2D image plane. Each point of the 2D pixel plane corresponds to a 2D vector. Hence, the size of the warping may be readily computed via total variation (TV) (Definition 2.41). This vector field is usually encoded by colors for visualization. The pixel intensity gives the 2D vector magnitude at the pixel, and the pixel color specifies the 2D vector direction at the pixel.

Example: Consider a ball flying across the sky. The ball pixels are translated across the frames by a tiny bit, but the L_2 distance between the frames is large. Our task is to find *pixel correspondences* between the two frames based on apparent motion. We set up a vector going from pixel (i, j) in frame t to the corresponding pixel (i', j') in frame $(t + 1)$. For example, if $(i, j) = (4, 5), (i', j') = (7, 2)$, then the forward flow is $(u, v) = (3, -3)$. We measure the distance between pixels by taking the L_2 norm of this vector and taking the average of these distances for every pixel. This is precisely what we do when calculating TV. This gives an idea of how much warping has taken place between the two frames. A small flow, however, can also correspond to human-perceptible changes: a small ball flying fast between two frames on a huge, otherwise static image will have a low TV value, but humans are able to point out the differences quickly. Nevertheless, the perturbed images are still deemed plausible by human inspection.

2.15.7 Adversarial Flow-Based Perturbation

How can we use optical flow to find adversarial patterns? Instead of estimating the flow between 2 consecutive frames, we *generate* a flow with a small total variation that fools our model, as done in the paper "Spatially Transformed Adversarial Examples" [215]. The three ingredients of their method are:

- **Adversarial Goal**: Reducing classification accuracy while being imperceptible to humans.
- **Strategy Space**: For every sample, the adversary may choose a flow f with $\|f\|_{\mathrm{TV}} \leq \epsilon$.
- **Knowledge**: Same as for FGSM. (White-box attack.)

This perturbation method is better aligned with human perception (i.e., it is a good proxy for it). It finds pixel-wise movement instead of additive perturbation. The adversary warps the underlying image mesh of image x according to f such that the classification result is

---

[18]This is a boundary between *intentions*. The method can be used to construct adversarial examples that also correspond to plausible OOD domains we might wish to generalize to.

[19]Warping refers to the pixel-wise displacements between the two image meshes.

wrong. If the vector field is aligned in the same direction (constant map), there is no total variation. On the contrary, if the vector field comprises vectors with large magnitudes that are closely spaced and point in different directions, it results in a large TV norm. These abrupt changes in nearby vectors correspond to steep gradients in the field. This is why penalizing total variation encourages images to be *smoother*.
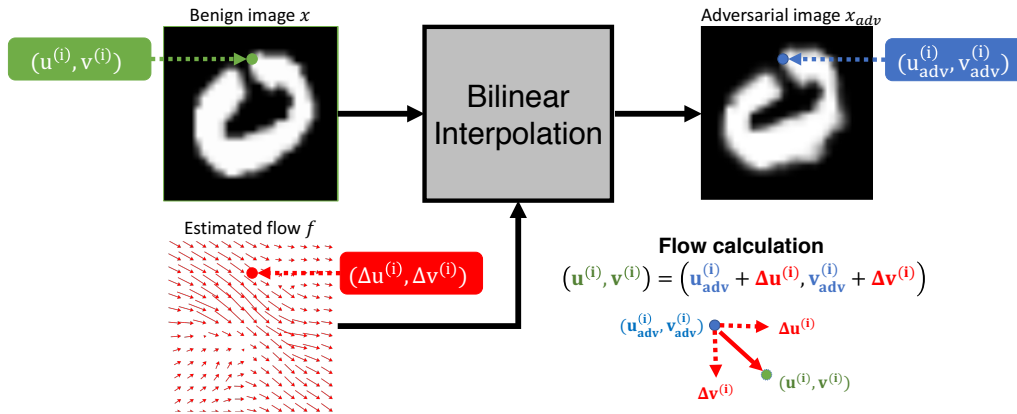

Figure 2.56: Overview of a flow-based adversarial attack using bilinear interpolation to obtain its final adversarial image from the *backward* flow, taken from [215]. See information 2.15.7 for details.

How can we obtain the final adversarial image from the adversarial flow? Figure 2.56 shows a possible way using the *backward flow* and *bilinear interpolation*.

> **Additional Information: Interpolation between source and adversarial images in flow-based adversarial attack**
>
> During the adversarial attack in Figure 2.56, the image x is fixed. The devil comes up with a *backward* optical flow that takes the target pixels to the original pixels. The reason to predict backward flow instead of forward flow is easier bilinear interpolation. When the backward flow is available, each pixel of the adversarial image can be computed after querying some known pixel of the original image (source). On the contrary, using the forward flow to obtain the adversarial example would result in "holes" in the image. The actual *magnitude* of the warps does not matter when calculating the TV norm, so translations of any kind are allowed. At borders, we might copy the pixels of the original image.

2.15.8   White-Box vs. Black-Box Attacks

So far, we have discussed white-box (Definition 2.40) and black-box (Definition 2.42) attacks. Let us discuss some pros and cons of these paradigms.

White-box attacks are powerful. The adversary can obtain the input gradients from the model. (Examples: FGSM, PGD, Flow-Based Perturbation.) White-box attacks are, however, not so realistic. For an ML model on the cloud/as an API, we are never allowed to look into the details of the model. It is intellectual property, and exposing it would make the model

vulnerable to various attacks. The quick solution that most companies follow is to not open source their model. *Black-box attacks are much weaker than white-box attacks but also much more realistic.*

Many real-world applications are based on API access. There are also further limitations to a realistic scenario:

- The number of queries within a time window is limited (rate limit).
- Malicious query inputs are possibly blocked.

For example, consider a face model recognizing the user in a photo album: If we start sending strange patterns like random noise or non-face images, it can easily be detected, and we can be blocked from the service.

Examples of black-box APIs. GPT-3.5/4 [26, 152] produces text output given text input. It is an interesting objective to attack GPT-N based on only input/output observation pairs. One example is Jailbreak Prompts [178]: Here, the adversarial goal is making the model tell us information about immoral or illegal topics; the strategy space of the devil is giving any prompt to the model; and the knowledge of the adversary is the observed answers of the model. The attack is black-box by definition because we do not have access to the model's internal structure. DALL-E [162] produces an image given a textual description.

In the following sections, we will discuss black-box attacks in more detail.

2.15.9   Black-Box Attack via a Substitute Model

> **Definition 2.43: Substitute Model**
>
> A substitute model is a network that is used to mimic a model we wish to attack. Prior knowledge about the attacked model is incorporated into the substitute model, such as the type of architecture, the size of the model, or the optimizer it was trained with.

Gather input-output pairs from the original model (API calls).

Train a substitute model *g* with the input–output pairs.
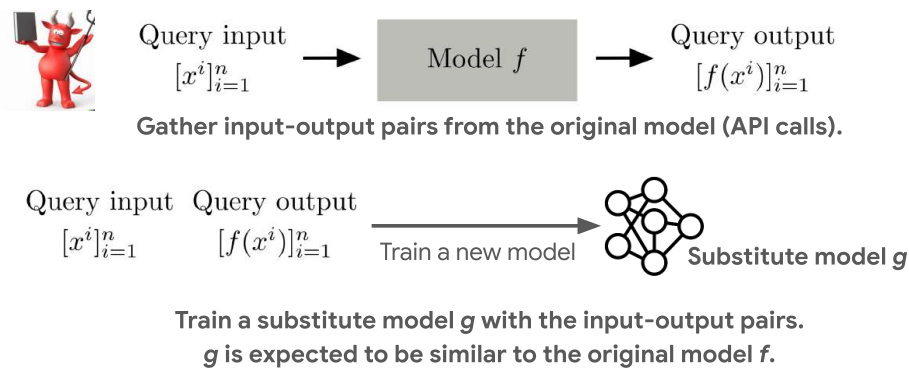*g* is expected to be similar to the original model *f*.

Figure 2.57: Illustration of a method for using substitute models to generate black-box adversarial attacks. We only need query inputs and outputs to train the substitute model.

We will start an overview of the black-box attacks with the seminal work "Practical Black-Box Attacks against Machine Learning" [154]. It introduces the idea of using a substitute model to attack the original model. An overview of the method is given in Figure 2.57. Using this approach, we might need a lot of input-output pairs from the original model, depending on

how complex the model is. Ideally, we want to follow the architecture of the target model. If we, e.g., know that the original model is a Transformer, we should also use one. We then attack the original model by creating adversarial inputs that attack the substitute model g. By using g, the adversary can generate *white-box* attacks. The hope is that this attack also works for f. Based on empirical observations, this can work quite well.

For smaller models, this method might be feasible to attack model f well. However, for larger models, we need tons of data and extreme computational effort to train the substitute model. In particular, only a handful of companies in the world could mimic GPT-3 with a substitute model. It would be easy to trace back who is responsible for the attacks. Usually, such large companies focus on problems other than these black-box attacks. They already have many other problems, e.g., private training data leakage by querying, bias issues, or explainability, that are way more realistic.

2.15.10 Black-Box Attack via a Zeroth-Order Attack

Another type of black-box attack is based on the approximation of the model gradient with a lot of API calls. One way to do this is described in the work "ZOO: Zeroth Order Optimization based Black-box Attacks to Deep Neural Networks without Training Substitute Models" [35]. The idea comes from the fact that one can approximate the gradient of the loss numerically using finite differences, so that for a small enough $h \in \mathbb{R}$:

$$\frac{\partial \mathcal{L}(x)}{\partial x_i} \approx \frac{\mathcal{L}(x + he_i) - \mathcal{L}(x)}{h}$$

where e_i is the ith canonical basis vector. One can also use a more stable symmetric version that gives better approximations in general (but requires more network evaluations):

$$\frac{\partial \mathcal{L}(x)}{\partial x_i} \approx \frac{\mathcal{L}(x + he_i) - \mathcal{L}(x - he_i)}{2h},$$

where $x \in \mathbb{R}^d$ is a flattened image and $\mathcal{L}(x) \in \mathbb{R}$ is the loss function of choice, based on our target class y that we want the model to classify x as.

Example: Consider a 200×200 image. In this case, $x \in \mathbb{R}^{120,000}$ and $\nabla \mathcal{L}(x) \in \mathbb{R}^{120,000}$. We need 120,001 API calls to approximate the gradient *of a single image*, or 240,000 if we consider the symmetric approximation. No API will let us do this in a manageable time. For this to work, we also need access to the logits $z(x)$ or the probabilities $f(x)$ from the model to compute the objective $\mathcal{L}$, not just the predicted class label. For example, we might use the objective $\mathcal{L}(x) = \max\{\max_{i \neq y} z(x)_i - z(x)_y, -\kappa\}$, as given in [35].

Still, the worst case with such black-box attacks for the model owners is that the performance drops. This is not a realistic goal for an adversary, as it only happens to the attacker and only on the adversarial samples they create. If the attacker wants to decrease performance for others, too, they need access to the data stream.

Note: The paper was published in 2017. Back then, these attacks were focused on theoretical possibilities. Nowadays, the field is focusing more on realistic threats we have to tackle. The focus has shifted.

We can even be more imaginative and train a local model that predicts the pixel location

most likely to generate the highest response by the attacked model. In this case, we need a dataset of (image, pixel) pairs where the pixel changes the prediction of many locally available models the most. Leveraging this dataset, we can get the model's pixel prediction and find the pixel's perturbation via API calls that result in the desired behavior (e.g., compute gradients for that pixel using finite differences and update that pixel of the image using the sign of the gradient).

We simplify the previous approach and pick random coordinates to perform stochastic coordinate descent, as shown in [35]:

Algorithm 3: Stochastic Coordinate Descent
---
while *not converged* **do**
    Randomly pick a coordinate $i \in \{1, \ldots, d\}$
    Compute an update dx_i^* by approximately minimizing

$$\underset{dx_i}{\arg\min}\, \mathcal{L}(x + dx_i e_i)$$

    $x_i \leftarrow x_i + dx_i^*$
end

This is, of course, not very efficient. It is better to pick i smartly and perturb that pixel using a few API calls to determine a suitable perturbation.

2.15.11 Defense against Attacks: Adversarial Training

We have discussed many adversarial attacks. Is there any way to defend against them? To answer this, we will touch upon one instructive defense method that gave rise to the research direction of defense methods, called adversarial training. This method was introduced in the paper "Towards Deep Learning Models Resistant to Adversarial Attacks" [133]. It is generally perceived as one of the best-working defenses against L_p attacks.

Adversarial training has a minimax formulation: Optimize θ w.r.t. the worst-case perturbation of x as

$$\min_\theta \mathbb{E}_{(x,y)\in\mathcal{D}} \left[\max_{dx\in S} \mathcal{L}(x + dx, y; \theta) \right]$$

with, e.g., $S = [-\epsilon, \epsilon]^N$ corresponding to an L_∞ attack. In practice, we do a few PGD steps for each x, generate an attack dx, and use that for training θ.

Results of Adversarial Training

We would like to discuss Figure 2.58 from several aspects, as this extensive benchmark sheds light on many interesting tendencies in adversarial generalization.

Impact of adversarial attacks. The impact is very strong for standard training. Any kind of attack pushes accuracy to 0. FGSM is slightly weaker than PGD. PGD gives 0 accuracy for any capacity scale.

Impact of adversarial training. (b) If we train on FGSM samples, we become very robust to FGSM-attacked test inputs. The model accuracy for natural images decreases compared
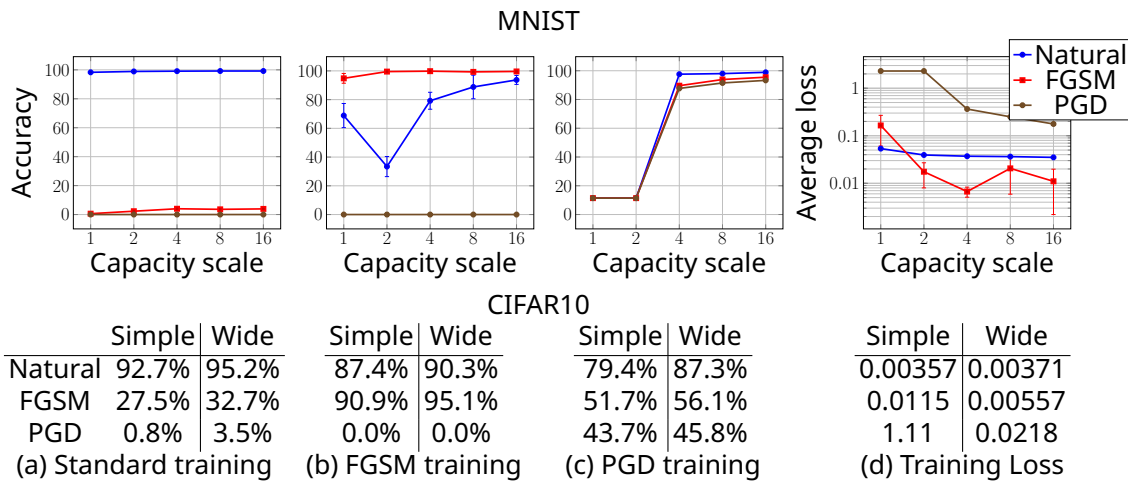
Figure 2.58: Results of adversarial training considering the FGSM and PGD attacks when attacked during testing. Adversarial training with PGD is very effective, given enough capacity. 'Natural' refers to test images that are not attacked. 'Simple' and 'wide' refer to the architecture used. A detailed analysis is given in the text. Figure adapted from [133].

to standard training, especially for smaller capacities. There is always a trade-off, even if we mix attacked and natural images during training (which is not done here): adversarial training leads to a drop in natural accuracy, as we cannot be perfect at both. The method breaks down completely for PGD-attacked inputs that are very strong. (c) If we train on PGD samples, we get fairly nice adversarial accuracy. The method seems to work but only for high-capacity models. For low-capacity models, PGD training is useless. The higher the capacity, the more effective PGD training is. The general lesson of adversarial training: We really need a high-capacity model.

Impact of model capacity. For adversarial training, we need to be able to take care of much more complexity in the data distribution, which results in very complex decision boundaries. We must defend against attacks while learning the task itself, which is quite challenging.[20] The stronger the attack, the more apparent the need for a high-capacity model is.

Transferability of attacks/defenses. In general, to be able to defend against a particular attack, we need to train against it. Otherwise, the method often completely breaks down. PGD transfers very well to natural and FGSM samples, but only for large-capacity models. (a): most accurate on natural images. (b): most accurate on FGSM images. (c): most accurate on natural (!) images.

Computational Complexity

Adversarial training is perceived as a costly defense mechanism. Let us highlight where the main overhead is coming from and discuss complexities from various viewpoints.

Per-batch complexity. We define per-batch complexity as the number of forward and backward steps per batch. In vanilla training, we have one forward step and one backward step per batch. In adversarial training, we need T forward and backward steps for finding the adversarial PGD samples (batched), and one (forward, backward) pair for updating θ. Therefore,

---

[20]In *curriculum learning*, we are doing the exact opposite, which can help stabilize training.

in total, we need $T + 1$ forward and backward steps. (For FGSM, we naturally have $T = 1$, but it does not work well for adversarial training.) *For a single batch, adversarial training is $T + 1$ times as expensive as vanilla training.*

Training complexity. Similarly to per-batch complexity, we define training complexity as the number of forward and backward steps *per training*. As adversarial training is $T + 1$ times as expensive as vanilla training for a single batch, the entire training is also *at least $T + 1$ times* more expensive. We might even need more epochs because we are solving a more complex problem. The apparent complexity is $T + 1$ times as much as for vanilla training. The added complexity can arise from having to add more epochs.[21] We generally control the complexity by varying the number of iterations T per attack.

Inference complexity. Inference complexity refers to the number of forward and backward steps per inference. We need one forward step and one backward step for both the vanilla-trained model and the adversarially trained model. After adversarial training, we use our model as usual. We can, e.g., open source it and let people use it without any overhead.

> **Additional Information: Data Augmentation vs. Adversarial Training**
>
> Is data augmentation considered adversarial training? Augmentation methods do not know anything about the model. The model can learn to be *invariant* to flipping and other perturbations with data augmentation, but these are not considered adversarial attacks. Adversarial attacks are much stronger than simple data augmentation methods. Adversarial training also does not make the model invariant to many (possibly large-scale) perturbations, only to the adversarial attacks.

2.15.12 Obfuscated Gradients: Breaking the defense Again!

As we have seen in the previous sections, people came up with adversarial attacks and others made defenses against them. Of course, this was not a one-round process. In the field of adversarial robustness, there is a constant loop between attacks and defenses trying to improve upon each other. Let us consider attacks against types of previously popular defenses such as *obfuscated gradients*.

> **Definition 2.44: Obfuscated Gradients**
>
> Obfuscated gradient methods are a class of defense mechanisms designed so that the constructed defense necessarily causes gradient masking. We consider three types of obfuscated gradient methods [30]:
> - **Methods shattering gradients**: Gradients are shattered when a defense makes the computation graph non-differentiable, introduces numeric instability, or otherwise causes a gradient to be non-existent or incorrect.
> - **Methods using stochastic gradients**: Stochastic gradients are caused by random-

---

[21] If we also take into consideration that we might need a model of much higher capacity, the relative cost of adversarial training is even more extreme.

> ized defenses, where either the network itself is randomized, or the input is randomly transformed before being fed to the classifier, causing the gradients to become randomized.
> • **Exploding and vanishing gradients**: These are often caused by defenses that consist of multiple iterations of neural network evaluation, feeding the output of one computation as the input of the next. This type of computation, when unrolled, can be viewed as an extremely deep neural network evaluation, which can cause vanishing/exploding gradients.

2018 was a loud year for adversarial attacks and defenses. In particular, ICLR'18 was teeming with defense methods. The paper "Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples" [13] from ICML'18 (which was three months later than ICLR'18) claimed:

> "7 of 9 ICLR'18 defenses do not work."

As shown in Table 2.8, after the authors' proposed attack, they mostly reached a very low

Table 2.8: Results of ICLR'18 defenses against the ICML'18 attack [13] taken from the paper. "Defenses denoted with * propose combining adversarial training; we report here the defense alone [...]. The fundamental principle behind the defense denoted with ** has 0% accuracy; in practice, imperfections cause the theoretically optimal attack to fail [...]."

Defense	Dataset	Distance	Accuracy
Buckman et al.	CIFAR	$0.031\ (\ell_\infty)$	0%*
Ma et al.	CIFAR	$0.031\ (\ell_\infty)$	5%
Guo et al.	ImageNet	$0.005\ (\ell_2)$	0%*
Dhillon et al.	CIFAR	$0.031\ (\ell_\infty)$	0%
Xie et al.	ImageNet	$0.031\ (\ell_\infty)$	0%*
Song et al.	CIFAR	$0.031\ (\ell_\infty)$	9%*
Samangouei et al.	MNIST	$0.005\ (\ell_2)$	55%**
Madry et al.	CIFAR	$0.031\ (\ell_\infty)$	47%
Na et al.	CIFAR	$0.015\ (\ell_\infty)$	15%

(0 - 9%) accuracy even when using the ICLR'18 defenses. Why were these obfuscated gradient methods so ineffective? These defenses are specifically targeted against gradient-based attacks. Obfuscated gradient methods break down the gradients (i.e., they make it malfunction). This misleads gradient-based attacks. The problem is that the model itself is *still vulnerable* to most of these attacks using a simple trick. One can use slight modifications of gradient-based attacks to attack such defenses again, usually with great success.

Breaking down gradients does not give us any guarantees.

Even if the previously mentioned methods worked in making gradient-based adversarial attacks impossible, the model being safe is not equivalent to no gradient-based algorithm being able to find an attack. There can still be some adversarial image within the L_p ball (or neighborhood in general). When we break down the gradients, PGD cannot attack the image in the right way *directly*. Using PGD naively results in a benign image, i.e., the network can still recognize it well.

The model is safe when there is absolutely no adversarial sample within the attack space. If this is not guaranteed, there can still be some algorithm that can find the working attack.

How can we make the gradients malfunction?

One way to make the gradients malfunction is to transform the inputs before feeding them to the DNN. These are called input-transformation-based defenses. They apply image transformations (and possibly random combinations thereof) to the original input image. The idea is that if we just transform our image in different ways using a discrete set of transformations, that does not change the content of the image much, and if we have many variations of possible transformations (of which we select one at test time), that is supposed to be very effective against adversarial attacks. This is because adversarial attacks are minimal changes in the image, and if we are killing these small changes using transformations, the attack will probably not harm the model anymore. We want to remove adversarial effects from the input image before feeding the result to the DNN. As we will soon see, this intuitive reasoning is *flawed* in most cases, as input-transformation-based defenses only work when considering chained random transformations with a combinatorial scaling of possibilities.
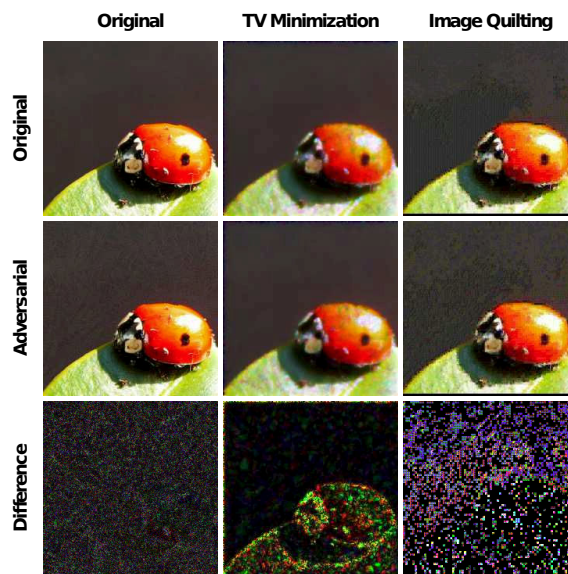
Examples for Input Transformations

Figure 2.59: Example input transformations that can be used for gradient obfuscation. Figure taken from [68].

Several examples of input transformations are shown in Figure 2.59 that are detailed below.

Cropping and rescaling of the original image. We crop the part that contains the gist of what is going on in the image or rescale the image to the input size of the network.

> **Definition 2.45: Bit Depth**
>
> The bit depth of an image refers to the number of colors a single pixel can represent. An 8-bit image can only contain $2^8 = 256$ unique colors. A 24-bit image can contain $2^{24} = 16,777,216$ unique colors.

Bit depth reduction. Reducing the bit depth kills some information, but by doing this denoising (from the perturbation's viewpoint), we can also remove critical adversarial perturbations.

JPEG encoding and decoding. JPEG uses Discrete Cosine Transform (DCT). This is a typical transformation included in image viewers – a natural way to defend against perturbations.

Removing random pixels and inpainting them The inpainting can be done, e.g., via TV (Definition 2.41) minimization. When removing a boundary region, such inpainting will not result in a constant region (having the average value of the neighboring pixels) but rather a very smoothed version of the original image. (The boundary will be followed to some extent.)

Image quilting. This method reconstructs images using small patches from *other images* in a database. The used patches are chosen to be similar to the original patches. These are also usually tiny. Before feeding it to the network, we replace the original image with the reconstruction.

Results of naive FGSM, DeepFool [143], and Carlini-Wagner [31] after input transformations

To see whether the input transformation defense works, we take a look at the results of FGSM, DeepFool, and the Carlini-Wagner method in Figure 2.60. The general message of these results is that applying the previously listed input transformation to an image protects it against gradient-based adversarial attacks. We will see that this is an *incorrect conclusion*.

Straight-Through Gradient Estimator

One of the reasons why previous methods using input transformations *still fail to defend our networks* is the fact that we can still "approximate" the gradient of the defended model by using a straight-through gradient estimator.

> **Definition 2.46: Straight-Through Gradient Estimator**
>
> The straight-through estimator generates gradients for a non-differentiable transformation as if the forward pass were the identity transformation; i.e., it lets the gradient flow through in the computational graph.

A successful application of the straight-through estimator is attacking JPEG encoding/decoding defenses. The *forward* pass is JPEG encoding and decoding, which is non-differentiable (because of quantization) but close to an identity mapping. In the *backward* pass, we compute the gradient as if the forward were the identity mapping. The fact that this is a successful
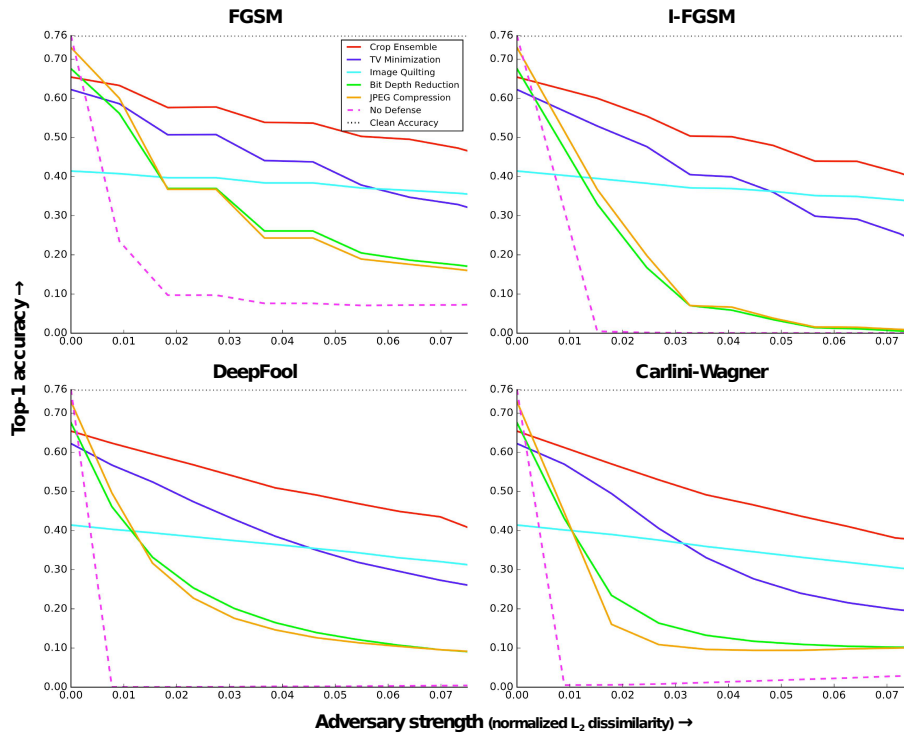
Figure 2.60: Top-1 classification accuracy of ResNet-50 on adversarial samples of various kinds. If we use no input transformations, the model's predictions break down completely. If we use the transformations listed in text *individually*, the methods start failing. The stronger the adversary (i.e., the more L_2 dissimilarity we allow), the better the attack methods do, but they still perform quite poorly. Figure taken from [68].

application of the estimator for an attack shows that this transformation only helped for gradient obfuscation because it made the computations non-differentiable. A Python example of a JPEG transformer is shown in Listing 2.1.

The problem with naive gradient obfuscation methods

When we attack models employing gradient obfuscation methods detailed above as a white box, we *also have access to the transformations*.[22] First, assume that there is a single deterministic transformation. We do not have to know what this transformation precisely is; we just need access to it.

Cropping and rescaling. This is a differentiable transformation (cropping is just indexing, which is differentiable; rescaling is linear), therefore, we can attack the joint network, i.e., the entire pipeline. The defense does not work at all – we can generate successful attacks again. This is depicted in Figure 2.61.

Other discrete transformations. For example, consider JPEG encoding and decoding. Such transformations are not differentiable. However, we can still "differentiate through" quantization layers, using the *straight-through gradient estimator* (Definition 2.46). We can generate successful attacks again, as depicted in Figure 2.62.

Mixture of random transformations. Now, assume that there are multiple transformations,

---

[22]Note, however, that BaRT [160] (2.15.14) works because it has such a large stochasticity internally.

```python
class JPEGTransformer(nn.Module):
    def forward(self, x):
        """JPEG encoding and decoding."""
        encoded_x = self.jpeg_encode(x)
        transformed_x = self.jpeg_decode(encoded_x)
        return transformed_x

    def backward(self, x, dy):
        """Straight-through estimator.
        Computes gradient as if self.forward = lambda x: x.
        """
        return dy
```

Listing 2.1: Python example of a JPEG transformer.

and one of them (or a mixture of them) is chosen randomly. When there is uncertainty in what transformation is used, the white box partially becomes a black box, as we do not know what is taking place in the random transformation. Still, for easier cases, the attacker can generate an attack that works for *any* of the transformations (defenses) by performing *Expectation over Transformations* (EoT).[23]

Expectation over Transformations (EoT). One can observe that

$$\nabla \mathbb{E}_{t \sim T} f(t(x)) = \mathbb{E}_{t \sim T} \nabla f(t(x)),$$

as the gradient and integral can be exchanged when a function is sufficiently smooth, which DNNs are. (For discrete transformations, we use the straight-through estimator anyway, which makes them also work.) The formula tells us that to attack the expected output of f w.r.t. $t \sim T$, we take the gradient for each transformation and then take the expectation w.r.t. the transformations. This procedure can be trivially Monte Carlo estimated. We update the input w.r.t. the expected gradient's approximation iteratively. Python code for a simple EoT attack is given in Listing 2.2.

With sufficient capacity for the attacker, the defense can become ineffective. This is pushing the limit of the capacity of the attacker. If the attacker has full capacity to address many possibilities for transformations at test time, we attack all of them simultaneously. The ICML'18 attack applies all the techniques we mentioned before. It destroys the defense that uses random transformations and makes the network have 0% adversarial accuracy.

2.15.13 Effectiveness of Adversarial Training

Adversarial training (AT) does not introduce obfuscated gradients. It was hard for the ICML'18 method to attack adversarially trained models with greater attack success rates. AT is, therefore, an effective defense. Notably, the authors of [13] use vanilla adversarial training without EoT. Performing EoT additionally would increase computation costs but would likely result in even stronger defenses.

Note: Even after adversarial training, there might still be some adversarial samples within

---

[23]Naive iterative gradient-based optimization would not work, as the gradients of the individual random transformations are simply too noisy.
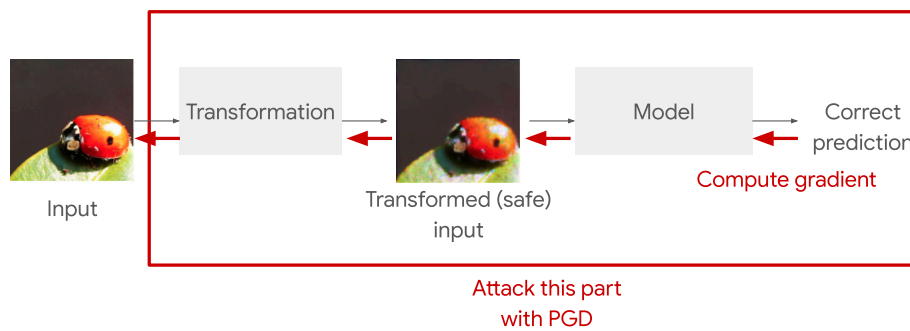
Figure 2.61: An easy way to circumvent obfuscated gradient defenses when the applied transformations are differentiable.
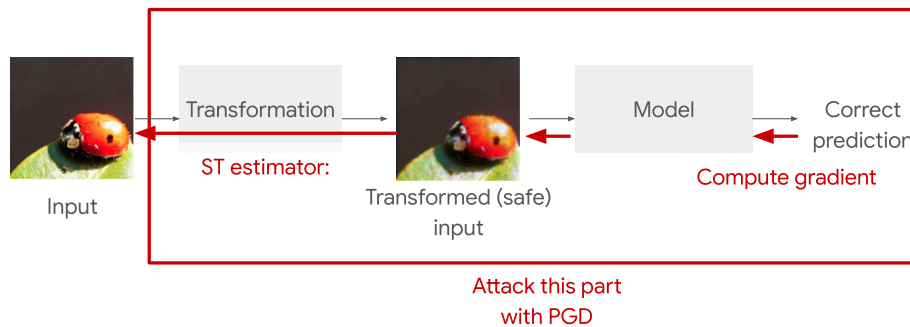

Figure 2.62: Circumventing obfuscated gradient defenses when the applied transformations are *non-differentiable*, using the straight-through gradient estimator.

the L_p ball – we get no guarantees. However, adversarial training is still understood as a solid defense. The critical caveat of AT is that it is complicated to perform at scale. If we are dealing with an ImageNet scale, it is possible but also very impressive. The training time increases notably: adversarial training takes at least $T + 1$ times as long as regular training (Subsection 2.15.11), but here we also have to perform EoT, resulting in a triple `for` loop.

2.15.14 Barrage of Random Transforms (BaRT)

As we have seen, we always have a loop of improvement in adversarial settings between attackers and defenders. Once a defense with a mixture of random transformations is broken (e.g., EoT effectively beats a defender with a reasonable number of candidate transformations), the question naturally arises: What happens when the set of transformations is gigantic on the defense side?

If the defender starts using random combinations of transformations, the number of possibilities grows exponentially as the number of individual transformations and the length of the transformation sequence grows.

The paper "Barrage of Random Transforms for Adversarially Robust Defense" [160] was a "reply" to the EoT paper that introduced an enormous set of possible transformations.

BaRT Method

The method introduces ten groups of possible image transformations listed below.

```python
def generate_eot_attack(x, model, transformation_list, num_samples):
    random_transformations = np.random.choice(transformation_list, num_samples)

    grad_eot = np.zeros_like(x)
    for transformation in random_transformations:
        y = model(transformation(x))
        grad_x = compute_input_gradient(y, x)

        # Approximate expectation by averaging.
        grad_eot += grad_x / num_samples

    return x + grad_eot
```

Listing 2.2: Python pseudocode for generating an EoT attack. This is not PGD or FGSM, but they would work the same way with EoT.

- Color Precision Reduction
- JPEG Noise
- Swirl
- Noise Injection
- FFT Perturbation
- Zoom
- Color Space
- Contrast
- Greyscale
- Denoising

Each group contains some number of transformations. In total, we have 25 transformations, each of which has parameters p that alter their behavior.

The choice of transformations is made as follows.

1. Randomly select k out of n transforms where each transform by itself is randomized.
2. Apply the selected transforms in a random sequence:

$$f(x) = f(t_{\pi(1)}(t_{\pi(2)}(\dots(t_{\pi(k)}(A(x)))\dots))),$$

where A is the adversary.

Selecting the transformations randomly and applying them in a random sequence generates an exponential number of possibilities ($n!/(n-k)!$) that still do not change the semantic meaning of the image. Even after applying all transformations, the model can still recognize the objects pretty well. However, the sheer number of possibilities makes it very hard for the attacker to prepare against all kinds of defenses. It must have a large enough capacity and many samples are required to Monte Carlo sample the expectation. To establish resilience against such input transformations, they are applied both during training and inference. Therefore, this is *not* a post-hoc algorithm.

The method has some overhead in the cost of training, but it boils down to selecting an input transformation sequence with can be done very efficiently on the CPU. The overhead is, therefore, similar to that of data augmentation. One can also influence this overhead by changing

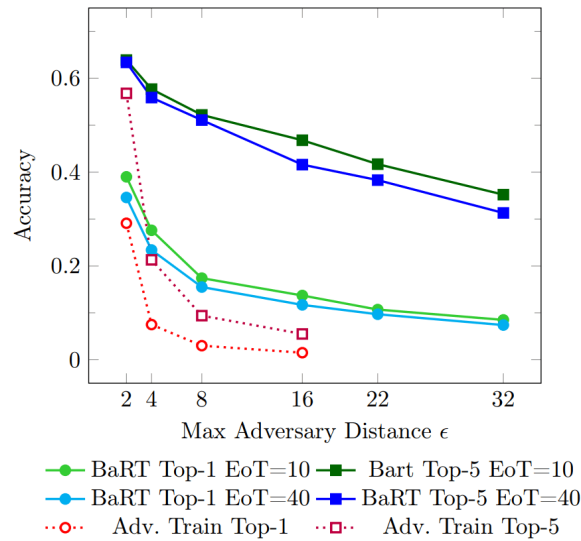how often the transformations are resampled.

BaRT Results

Figure 2.63: BaRT defends a model against PGD (which is not surprising). BaRT also defends a model against the ICML'18 methods with EoT (10 or 40 samples), designed to break gradient obfuscations. Using BaRT, performance does not drop too much by increasing the max adversary distance ϵ. It is even more effective than adversarial training – the attacker cannot push the scores down to 0, not even for $\|x - \hat{x}\|_\infty < 32$. (!) Top-k refers to top-k accuracy. Figure taken from [160].

The results of BaRT are shown in Figure 2.63. The key message here is that *BaRT is one of the SotA adversarial defense methods even in 2023.*

2.15.15   Certified defenses

Let us discuss *certifications of robustness*. Certified defense methods make sure there is *no successful attack* in the strategy space (e.g., the L_p ball) under some assumptions. The "Certified Defenses against Adversarial Examples" [161] paper can give certifications of robustness by considering many simplifying assumptions for the network and the adversarial objective.

The typical chain of thought for certified defenses is to come up with a trainable objective, and then show that solving this trainable objective will ensure that there is no worse attack than a certain type.

The authors consider a binary classification setting and a two-layer neural network where the score is calculated as

$$f(x) = V\sigma(Wx).$$

Here, $V \in \mathbb{R}^{2 \times m}$, $W \in \mathbb{R}^{m \times d}$, and σ is an elementwise non-linearity with bounded gradients to $[0, 1]$, e.g., ReLU or sigmoid. Notably, the authors calculate the score of both positive and negative classes instead of considering a single score for the ease of formalism. A certificate of defense is given by bounding the margin of the incorrect class over the correct one for any adversarial perturbation inside the L_∞ ϵ-ball centered at a particular input x, denoted by $B_\epsilon(x)$. Further details are discussed in Information 2.15.15.

Definition 2.47: Fundamental Theorem of Line Integrals

Consider a parametric curve $r : [a, b] \to \mathbb{R}^d$ and a differentiable function $f : \mathbb{R}^d \to \mathbb{R}$. Then

$$\int_a^b \underbrace{\langle \nabla f(r(t)), r'(t) \rangle \; dt}_{\langle \nabla f(r(t)), dr \rangle, \; dr = r'(t)dt} = f(r(b)) - f(r(a)).$$

In words: The integral of directional derivatives along the curve r of the function f is equal to the difference of boundary values of f. In short, the shape of the curve r does not matter.

Connection to single variable calculus: The fundamental theorem of integrals states that for a differentiable $f : \mathbb{R} \to \mathbb{R}$:

$$\int_a^b f'(x) \; dx = f(b) - f(a).$$

In this case, we have a single possible way from a to b, which is generalized for line integrals.

Additional Information: Formulation of the Certified Defenses Method

The authors consider the following worst-case adversarial attack: $A_{\mathrm{opt}}(x) = \arg\max_{\tilde{x} \in B_\epsilon(x)} \tilde{f}(\tilde{x})$, where

$$\tilde{f}(x) := \underbrace{f^1(x)}_{\text{score of incorrect label}} - \underbrace{f^2(x)}_{\text{score of correct label}}.$$

The attack is successful if $\tilde{f}(A_{\mathrm{opt}}(x)) > 0$ as the incorrect class is predicted.

We derive the following upper bounds on the severity of any adversarial attack $A(x)$:

$$\tilde{f}(A(x)) \overset{(i)}{\leq} \tilde{f}(A_{\mathrm{opt}}(x)) \overset{(ii)}{\leq} \tilde{f}(x) + \epsilon \max_{\tilde{x} \in B_\epsilon(x)} \|\nabla \tilde{f}(\tilde{x})\|_1 \overset{(iii)}{\leq} \tilde{f}_{\mathrm{QP}}(x) \overset{(iv)}{\leq} \tilde{f}_{\mathrm{SDP}}(x).$$

(i) Arises from the optimality of A_{opt}. (ii) leverages the fundamental theorem of line intergrals (Definition 2.47):

$$\tilde{f}(\tilde{x}) = \tilde{f}(x) + \int_0^1 \nabla \tilde{f}(t\tilde{x} + (1-t)x)^\top (\tilde{x} - x)dt$$
$$\leq \tilde{f}(x) + \max_{\tilde{x}' \in B_\epsilon(x)} \epsilon \|\nabla \tilde{f}(\tilde{x})\|_1$$

where $\tilde{x} \in B_\epsilon(x)$ and the inequality holds because the linear interpolation $t\tilde{x} + (1-t)x$ of two elements x and $\tilde{x}$ of $B_\epsilon(x)$ is also an element of $B_\epsilon(x)$ for any $t \in [0, 1]$. In (iii), $\tilde{f}_{\mathrm{QP}}(x)$ denotes the optimal value of a (non-convex) quadratic program. This is a specific bound for two-layer networks where $\tilde{f}(x) = f^1(x) - f^2(x) = v^\top \sigma(Wx)$ with $v := V^1 - V^2$ being the difference of last-layer weights of the correct and incorrect class. In this specific case,

we upper-bound $\tilde{f}(x) + \epsilon \max_{\tilde{x} \in B_\epsilon(x)} \|\nabla \tilde{f}(\tilde{x})\|_1$ by noting that for $\tilde{x} \in B_\epsilon(x)$:

$$\|\nabla \tilde{f}(\tilde{x})\|_1 = \|W^\top \operatorname{diag}(v) \sigma'(W\tilde{x})\|_1 \leq \max_{s \in [0,1]^m} \|W^\top \operatorname{diag}(v) s\|_1 = \max_{s \in [0,1]^m, t \in [-1,1]^d} t^\top W^\top \operatorname{diag}(v) s$$

where the last equality shows a different way to write the L_1 norm. Therefore,

$$\tilde{f}(x) + \epsilon \max_{\tilde{x} \in B_\epsilon(x)} \|\nabla \tilde{f}(\tilde{x})\|_1 \leq \tilde{f}(\tilde{x}) + \epsilon \max_{s \in [0,1]^m, t \in [-1,1]^d} t^\top W^\top \operatorname{diag}(v) s =: \tilde{f}_{\text{QP}}(x).$$

The reason why we do not stop here is that this quadratic program is still a non-convex optimization problem. This is why we turn to (iv), which gives a *convex* semidefinite bound. First, the authors of [161] reparameterize the optimization problem in s as

$$\tilde{f}_{\text{QP}}(x) := \tilde{f}(x) + \epsilon \max_{s \in [-1,1]^m, t \in [-1,1]^d} \frac{1}{2} t^\top W^\top \operatorname{diag}(v)(\mathbf{1} + s)$$

where $\mathbf{1}$ is a vector of ones. Then, one needs to define auxiliary vectors and matrices to obtain the form of a semidefinite program:

$$y := \begin{pmatrix} 1 \\ t \\ s \end{pmatrix}$$

$$M(v, W) := \begin{bmatrix} 0 & 0 & \mathbf{1}^\top W^\top \operatorname{diag}(v) \\ 0 & 0 & W^\top \operatorname{diag}(v) \\ \operatorname{diag}(v)^\top W \mathbf{1} & \operatorname{diag}(v)^\top W & 0 \end{bmatrix}.$$

Now, we rewrite $\tilde{f}_{\text{QP}}(x)$ as

$$\tilde{f}_{\text{QP}}(x) = \tilde{f}(x) + \epsilon \max_{y \in [-1,1]^{(m+d+1)}} \frac{1}{4} y^\top M(v, W) y = \tilde{f}(x) + \frac{\epsilon}{4} \max_{y \in [-1,1]^{(m+d+1)}} \langle M(v, W), yy^\top \rangle.$$

Finally, we note that $\forall y \in [-1,1]^{(m+d+1)}$, $yy^\top$ is a positive semidefinite matrix[a] and the diagonal of $yy^\top$ is a vector of ones. Defining $P = yy^\top$, we obtain the convex semidefinite program

$$\max_{y \in [-1,1]^{(m+d+1)}} \frac{1}{4} \langle M(v, W), yy^\top \rangle \leq \max_{P \succeq 0, \operatorname{diag}(P) \leq 1} \frac{1}{4} \langle M(v, W), P \rangle$$

where the notation $P \succeq 0$ refers to P being positive semidefinite, which allows us to define $\tilde{f}_{\text{SDP}}(x)$ as

$$\tilde{f}_{\text{SDP}}(x) := \tilde{f}(x) + \frac{\epsilon}{4} \max_{P \succeq 0, \operatorname{diag}(P) \leq 1} \langle M(v, W), P \rangle.$$

Notably, the optimization problem in $f_{\text{SDP}}(x)$ is fixed in the neural network weights v and W and does not depend on x. Therefore, obtaining it is very much feasible, as we only need to calculate the input-agnostic upper bound once for each model.

Sadly, our story does not end here. One may assume that the post-hoc application of the above upper bound is enough. While we can indeed calculate such a certificate post-hoc, it might be arbitrarily loose. Regular cross-entropy training encourages $\tilde{f}(x)$ to be large

in magnitude on training samples. However, the term $\frac{\epsilon}{4} \max_{P \succeq 0, \mathrm{diag}(P) \leq 1} \langle M(v, W), P \rangle$ is *not encouraged to be small* to tighten the bound. One might naively consider the following, non-post-hoc objective instead to obtain tighter bounds:

$$W^*, V^* = \arg\min_{W,V} \sum_n \mathcal{L}(V, W; x_n, y_n) + \lambda \max_{P \succeq 0, \mathrm{diag}(P) \leq 1} \langle M(v, W), P \rangle$$

where λ controls the regularization strength. Using this objective is clearly infeasible, however: For each gradient step, we need the solution to the inner semidefinite program. Without going too much into detail, one can obtain a *dual formulation* of the semidefinite program to eliminate the inner optimization problem. First, we state the dual formulation:

$$\max_{P \succeq 0, \mathrm{diag}(P) \leq 1} \langle M(v, W), P \rangle = \min_{c \in \mathbb{R}^{(d+m+1)}} (d+m+1) \cdot \lambda^+_{\max}(M(V, W) - \mathrm{diag}(c)) + \sum_i max(c_i, 0)$$

where $\lambda^+_{\max}(\cdot)$ calculates the maximal eigenvalue of the input matrix or returns zero if all eigenvalues are negative. How can we use this to eliminate the inner optimization problem? As the inner problem becomes the unconstrained minimization of an objective in $c \in \mathbb{R}^{(d+m+1)}$, we optimize c in the same optimization loop as parameters V and W. Therefore, we only have an additional parameter we have to optimize over and we can still use gradient-based unconstrained optimization.

This leads us to the final objective: We optimize

$$(W^*, V^*, c^*) =$$

$$\arg\min_{W,V,c} \sum_n \mathcal{L}(V, W; x_n, y_n) + \lambda \cdot \left[(d+m+1) \cdot \lambda^+_{\max}(M(V, W) - \mathrm{diag}(c)) + \sum_i max(c_i, 0) \right]$$

which can be done quite efficiently. This encourages the network to be robust while also allowing us to provide a certification of robustness.

Given $V[t], W[t]$ and $c[t]$ values at iteration t solving the above optimization problem, one obtains the following guarantee for any attack A:

$$\tilde{f}(A(x)) \leq \tilde{f}(x) + \frac{\epsilon}{4} \left[D \cdot \lambda^+_{\max}(M(V[t], W[t]) - \mathrm{diag}(c[t])) + + \sum_i max(c[t]_i, 0) \right].$$

We get a certificate of the defense: Whatever perturbation there is in the L_∞ ϵ-ball, the loss is bounded from above. This is theoretically meaningful but not yet in practice: The study is confined to 2-layer networks.

---

[a]This statement holds for arbitrary vectors $y \in \mathbb{R}^n$.

2.15.16 History and Possible Future of Adversarial Robustness in ML

A Coarse Timeline of Adversarial Robustness is listed below.

- **First attack**: L-BFGS attack (2014) [192]. This is a complicated method that does not

work too well.

- **First practical attack**: FGSM attack (2014). As discussed previously, this is a straight-forward method that works reasonably well.
- **Stronger iterative attack**: DeepFool (2015) [143].
- **First defense**: Distillation (2015) [79]. Training labels of the distilled network are the predictions of the initially trained network. Both networks are trained using temperature T.
- **First black-box attack**: Substitute model (2016).
- **Strong attack**: PGD (2017).
- **Strong defense**: Adversarial Training (2017).
- **First detection mechanisms**: Adversarial input detection methods (2017). Instead of making the model stronger, we train a second model to detect adversarial patterns. However, the attackers can also generate patterns that avoid these detections (fool the detector in a white-box fashion).
- "It is easy to bypass adversarial detection methods." (2017) [30].
- Defenses at ICLR'18 (2018): input perturbation, adversarial input detection, adversarial training, etc.
- "Defenses at ICLR'18 are mostly ineffective.": Obfuscated gradients (2018).
- Barrage of Random Transforms (2019). One only needs to apply many transformations sequentially in a random fashion.

2020- : We should *stop* the cat-and-mouse game between attacks and defenses. It is a dead end. We are spiraling around attack and defense. Diversifying and randomizing (e.g., BaRT) is a promising approach. However, the constant spiral of whether the attacker or defender has more capacity to generate attacks/defenses is not very interesting from an academic perspective.

There are two main alternatives one may choose to work on:

- Certified defenses making sure there is *no attack* in the L_p ball.
- Dealing with realistic threats rather than unrealistic worst-case threats.

2.15.17   Towards Less Pessimistic defenses

Usually, the considered attacks are way too strong. Instead, we should work more on (1) defenses against black-box attacks, which is an exciting subfield of adversarial attacks, or (2) defenses against non-adversarial, non-worst-case perturbations (OOD gen., domain gen., cross-bias gen.). These are what we have learned in the previous chapters. Many researchers who used to study adversarial perturbations are now working on general OOD generalization and naturally shifting distributions.

Chapter 3
Explainability

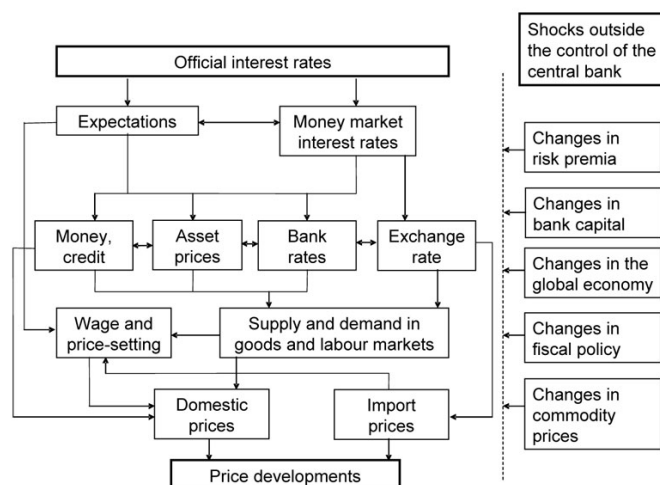# 3 Explainability

3.1 Introduction

Figure 3.1: Schematic illustration of the main transmission channels of monetary policy decisions, taken from [19]. Controlling price developments requires fine-grained control.

If we understand a system and its underlying mechanisms well, we can use the system to control something. An example is the economy: As we control official interest rates,[1] we control the amount of money in the market. Official interest rates affect many components of the economy (e.g., bank rates, exchange rates, or asset prices) and finally also affect price developments (e.g., domestic prices and import prices), all through a highly complex procedure. This is illustrated in detail in Figure 3.1. There is always a new situation coming (shocks outside the central bank's control). One cannot solely rely on experience, as we do not have so much history of the market to base our decisions on previous experiences when we face new situations. It is essential to *know what is happening in the system* to perform control.

When faced with a black box system, we do not understand its inner workings. For example, we do not exactly understand why a self-driving car is following the road in one case but not following it in the other. As we cannot control the system precisely, we cannot fix it when it is malfunctioning.

---

[1]The central bank is directly in control of this through determining official interest rate policies. Similarly, other policy rates and asset purchases have a large effect on how prices develop.

3.1.1 Ways to Control Undefined Behavior

> **Definition 3.1: Explanation**
>
> An explanation is an answer to a *why*-question. [139]

> **Definition 3.2: Interpretability**
>
> Interpretability is the degree to which an observer can understand the cause of a decision. [23]

> **Definition 3.3: Explainability**
>
> Explainability is post-hoc interpretability. [124] It is the degree to which an observer can understand the cause of a decision after receiving a particular explanation.

> **Definition 3.4: Justification**
>
> A justification explains why a decision/prediction is good but does not necessarily aim to explain the actual decision-making process. [23] It is also not necessarily sound.

Two general ways exist to control undefined behavior in OOD (novel) situations: using unit tests and fixing only after understanding.

An infinite list of unit tests and data augmentation. We were looking into this in previous sections. In particular, in Section 2.10, we saw how we can identify spurious correlations in our model, and in Section 2.11, we saw how we can *incorporate* samples from different domains (e.g., unbiased samples) into the training procedure to obtain more robust models. Our goal is to let a model work well in any new environment. For evaluation purposes, we introduce a new evaluation set every time, e.g., introduce ImageNet-{A, B, C, D, ...}. A natural next step is to augment our network's *training* with samples from ImageNet-{A, B, C, D, ...} and seek new evaluation sets. We are sequentially conquering different unit tests, hoping that we eventually get a strong system that works well in any situation. But is that *really* going to happen?

Understand first, fix after. The goal here is the same as before: make a model that works well in any new environment. For evaluation, we examine cues utilized by the model (explainability). If we understand that the model is not utilizing the right cue for recognition, then we have a way to control this. We regularize the model later to choose cues that are generalizable. We do *not* evaluate whether the model works well on ImageNet-{A, B, C, D, ...}, as we directly control the used cues. We regularize the model to choose generalizable cues (using feature selection). This seems to be the more scalable approach. An infinite number of unit tests will probably not solve all our problems.

3.1.2 Explainability as a Base Tool for Many Applications

There are numerous applications that require the *selection of the 'right' features*. In fairness, we wish to eliminate, e.g., demographic biases which requires us to select features that do not take demographic aspects into account. In the field of robustness, we also have to select powerful features to combat distribution shifts.

There are also many applications that require *better understanding* and *controllability*. One example is ML for science where the aim is to discover scientific facts from (usually) high-dimensional data. Here, understanding and control is the *end goal*. We can also consider the task of quickly adapting ML models to downstream tasks (e.g., GPT-3 and other LLMs). If we understood what GPT-3 or other LLMs do/know, we could probably quickly adapt them to downstream tasks by only choosing the parts or subsystems responsible for useful utilities for the downstream task. In that case, we might not even need any fine-tuning.

> **Definition 3.5: Attribution**
>
> Attribution can be understood as the assignment of a reason for a certain event. It is often used in the field of Explainable Artificial Intelligence (XAI) to describe attributing factors to a model's behavior in an explanation. Such factors are often selected from (1) the input's features we are explaining (be it the raw input features or intermediate feature representations of NNs), (2) the elements of the training set, or (3) the model's parameters.

> **Definition 3.6: Explanation by Attribution**
>
> An explanation by attribution method is a function that takes an input x and a model f and outputs the explanation of which features/training samples/parameters contribute most to the prediction $f(x; \theta)$ of the model for input x.

Several applications require a *better understanding of the training data*. Consider the detection of private information in a training dataset. Instead of attributing to the test data, we can also reason back to the training data. For example, if our model seems to have learned something private from user data and users can even be identified based on this information, it would be very informative to be able to trace back specific predictions to the training data (and remove private data from the training set or make sure that such information cannot leak). If an LLM outputs something that looks like someone's home address, tracking down where this information came from in the training set is very informative for those who audit the training data. Attributing to the original authors in the training data is an increasingly popular and useful task. Example questions include "What prior art made DALL-E generate a certain image? What authors can be attributed?" XAI can give answers to such questions.

Finally, let us discuss applications requiring *greater trust*. One example is ML-human expert symbiosis where a human expert is working with ML to generate better outcomes. Trust is also needed in high-stakes decision areas: for example, finance, law, and medical applications.

3.1.3 Explainability as a Data Subject's Right

Nowadays, explanations are stipulated in law [64]. XAI has close ties to national security – the research field originated from the DARPA XAI program of the US, in 2016. The EU also considered AI legislation crucial – GDPR has an article about automated decision-making, and the AI act is an even newer *proposal* of harmonized rules on general artificial intelligence systems. A common theme of AI legislation is that suitable measures are needed to safeguard a data subject's rights, freedom, and legitimate interests. Data subjects have a right to request explanations in automated decision-making and to obtain human intervention. Critical decisions are made about humans based on automatized systems (ML) using their personal data, e.g., in CV preselection or loan applications. The data owner has the right to know which feature has caused, e.g., a loan rejection.

Three Key Barriers to Transparency

There are mainly three barriers to transparency. Let us briefly discuss these.

Intentional concealment. For example, a bank might intentionally conceal their decision procedure for loan rejection. Decision-making procedures are often kept from public scrutiny.

Gaps in technical literacy. For example, even if the bank is enabling insight into its decision-making procedure, people may not be able to understand the raw code. For most people, reading code is insufficient.

Mismatch between actual inner workings of models and the demands of human-scale reasoning and styles of interpretation. This is perhaps the most technical aspect this book seeks answers to. Human-comprehensibility was highlighted as a crucial aspect of XAI methods by several researchers [139, 142, 21]. If we are showing the weights of a model to a human, it is unlikely that they see some meaning. We need summarization, dimensionality reduction, and attachment of human-interpretable concepts.

We have answered "Why is an explanation needed?" Let us turn to "When is an explanation needed?"

3.1.4 When is an explanation needed?

The following points are inspired by [47, 95]. Explanations may highlight an incompleteness/problem. In particular, explanations are typically required when something does not work as expected. When everything is working well, we usually do not question why something is working. When something does not work, we start raising questions.

3.1.5 When is an explanation *not* needed?

First, we discuss a list of examples from [47]. We also argue why this list of examples might not be descriptive enough.

- **Ad servers.** Our remark is that it is a request of society in general that they should be prompted for consent if they want to see targeted ads, and also to gain insight into how profiling works.

- **Postal code sorting.** Even though in general, society might not care much about the inner workings of post offices, explanations might still be needed for debugging such sorting systems or for unveiling potential security risks.
- **Aircraft collision avoidance systems.** Again, explanations for such systems are generally not requests of society. Still, the aviation company must be in control of all situations that might arise, and for that, explanations are great tools.

The above list lacks *recipients*. Whether an explanation is needed in a certain situation depends on the *explainee*. A person not using the internet might, indeed, not care about ad servers. Similarly, a person who does not work as a developer for a post office might not need explanations about the sorting algorithm. Still, we can almost always find target groups for explanations about any topic.

The general reasoning of [47] is sound: generally, we might not need explanations when

1. there are no significant consequences for unacceptable results, or when
2. the problem is sufficiently well-studied and validated in real applications that we *trust* the system's decision (even if the system is not perfect).

However, explanations are always great tools for exploratory analysis.

3.2 Human Explanations

3.2.1 How do humans explain to each other?

We discuss Tim Miller's work, titled "Explanation in Artificial Intelligence: Insights from the Social Sciences" [139].[2] According to [136], people ask for explanations for two main reasons:

1. **To find meaning.** To reconcile the contradictions or inconsistencies between elements of our knowledge structures. We are trying to figure out at which principle we have contradictions. There are often contradictions between our understanding and the status quo in the outside world.
2. **To manage social interaction.** To create a shared meaning of something, change others' beliefs and impressions, or influence their actions. Example questions include "Why am I doing this? Why are you doing this?" But also: "If I believe what you are doing has a greater cause, I can also align my action to what you are doing."

Both are important for XAI systems.

- **Finding meaning in XAI.** "Why is this model not doing as I expect? Where is this inconsistency coming from?"
- **The social aspect of XAI.** We want to be able to share our way of thinking with the machine, and we expect it also to be able to do the same.[3]

---

[2]This is a highly recommended work for those working or wishing to work in XAI.

[3]The field of Human-AI Interaction works on such methods. One possible way of knowledge exchange is through textual discussions, as seen in LLMs.

Human-to-human explanations are . . .

. . . contrastive. Explanations are sought in response to particular counterfactual cases. People usually do not ask, "Why did event P happen?", they ask, "Why did event P happen instead of some event Q?" Even if the apparent format is the former, it usually *implies* a hidden foil (i.e., the alternative case). **Example**: For the question "Why did Elizabeth open the door?", there are many possible foils. (1) "Why did Elizabeth open the door, rather than leave it closed?" – a foil against the action. (2) "Why did Elizabeth open the door rather than the window?" – a foil against the subject. (3) "Why did Elizabeth open the door, rather than Michael opening it? – a foil against the actor. A brief criticism of XAI is that the questions asked often do not have any foil in mind in general. We ask questions like, "Why was this image categorized as A?" It would be perhaps less ambiguous to ask, "Why is this image categorized as A, not B?" This formulation makes the foil clear. Another way to extend the question: "Will this image still be categorized as A even if the image is modified?" We will see that this kind of question is implied in many XAI systems. In a sense, input gradients are asking such contrastive questions.

. . . selective. When someone asks for an explanation for some event, they are usually not asking for a complete list of possible causes but rather a few important reasons and causes relevant to the discussion at hand. Humans are adept at selecting one or two relevant key causes from a sometimes infinite number of causes as the explanation. If we generate all kinds of causes for explaining a single event, the causal chain can be too large and hard to handle for the explainee. The principle of simplicity dictates that the explainer should not overwhelm the explainee.

. . . social and context-dependent. Philosophy, psychology, and cognitive studies suggest that we are not explaining the same thing to everyone – we change the way we explain based on whom we are talking to. The way we explain depends greatly on our model of the other person. People employ cognitive biases and social expectations. Explanations are a transfer of knowledge, presented as part of a conversation or interaction. If a person we are talking to does not know something, we are filling in the gap in their understanding. If they seem to understand the subject well, we can share less obvious causes for an event too. Explanations are thus presented *relative* to the explainer's beliefs about the explainee's beliefs.

. . . interactive. Through the exchange of explanations and confirmation of understanding, we can continuously stay on the same page. The explainee can let the explainer know what subset causes they care about that are relevant for them. The explainer can then select a subset of that subset based on other criteria. The explainer and the explainee can interact further and argue about explanations. In XAI, there have been relatively few works on interactive explanations so far. Typically, we generate human-agnostic explanations that should work for everyone. Based on human interactions, we should be able to generate personalized explanations.

Because of these properties, there is no single correct answer to "Why?".

3.3 Properties of Good Explanations

3.3.1 What are good explanations?

From now on, we will be using more refined terminologies that are also used in XAI. The properties of a good explanation we deem most important are listed and explained below.

Soundness/faithfulness/correctness. The explanation should *identify the true cause for an event*. This is the primary focus of current XAI evaluation: The attributions should identify the true causes (a model used) for predicting a certain label. It is also high on the list of desiderata from domain experts [114]. ("What do you need from explainability methods?") However, it is important to highlight that this is not the only criterion for a good explanation. **Example**: "Why did you recognize a bird in this image?" If the model points to a feature that does not contribute to its prediction of 'bird', then its explanation is not sound/faithful. Another example is the following. A doctor wants to know the actual thought process of the system rather than just a likely reasoning from a human perspective. By understanding what is going on inside recognition mechanisms, we might be able to learn more than what humans are currently capable of extracting from an image.

Simplicity/compactness. The explanation should *cite fewer causes*. A good balance is needed between soundness and simplicity (such that humans can handle the explanation).

Generality/sensitivity/continuity. The explanation should *explain many events*. They should not only explain very specific events – one usually seeks a general explanation. In XAI, generality means that the explanation should apply to many (similar) samples in the dataset.

Relevance. The explanation should be *aligned with the final goal*. This criterion asks "What do we need the explanation for?" If we need it for fixing a system, the explanation should help us fix it. If we need explanations for understanding, a good explanation should then let us understand the event in question.

Socialness/interactivity. Explaining is a social process that involves the explainer and the explainee. In XAI, the explainer is the XAI method, and the explainee is the human. As mentioned before, the explanation process is dependent of the explainee. A good explanation could consider the social context and adapt and/or interact with the explainee. It should not always cite the most likely cause but also retrieve causes that are interesting for the user. These do not necessarily coincide.

Contrastivity. The foil needs to be clearly specified. Ideally, a method should be able to tell how the model's response changes when we change something from A to B. An example of contrastivity is shown in the paper "Keep CALM and Improve Visual Feature Attribution" [99], which we are going to discuss later in more detail. An illustration of contrastive explanations is given in Figure 3.2.

Human-comprehensibility/coherence/alignment with prior knowledge of the human. The given explanation should fit the understanding and expectations of the human. It should also be presented in a format that is natural to humans.

The first property, soundness, is much more often used in XAI for evaluation. Nevertheless, the others are just as important.
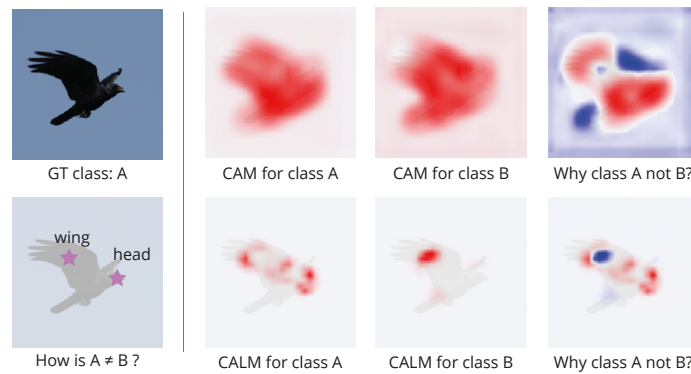
Figure 3.2: Example of contrastive explanations in XAI in the last column, taken from [99].

3.3.2 Intrinsically Interpretable Models

Intrinsically interpretable algorithms are generally deemed interpretable and need no post-hoc explainability [47]. A DNN is not like this. We can generate an explanation for a prediction post-hoc, making it explainable, but the DNN itself is not interpretable.

Sparse Linear Models with Human-Understandable Features

Sparse linear models make a prediction x using the following formula:

$$x = \sum_{i \in S} c_i \phi_i \qquad |S| \ll m.$$

The prediction is a simple linear combination of features ϕ_i with coefficients c_i. We also understand the individual ϕ_is very well, as they are human-understandable. Sparse linear models further contain a small number of coefficients: There are few enough coefficients for humans to understand the way the model works. Every feature ϕ_i is contributing to the prediction by giving a factor $c_i \phi_i$ to the sum. We know the exact contributions. We are citing every ϕ_i as the cause of the outcome x. This explanation is very general – it works for any ϕ_i and any value thereof.

> **Additional Information: Feature Attribution vs. Feature Importance**
>
> By construction, when $\phi_i = 0$, the feature contributes with a factor of zero to the final prediction. By treating $c_i \phi_i$ as our attribution score for feature i, we cannot give a non-zero attribution score to features whose value is zero. Depending on what ϕ_i is encoding, this might have surprising consequences. For example, when the individual features are pixel values ($\phi_i \in [0, 1]$), we cannot attribute to black pixels. To resolve this, we note that attribution scores are not synonymous with *feature importance* – the score we obtain is a reason for the prediction x w.r.t. feature ϕ_i, which is exactly $\phi_i c_i$. This does not mean that this feature is unimportant, it just means that this term contributed the value 0 to the final prediction. The importance of the feature is better measured by the coefficient c_i, but one can only directly interpret it as the (signed) importance of the feature in the prediction of the model if (1) the features are uncorrelated and (2) they are

on the same scale (e.g., by standardization). The coefficient c_i can also be used as an attribution score: while $c_i\phi_i$ measures the *net* contribution of the feature for the prediction of x, the coefficient c_i measures the *relative* contribution of the feature if we slightly change that feature.

Do these sparse linear models explain well according to our criteria?

Let us consider the aforementioned criteria and evaluate sparse linear models according to them.

Soundness/faithfulness/correctness. By definition, every feature is a sound cause for the outcome, with contribution given by the terms $c_i\phi_i$.

Simplicity/compactness. One can control this aspect with the number of features. Sparsity enforces the model to use few causes. We could not understand the model's decision-making process if we had millions of features.

Generality/sensitivity/continuity. By definition, whenever the cited causes happen, similar outcomes follow.

Relevance. This criterion always depends on the final goal. Is it to debug? Or to understand? For debugging purposes, we measure the quality of the explanation in terms of whether it is actually helping a human find features that are not working and whether the human can fix the system based on the explanation.

Socialness/interactivity, Contrastivity. One can simulate contrastive reasoning from the ground up. Instead of having $\phi_1,\ldots,\phi_{|S|}$, we can leave one out and see what happens afterward. This is the counterfactual answer for the effect of leaving a feature out. However, the models are not social and interactive by default. An additional module is required for that. The explanation is also not personalized; it does not consider the level of knowledge of the explainee. We can attach a chatbot or interactive system to make interactivity possible.

Decision Trees with Human-Understandable Criteria

In a decision tree that follows human-understandable criteria, all deciding features (criteria) follow a human-understandable concept. Unless the depth of the tree is too large, or we have too many branches, the entire tree is human-interpretable. **Example**: The task is to predict the animal breed. Features are not too many well-known properties of animals. Humans can then directly understand how the decision tree makes a particular prediction.

3.4   Taxonomies of Model Explainability

There are different ways to divide the set of explainability methods. The correlated "axes" of variation are as follows.

Intrinsic vs. Post-hoc

Intrinsic interpretability means the model is interpretable by design (i.e., sound, simple, and general at once). The way the input is transformed into the output is interpretable. **Examples**: sparse linear models and decision trees.

Post-hoc explainability means the model lacks interpretability, and one is trying to explain its behavior post-hoc. **Example**: turning a DNN into a more interpretable system. The intrinsically interpretable models discussed before are very useful throughout our studies of explainability: We often turn a part of our network into a linear model and analyze that (e.g., Grad-CAM [174] in 3.5.17) or approximate the whole model using a sparse linear model and explain it in a local input region (e.g., LIME [165] in 3.5.9).

Global vs. Local

Global explainability means the given explanation is not on a per-case basis. We do not want to understand one particular event but rather the entire system, allowing us to understand per-case decisions as well. An example of global explainability is the SIR epidemic model [97], shown in Figure 3.3. Here, β, γ are rates of transitions. The system is based on the differential equations in Figure 3.3 that explain the entire system. We simulate the future based on our choice of β, γ. This gives an overall understanding of the mechanism but is often impossible to give for complex, deep black-box models. It is particularly useful for scientific understanding and simulation of counterfactuals. ("What would happen if we changed some parameters?")

$$\frac{dS}{dt} = -\beta SI$$

$$\frac{dI}{dt} = \beta SI - \gamma I$$

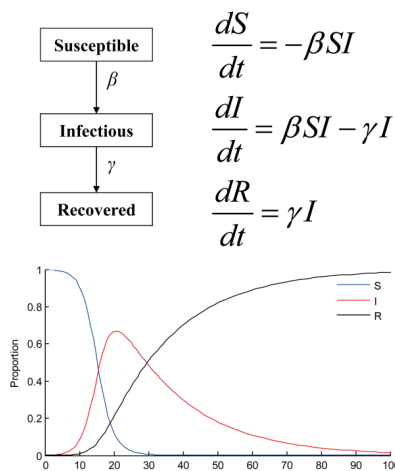$$\frac{dR}{dt} = \gamma I$$

Figure 3.3: The SIR epidemic model is an example of a global explainability tool. The differential equations above determine the behavior of the system. Having chosen the parameters β, γ, we simulate the future. Figure taken from [128].

Local explainability means we want to understand the decision mechanism behind a particular case/for a particular input. **Example**: "Why did my loan get rejected?" – explanations for this do not lead to a global understanding of the system. Local explainability is the main focus/interest of the book. This type of analysis is feasible in somewhat sound ways even for complex models.

Attributing to Training Samples vs. Test Samples

A model is a function approximator. It is also an output of a training algorithm. The input to the training algorithm is the training data and other ingredients of the setting. We write the model prediction as a function of two variables:

$$Y = \mathsf{Model}(X; \theta) = \mathsf{Model}(X; \theta(X^{\mathsf{tr}}))$$

where Y is our prediction, X is the test input, θ are the model parameters, and X^{tr} is the training dataset. The prediction of our model is implicitly also a function of the training data.

We can trace back (i.e., attribute) the output Y for X to either (1) particular features X_i of the test sample $X = [X_1, \ldots, X_D]$, or (2) particular training samples X_i^{tr} in the training set $X^{\mathsf{tr}} = \{X_1^{\mathsf{tr}}, X_2^{\mathsf{tr}}, \ldots, X_N^{\mathsf{tr}}\}$.

One may also attribute the prediction to a particular parameter θ_j or a layer, but individual parameters are often not very interpretable to humans. Usually, we "project parameters" onto the input space by gradient-based optimization.

> **Additional Information: Correlations in the axes of variation**
>
> Models that are intrinsically interpretable are interpretable on a global scope – they give an understanding of the whole model. But they can also be used to explain particular decisions based on an input. While explaining local decisions is also possible, the focus is rather on the global scale.
>
> Methods with intrinsic interpretability also do not *have to* directly attribute their predictions to anything. However, this is still often possible, e.g., in the case of sparse linear models.

3.4.1   Soundness-Explainability Trade-off

Explanations try to linearize a model in some way. What humans can naturally understand is a summation of a few features (i.e., a sparse linear model). There is an inherent *soundness-explainability trade-off*. One extreme is the original DNN model by itself: It is by definition sound but not interpretable. Another extreme is creating a sparse linear model as the global linearization of a DNN around a particular point of interest. It cannot be sound as a global explanation but is very interpretable.

Between the two extremes, explanation methods try to linearize different bits of the model for either the entire input space (generic input) or for a small part of it. It is relatively easy to linearize the full model for a small part of the input space. (For example, LIME, discussed in Section 3.5.9.) It is also relatively easy to linearize a few layers of the model for generic input. (For example, Grad-CAM, discussed in Section 3.5.17.) However, it is impossible to faithfully linearize the full model for generic input. Then, we are back to global linearization.

3.4.2 Current Status of XAI Techniques

XAI research is often harshly criticized for being useless.

"Despite the recent growth spurt in the field of XAI, studies examining how people actually interact with AI explanations have found popular XAI techniques to be ineffective, potentially risky, and underused in real-world contexts." [49]

People working in human-computer interaction are very critical of XAI techniques in ML conferences, as they often do not take humans into account appropriately yet. It is essential to condition our mindset to help those who wish to actually use XAI techniques rather than working on techniques that look fancy or theoretically beautiful (e.g., completeness axioms, 3.5.5).

"The field has been critiqued for its techno-centric view, where "inmates [are running] the asylum", based on the impression that XAI researchers often develop explanations based on their own intuition rather than the situated needs of their intended audience (final goal is not taken into account). Solutionism (always seeking technical solutions) and Formalism (seeking abstract, mathematical solutions) are likely to further widen these gaps." [49]

We want to move away from developing such XAI techniques and focus on the demands of those needing XAI systems.

Note: Formalism *is* helpful (both in method descriptions and evaluation), but the most important aspect should be whether these methods actually help people. Formalism is not the end goal.

3.5 Methods for Attribution to Test Features
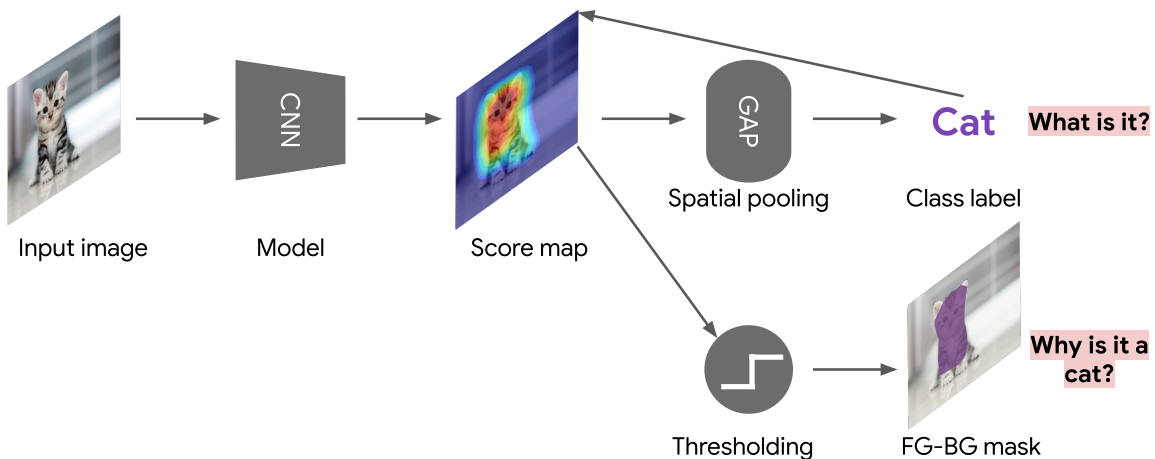
Figure 3.4: Simplified high-level overview of the CAM method. The model makes a prediction ('cat'), which is then used to select the appropriate channel of the Score Map that describes the attribution scores for class 'cat'. Finally, an optional thresholding can be employed to make a binary attribution mask.

So far, we have laid down what we desire from explainable ML. In this section, we discuss actual methods for extracting explanations from DNNs. In particular, we will look at methods that attribute their predictions to test features. Instead of the "What is in the image?"

question, explanation methods seek to answer the "Why does the model think it is the predicted object?" question.[4] For example, CAM [223] produces a score map w.r.t. the predicted label, as illustrated in Figure 3.4. We can threshold it to get a foreground-background mask as an explanation (which is not necessarily a mask w.r.t. the GT object location, as the network being explained might have, e.g., background or texture biases). We can also leave out thresholding and keep continuous values in the map.

3.5.1   What features to consider in attribution methods to test features?

Definition 3.7: Superpixel

Superpixels are groupings of pixels respecting color and edge similarity (that very confidently belong to the same object instance). It gives us a finer grouping than semantic segmentation (in the sense that the pixels are not grouped into only a couple of categories, but rather into many patches of pixels that closely belong together) but a coarser one than the raw pixels. An illustration is given in Figure 3.5. There have been many improvements in superpixel technology until a few years ago. Nowadays, not many people are looking into superpixel methods. These are often used as features in explainability for images. They reduce the number of features we have to deal with without sacrificing soundness.
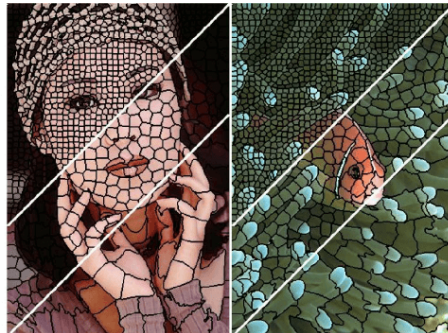
Figure 3.5: Illustration of superpixels of various granularities, which is a popular choice of features for attribution maps. Figure taken from [1].

We generalize the notion of a feature to any aggregation or description of the input to the model. Possible features are listed below for visual models taking an image as input. These are also illustrated in Figure 3.6.

1. **Single pixels.**
2. **Image patches.** We can aggregate pixels into image patches, considering each patch as a feature.
3. **Superpixels.**
4. **Instance mask(s).**
5. **High-level attributes.** For example, attributes for a cat image input can be Cute, Furry, Yellow eyes, Two ears, Animal, and Pet. The values for each of these attributes can be

---

[4]Curiously, some explanation methods are also good at Weakly Supervised Object Localization (WSOL), that aims to answer the "Where is the object in the image?" question.
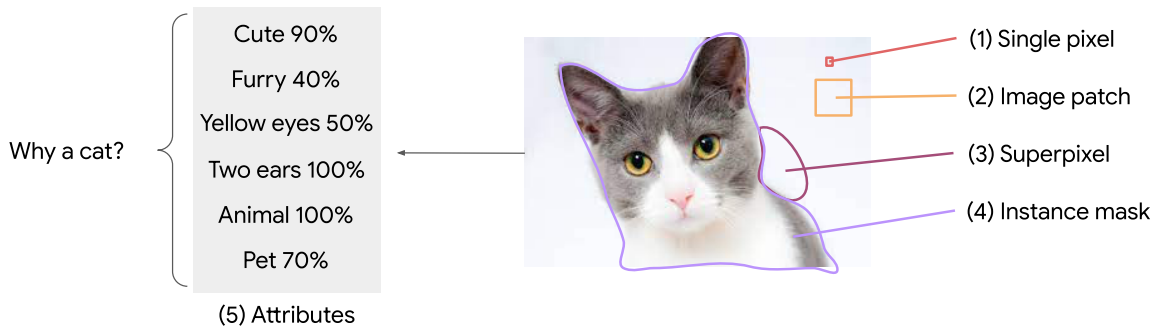
Figure 3.6: Illustration of several feature representations for the same image. There is a wide range of features we can attribute to.

percentages representing how fitting a certain attribute is for the input. For example, Two ears → 100% means the feature is maximized, i.e., the attribute perfectly fits the input. **Note**: These are *not* the attribution scores corresponding to the individual attributes. Attribution scores are values describing how each of the features influences the network prediction, whereas the attributes *describe* the input. The attributes can be subjective, can point to specific regions of the image, and can also describe, e.g., the general feeling of the image.

For natural language models taking a token sequence as inputs, we often use individual tokens/words as features. We can think about the contribution of each token (or word) towards the final prediction (e.g., sentiment analysis), as considered in the paper "A Song of (Dis)agreement: Evaluating the Evaluation of Explainable Artificial Intelligence in Natural Language Processing" [146]. Examples of attributing to individual tokens can be seen in Figure 3.7.

Note: Explanation methods can give significantly different results for the same input, as shown in Figure 3.7. This has also been reported in the paper titled "The Disagreement Problem in Explainable Machine Learning: A Practitioner's Perspective" [109]: local methods approximate the model at a particular test point x in local neighborhoods, but there is no guarantee that they use the same local neighborhood. Indeed, since different methods use different loss functions (e.g., LIME with squared error vs. gradient maps with gradient matching), it is likely that different methods produce different explanations.

Additional Information: Choice of Features

If we gather all the features of an image, do we have to obtain the original image by definition? The answer is *yes*; we generally wish to *partition* the image with features.
- For partitioning, one may choose panoptic segmentation [103], a combination of instance segmentation and semantic segmentation. This considers both object and stuff masks (where stuff refers to, e.g., 'road', 'sky', or 'sidewalk'). Another option is regular semantic segmentation, which can also handle various stuff categories. The COCO-Stuff [29] dataset gives many examples of how semantic segmentation can partition images in a detailed way.
- Considering only image parts corresponding to different instance masks as fea-
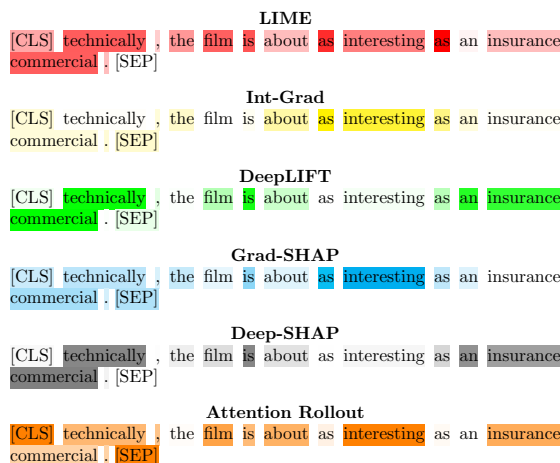
LIME

[CLS] technically , the film is about as interesting as an insurance commercial . [SEP]

Int-Grad

[CLS] technically , the film is about as interesting as an insurance commercial . [SEP]

DeepLIFT

[CLS] technically , the film is about as interesting as an insurance commercial . [SEP]

Grad-SHAP

[CLS] technically , the film is about as interesting as an insurance commercial . [SEP]

Deep-SHAP

[CLS] technically , the film is about as interesting as an insurance commercial . [SEP]

Attention Rollout

[CLS] technically , the film is about as interesting as an insurance commercial . [SEP]

Figure 3.7: Sentiment analysis example. Explanation methods give significantly different attributions. "The average Kendall-τ correlation across all methods for this example is 0.01." [146] Figure taken from [146].

> tures is problematic, as stuff information is thrown away (we get rid of stuff categories), and we do not have a partition of the original image anymore.

A feature is thus a general concept. The task for feature attribution methods is to determine which feature contributes how much to the model's prediction.

In the last section, we have seen that counterfactual (i.e., contrastive) reasoning matters a lot in explaining to humans. The most basic way to explain a model's decision in a counterfactual way is by asking a question of the form "Is the input image still predicted as a cat if this feature is missing?" We remove a particular set of pixels from the image and see how the model's prediction changes. We have many possibilities to encode what we mean by a "missing" pixel. For example, we can fill them with black, gray, or even pink pixels (which are rarely seen in natural images but do not intuitively encode a baseline image). We can even choose to inpaint them based on the context. One could also ask counterfactual questions of the form "Is the input image still predicted as a cat if this feature is replaced with something else?" In this case, we can, e.g., insert an image of a dog in the "missing" patch, illustrated in Figure 3.8. After carrying this out for all pixels, we get an answer to "Which features contribute most to predicting a cat rather than a dog?" Even in the simple setting of removing a square patch from an image, many things must be considered.

3.5.2 Intrinsically interpretable models support counterfactual evaluation by design.

In a DNN, when we change something in the input, it is highly unclear how the forward propagation is influenced to obtain the final answer. In decision trees, we can just change one attribute in any way and check how the result changes (by selecting the other branch at a corresponding attribute). We can do a full simulation quickly where we understand each part of the decision-making process. We can still do the simulation for a DNN, but we only observe the outputs (before and after the change) in an interpretable way. We have no good intuition
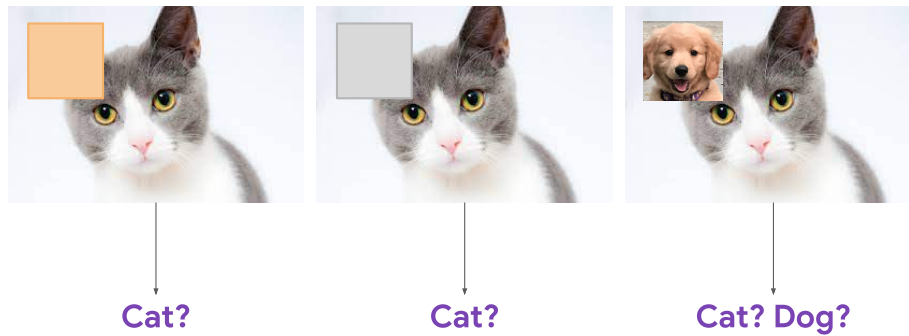
Figure 3.8: Three possibilities of counterfactual explanations. The left image encodes the missingness of a patch by orange pixel values. The center image encodes missingness by gray pixel values. These give answers to the question "How would the prediction change if we removed a patch of the image?" The right image asks a slightly different question: "Which features contribute to predicting a cat rather than a dog?", as the dog image does not aim to encode missingness, i.e., we cannot talk about removing the patch.

about what changes *internally*.

A sparse linear model is just a summation. Every feature contributes linearly to the final output. It is easy to interpret the relationship between the features and outputs. We know the full effect on the output of changing (or removing) one or many features. Our implicit aim is to linearize our complex models in some way for interpretability.[5] This is a common mindset of attribution methods. Because of the linear relationship between inputs and outputs, we do not have to compute differences between outputs to study counterfactual evaluation. We already know how the output changes by changing some inputs. This is highly untrue for DNNs, requiring a forward pass each time. We will see that under some quite strong assumptions, we can use the input gradient and derivative quantities for counterfactual evaluations.

3.5.3 Infinitesimal Counterfactual Evaluation in Neural Networks: Saliency Maps

We can perform the removal analysis for all input features for neural networks, e.g., using a sliding window of patches as features. This, however, takes very long for DNNs. For each image, one needs to compute N forward operations through a DNN, where

$N =$ number of sliding windows per image $\times$ number of ways to alter the window content.

Doing this in real-time during inference on a single sample is infeasible without sufficient computational resources for parallelization. Doing it offline for an entire dataset also takes very long if the dataset is large. One can use batching, but only a small number of samples fit on the GPU usually.

However, we can consider a special case where our *features are pixels* and the *perturbation is small* (infinitesimal). In this case, we can compute counterfactual analysis quickly, at the cost of the huge restriction of the perturbation size being small.

Example: Consider pixel $(56, 25)$ with original pixel value: $(232, 216, 231)$. Suppose that all

---

[5]We might not want to linearize the entire model. Partial linearization is often used, e.g., in Grad-CAM (3.5.17) and TCAV (3.5.13).

pixels are left unchanged except this one where the new pixel value is set to $(233, 216, 231)$.[6] Further, suppose that the original cat score was 96.5%, but after the change, the cat score for the perturbed image decreases to 96.4%. This seems familiar: That is exactly how we approximate the gradient numerically:

$$\frac{\partial f(x)}{\partial x_i} \approx \frac{f(x + \delta e_i) - f(x)}{\delta}, \qquad \delta \text{ small.}$$

In the ordinary sense, the gradient of the score (or probability) of the predicted class w.r.t. an RGB pixel is a 3D vector (as the RGB pixel itself is also a 3D vector). However, we will consider $x \in \mathbb{R}^d$ as a flattened version of an image and will also collapse the color channels. We treat the elements of the resulting vector as pixels. Therefore, the ith pixel direction does *not* correspond to the general definition of a pixel in the following sections.

> **Additional Information: Discrete Representations of Color**
>
> The 8-bit representation is just a convention for RGB images. There exist 16- and 32-bit representations too. The RGB scale is continuous.

A very inefficient way to compute the attribution of each pixel is to compute the forward pass (number of pixels + 1) times (perturbed images plus original image, as the latter is shared in all gradient approximations) to measure pixel-wise infinitesimal contribution. A large approximate gradient signals a significant contribution of the corresponding pixel for an infinitesimal perturbation (because of a significant change in the score of the predicted class).

Note: Here, we consider the *relative* contribution of a pixel (as we equate high contribution to a high *relative change* in the network output for an infinitesimal perturbation), similarly to the sparse linear model case where the relative contribution of feature ϕ_i was given by the coefficient c_i. Of course, this was just a special case of the gradient for the sparse linear model case: if we differentiate $x = \sum_{j \in S} c_j \phi_j$ w.r.t. ϕ_i, we get back c_i again.

> **Additional Information: On the Properties of Gradients**
>
> The derivative
> $$f'(x) = \lim_{h \to 0} \frac{f(x + h) - f(x)}{h}$$
> is a normalized quantity. It gives the *relative* change in the function output, given an infinitesimal change in the input.

The smart way to compute changes in the output w.r.t. infinitesimal perturbations: Compute one forward and one backward pass w.r.t. the score of the predicted class to measure attributions for this infinitesimal perturbation. This answers the question "What will be the *relative change* in the predicted score if we change a particular pixel by an infinitesimal amount?"

This leads us to the definition of *Saliency/Sensitivity maps*.

---

[6]This is the smallest possible perturbation with bit depth 8 – a coarse approximation of the gradient.

Definition 3.8: Saliency/Sensitivity Map

The saliency or sensitivity map visualizes a counterfactual attribution for an input corresponding to infinitesimal independent per-pixel perturbations. It gives us a local explanation of the model's prediction. There are two usual ways to compute it.

Denoting the saliency map for input $x \in \mathbb{R}^{H \times W \times 3}$ and class $c \in \{1, \ldots, C\}$ by $M_c(x) \in [0,1]^{H \times W}$, the SmoothGrad [184] paper computes it as

$$(\tilde{M}_c(x))_{i,j} = \sum_k \left| \frac{\partial S_c(x)}{\partial x} \right|_{i,j,k}$$

(we take the L_1 norm of each pixel), and

$$(M_c(x))_{i,j} = \min\left(\frac{(\tilde{M}_c(x))_{i,j} - \min_{k \in \{1,\ldots,H\}, l \in \{1,\ldots,W\}}(\tilde{M}_c(x))_{k,l}}{P_{99}(\tilde{M}_c(x)) - \min_{k \in \{1,\ldots,H\}, l \in \{1,\ldots,W\}}(\tilde{M}_c(x))_{k,l}}, 1 \right)$$

where $S_c(x)$ is the score for class c given input x and P_{99} is the 99th percentile. This post-processing normalizes the saliency map to the $[0,1]$ interval and clips outlier pixels by considering the 99th percentile. Not clipping the outlier values could result in a close-to-one-hot saliency map.

In Simonyan et al. (2013) [182] (the original saliency paper), the authors compute it as

$$(\tilde{M}_c(x))_{i,j} = \max_k \left| \frac{\partial S_c(x)}{\partial x} \right|_{i,j,k}$$

and the normalization method is not disclosed.

Definition 3.9: First-Order Taylor Approximation

Consider a function $f : \mathbb{R}^d \to \mathbb{R}$. The first-order Taylor approximation of the function f around $x \in \mathbb{R}^d$ is

$$f(x + h) \approx f(x) + \langle h, \nabla f(x) \rangle.$$

Backpropagation linearizes the whole model around the test sample. To see this, observe that the gradient is used to construct the first-order Taylor approximation of the model around a particular test sample, which is the tangent plane of the model around the test sample:

$$f(x + \delta e_i) - f(x) \approx \left\langle \delta e_i, \frac{\partial f(x)}{\partial x} \right\rangle = \delta \frac{\partial f(x)}{\partial x_i}$$

where f gives the score for a fixed class c that is omitted from the notation. This tangent plane guarantees that the function output with this linearized solution will be as close as possible (in the set of linear functions) to the original function's output around the test input of interest in an infinitesimal region. We give a local (counterfactual) explanation with this linear surrogate model, as we only consider an explanation for a single test input. With this surrogate model, one can very cheaply compute input-based counterfactuals. However, these

will only be faithful to the original model in a tiny region around the test input of interest.

Note: Our surrogate model is linear but is not guaranteed to be sparse! It can still be hard to interpret when the input dimensionality is huge. This is primarily the reason why, instead of looking at actual gradient values, we visualize the dense gradient tensors in the form of saliency maps.

Summary of Infinitesimal Counterfactual Attribution

With local gradients, we obtain

$$f(x + \delta e_i) - f(x) \approx \left\langle \delta e_i, \frac{\partial f(x)}{x} \right\rangle = \delta \frac{\partial f(x)}{x_i}$$

which measures contribution of each pixel i with an infinitesimal (δ) counterfactual.

Problem with Saliency Maps
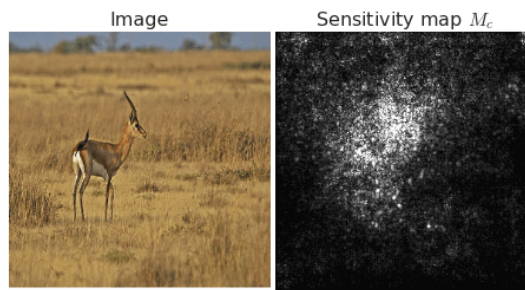
Figure 3.9: Example saliency map of image x w.r.t. the class 'gazelle', taken from [184]. Saliency maps can be challenging to interpret.

We visualize input gradients using saliency maps. These visualizations are not particularly helpful, as they are very noisy and hard to interpret further than a very coarse region of interest. An example is shown in Figure 3.9. **Note**: saliency maps are always w.r.t. a class c. We almost always compute it w.r.t. the DNN's predicted class. We might ask ourselves, "What do we actually get out of this?" We do not even see the object in these input gradient maps. Gradient maps only represent how much *relative* difference a tiny change in each pixel of x would make to the classification score for class c. It is debatable whether one should measure attribution values based on such infinitesimal changes.

Negative contributions are counted as contributions here. This varies from method to method, and no "good" answer exists. The ith element of the gradient measures the relative response of the classification score of class c to a perturbation of the image in the ith pixel direction. If it is positive, making the pixel more intensive results in a locally positive classification score change. If it is negative, it means we reach a higher classification score if we dim the pixel. Sometimes we only want to attribute to pixels that induce a positive change in the score when made more intensive. Sometimes we also want to take negative influences into account.

Additional Information: Gradients and Soundness

It is a *fact* that the gradient gives us the *true* relative changes in the prediction considering per-pixel, independent infinitesimal counterfactuals. It is very important to not confuse this fact with the statement that gradients give perfectly sound attributions in the sense that they flawlessly enumerate the true causes for the network making a certain prediction.

Soundness, by definition, measures whether the attribution method recites the true causes for the model to predict a certain class. As the true causes are encoded in the model weights and the forward propagation formula, which is not at all interpretable, it seems clear that *no attribution method that presents significantly simpler reasoning can be perfectly sound*. Saliency maps – that *seek* to give sound counterfactual explanations w.r.t. an infinitesimal perturbation – make use of such simple reasoning: linearizing the network around the input of interest and taking the rates of change as attribution scores. The linearization, the independent consideration of inputs (with which we discard the possible influence of input feature correlations on the network prediction), and the "arbitrary" normalization and aggregation techniques of the 3D gradient tensor are all significant simplifications that make saliency maps *impossible to be completely sound*.[a] Even if they *were* sound explanations w.r.t. infinitesimal perturbations, the question itself already seems oddly artificial: "Why did the network make a certain prediction for input x compared to an infinitesimally perturbed version of it?". There is no reason why society would demand explanations for such answers.

Moreover, feature attribution methods restrict the explanations to the features in the test input, but the true causes for a network to predict a certain output can also lie in the training set samples and the resulting model weights. Feature attribution methods only consider the test input features as possible causes and make crude assumptions to compute attribution scores. For the soundness of the *explanation*, the attribution method has to give the exact causes of why the network made a certain prediction for an input x. We argue that these exact causes cannot be encoded in general into a map measuring infinitesimal perturbations. Of course, most feature attribution explanations do not claim to provide sound explanations. Instead, they aim to highlight that, given an input, some features were more important in a certain decision than others.

To summarize, the saliency map does not give a perfectly sound attribution map for the predictions of the model on the input of interest because it uses abstractions and simplifications to make the explanation human-understandable.

Note: The soundness of an attribution method and the counterfactual or non-counterfactual nature of explanations it gives are completely independent. For non-counterfactual explanations, a sound attribution method simply aims to give the true influence of each feature of the original input on the network prediction *without* comparing to other predictions.

---

[a]Strictly speaking, the linearization is not a simplification when considering infinitesimal perturbations. However, such perturbations are fictitious, and if one wants to obtain the *net* changes in the network output, they have to consider small δ values that are not exact anymore.

3.5.4 SmoothGrad – Smoother Input Gradients

The natural question is: Can we get smoother maps of attributions that are more inter-pretable? To obtain smoother attribution maps than saliency maps, *SmoothGrad*, introduced in the paper "SmoothGrad: removing noise by adding noise" [184], computes gradients in the vicinity of the input x. It follows three simple steps:

1. Perturb the input x by additive Gaussian noise.
2. Compute the gradients of the perturbed images.
3. Average the gradients.

This gives us slightly less precise local attributions than the vanilla gradient (which is as local as possible). It also results in much clearer attribution maps because the added Gaussian noise and the gradient noise cancel out by averaging while the main signal remains in place. Examples are shown in Figure 3.10. Combining gradients of different perturbations can reduce the noise and perhaps allow us to see more relevant attribution scores. Formally,

$$\hat{M}_c(x) = \frac{1}{n}\sum_{i=1}^{n} M_c(x + \epsilon_i) \qquad \epsilon_i \sim \mathcal{N}(0, \sigma^2 I)$$

where care is also taken for each perturbed image $x + \epsilon_i$ to stay in the $[0,1]^{H \times W \times 3}$ space, as we are averaging across normalized saliency maps.
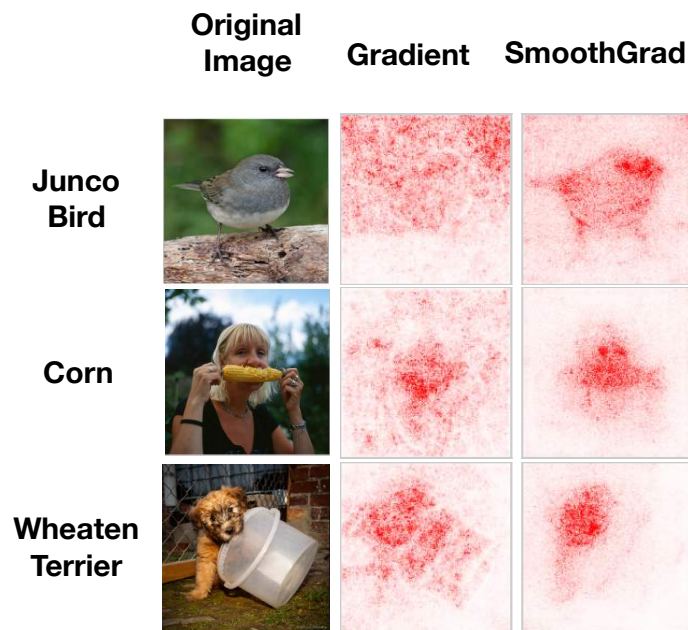
Figure 3.10: Qualitative comparison of SmoothGrad and saliency maps, taken from [2]. SmoothGrad gives attribution maps that are more aligned with human expectations and more interpretable. One has to be careful with confirmation bias, though (Section 3.7.3).

Summary of SmoothGrad

With "less local" gradients, we obtain

$$\mathbb{E}_z\left[f(x+z+\delta e_i) - f(x+z)\right] \overset{\text{Taylor}}{\approx} \mathbb{E}_z\left[\delta\langle e_i, \nabla_x f(x+z)\rangle\right]$$
$$= \delta\langle e_i, \mathbb{E}_z\left[\nabla_x f(x+z)\right]\rangle$$
$$= \delta\left(\mathbb{E}_z\left[\nabla_x f(x+z)\right]\right)_i$$
$$= \delta\mathbb{E}_z\left[\frac{\partial}{\partial x_i}f(x+z)\right]$$

which measures the contribution of each pixel i with an infinitesimal[7] counterfactual at multiple points $x+z$ around x. This expands the originally very local computation of the gradient to a slightly more global region around x.

3.5.5 Integrated Gradients

We will now go from local changes (simple gradients) to the inputs to more and more global changes in the hope that we obtain more sound attribution scores this way. *Integrated gradients* is the middle ground between local and global perturbations. It averages over local *and* global perturbations instead of perturbing only around a single point. We are linearly interpolating between two points in the input space.

In *Integrated Gradients*, introduced in the paper "Axiomatic Attribution for Deep Networks" [191], we choose a base image that contains no information, x^0, and consider our input image, x. We linearly interpolate between x^0 and x in the pixel space by slowly going from an image with no information (x^0, the *baseline image*) to the original image (x). We do the gradient computation at every intermediate point along the line, then average them (without weights, as the expectation is over a uniform distribution). This *nearly* gives us the integrated gradients method:

$$\mathbb{E}_{\alpha\sim\text{Unif}[0,1]}\left[f(x^0 + \alpha(x - x^0) + \delta e_i) - f(x^0 + \alpha(x - x^0))\right] \overset{\text{Taylor}}{\approx} \mathbb{E}_\alpha\left[\langle\delta e_i, \nabla_x f(x^0 + \alpha(x - x^0))\rangle\right]$$
$$= \delta\langle e_i, \mathbb{E}_\alpha\left[\nabla_x f(x^0 + \alpha(x - x^0))\right]\rangle$$
$$= \delta\left\langle e_i, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0))\, d\alpha\right\rangle.$$

This estimates the pixel-wise contribution with an infinitesimal counterfactual (δ), averaged over an entire line between the original input and the baseline image containing "no information".[8]

However, in the integrated gradients method, the contribution of pixel i is computed as

$$(x_i - x_i^0)\left\langle e_i, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0))\, d\alpha\right\rangle,$$

---

[7]In practice, we just choose a small δ value for the tangent plane to stay faithful to the function. Another choice, as we have seen before, is to consider $\mathbb{E}_z\left[\frac{\partial}{\partial x_i}f(x+z)\right]$ as the attribution score for pixel i.

[8]Again, we can consider the attribution score with or without δ. If one includes it, one must keep it a very small number in practice for the tangent plane to stay faithful to the function. This measures the approximate absolute expected change in the output. If one does not include it (this is the usual choice), the score measures the *relative* expected change in the output.

and we derived

$$\left\langle e_i, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0)) \, d\alpha \right\rangle.$$

We seemingly multiply a nicely motivated formula with pixel values. However, the integrated gradients formulation is actually the "prettier" formula, as it satisfies the *completeness axiom*. If we sum over the contribution of all pixels i, we obtain

$$\sum_i (x_i - x_i^0) \left\langle e_i, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0)) \, d\alpha \right\rangle$$

$$= \left\langle (x_i - x_i^0)e_i, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0)) \, d\alpha \right\rangle$$

$$= \left\langle x - x^0, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0)) \, d\alpha \right\rangle$$

$$= \int_0^1 \left\langle \nabla_x f(x^0 + \alpha(x - x^0)), x - x^0 \right\rangle \, d\alpha$$

$$= f(x) - f(x^0),$$

using the fundamental theorem of line integrals (Definition 2.47) with $r(\alpha) = x^0 + \alpha(x - x^0)$. In words: if we sum the pixel-wise contributions of all pixels (integrated gradients in the ith direction, multiplied by pixel differences), we get the difference between the original prediction and the baseline prediction.

The authors of [191] argue that the completeness axiom is a necessary condition for a sound attribution. This axiom states that pixel-wise contributions for input x must sum up to the difference between the current model output $f(x)$ and baseline output $f(x^0)$. Here, the baseline image is an image without "any information". It represents the complete absence of signal. We measure what kind of additional information we add per pixel on top of this baseline image. The baseline image can be, e.g., an image consisting of noise or a completely black image.[9]

Important downside of a black image baseline. If we choose our baseline to be a black image, black pixels (e.g., pixels of a black camera) cannot be attributed at all, as $x_i - x_i^0 = 0$. This does not seem right. The black pixels of the camera are very likely also contributing to the model prediction of the class camera. This is different from the sparse linear model case: $x = \sum_{i \in S} c_i \phi_i$. There, whenever an input feature ϕ_i was 0 (e.g., a black pixel), it contributed to the prediction with a factor of 0, and this was the *GT contribution* of this feature to the prediction. This was also a sound attribution. DNNs, however, are much more complex, and we no longer have this GT correspondence. Here, it is almost surely the case that the black pixels also contributed to the model prediction of a black camera. This problem is known as the "missingness bias" which we will further detail in later sections. Generally, the choice of the baseline value can be quite important. In many cases, random noise seems to be a better option. For the interested reader, the following resource describes other options for the choice of the baseline.

---

[9]A black image baseline is used in the paper. According to the authors, using a black image results in cleaner visualization of the "edge" features than using random noise.

Results of Integrated Gradients

The paper [191] only provides an empirical evaluation of the method's soundness. Example attribution maps are shown in Figure 3.11. According to the results, the integrated gradients method nicely attributes (focuses) to the actual object regions, whereas gradients alone do not give us the "focus" we would expect.

We as humans deem the results sensible (which coincides with the 'coherence with human expectations' property of a good explanation), as we would also focus on the regions that the method highlights. This is, however, a severe case of confirmation bias. We will discuss such biases in Section 3.7.3.

The attribution maps of the integrated gradients method are certainly more *interpretable* than saliency maps. These show more continuous regions; thus, the explanations are more selective. However, this is just one of the evaluation criteria for a good explanation. The soundness of the explanations is only measured qualitatively, even though quantitative analysis would have been critical.
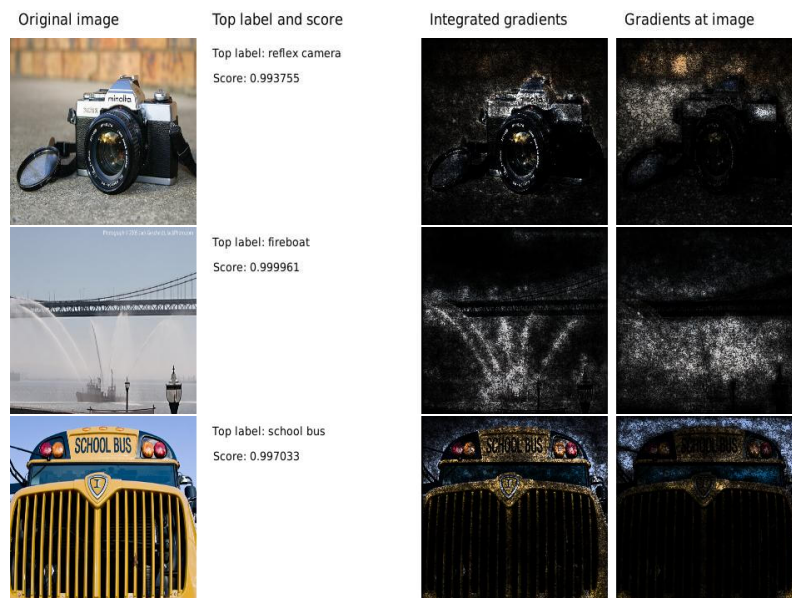
Figure 3.11: Qualitative comparison of Integrated Gradients and saliency maps, taken from [191]. The explanations given by Integrated Gradients are more visually appealing.

3.5.6 Comparing Local and Global Perturbations – Two Ways of Measuring Contribution

We consider two extremes in the domain of local explanation methods that aim to give counterfactual explanations: those that make *local perturbations* to the input x and those that perturb the input *globally*. We also consider an entire *spectrum* between these two extremes. This spectrum is depicted in Figure 3.12.

Local perturbations make very local changes to the input and measure the network's response.

- **Pro**: It has well-understood properties. (The concept of a gradient.) It has no depen-
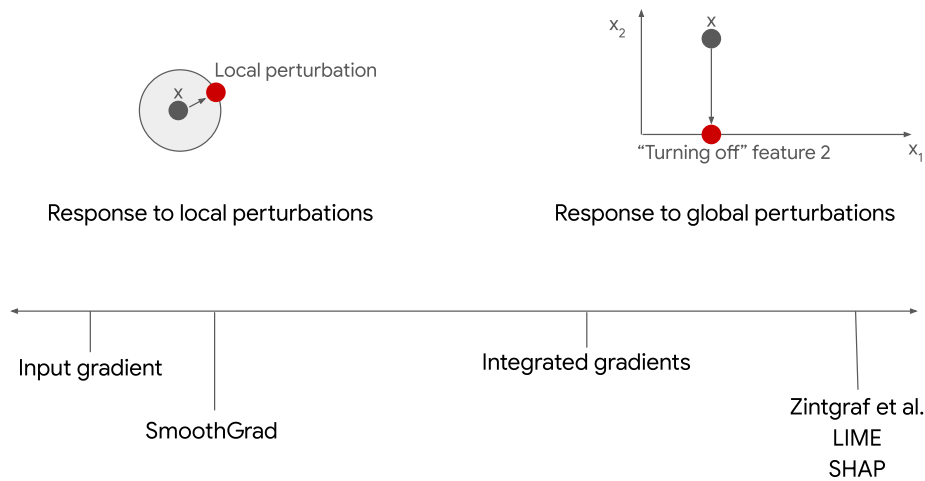
Figure 3.12: Spectrum of local explanation methods w.r.t. the nature of perturbations they employ.

> dence on reference values.
> - **Contra**: We only employ infinitesimal counterfactuals.

Global perturbations measure counterfactual responses by turning off features entirely in various ways.[10]

- **Pro**: This can lead to meaningful counterfactual analysis. This is also a much more natural question to seek explanations for.
- **Contra**: Setting the reference values is hard. Such methods are computationally heavy and need further assumptions/approximations to make them efficient.

The method of integrated gradients gives a smooth interpolation between local changes and turning off features completely. **Note**: These are all still *local explainability methods*. Whether the perturbation is local or global is an independent axis of variation.

3.5.7   Local $=$ Global for (Sparse) Linear Models

Consider a linear model

$$ x = \sum_{i \in S} c_i \phi_i \qquad |S| \ll m. $$

When responding to local perturbations, the gradient of the output w.r.t. the feature i is c_i. When responding to global perturbations, the effect of turning off feature i is $c_i \phi_i$. (Here, we actually set the feature to zero.)

Therefore, the spectrum in Figure 3.12 collapses into a single point for linear models: we do not have any distinction between the two methods. We often try to turn some complex non-linear models into linear ones locally. Therefore, it is of crucial importance to understand linear models.

---

[10]Turning a feature off means that the features receive the baseline value, which is not necessarily zero.

3.5.8   Zintgraf et al.: Inpainting + Black-Box Computation

The Zintgraf et al. (2017) [226] attribution method employs global perturbations – they measure missingness by imputation. It uses the "naive way" of computing the forward pass several times for computing counterfactual attributions. The proposed *prediction difference analysis* reflects the fundamental notion of a counterfactual explanation very well. We want to obtain

$$P(c \mid x_{\setminus i}) = \sum_{x_i} P(x_i \mid x_{\setminus i}) \underbrace{P(c \mid x_{\setminus i}, x_i)}_{\text{trained network}}$$

$$= \mathbb{E}_{P(x_i \mid x_{\setminus i})}\left[P(c \mid x_{\setminus i}, x_i)\right],$$

which is the probability of class c according to the network after removing feature i. As we do not know the true posterior $P(x_i \mid x_{\setminus i})$[11] over the missing feature, we approximate it using an inpainting model

$$Q_{\text{inpainter}}(x_i \mid x_{\setminus i}).$$

Therefore,

$$P(c \mid x_{\setminus i}) \approx \mathbb{E}_{Q_{\text{inpainter}}(x_i \mid x_{\setminus i})}\left[P(c \mid x_{\setminus i}, x_i)\right]$$

$$\approx \frac{1}{M}\sum_{m=1}^{M} P(c \mid x_{\setminus i}, x_i^{(m)})$$

where $x_i^{(m)} \sim Q_{\text{inpainter}}(x_i \mid x_{\setminus i})$. Finally, we calculate the counterfactual before and after removing feature i using the *weight of evidence* value:

$$\text{WE}_i(c \mid x) = \log_2(\text{odds}(c \mid x)) - \log_2(\text{odds}(c \mid x_{\setminus i})),$$

where

$$\text{odds}(c \mid x) = \frac{P(c \mid x)}{1 - P(c \mid x)}.$$
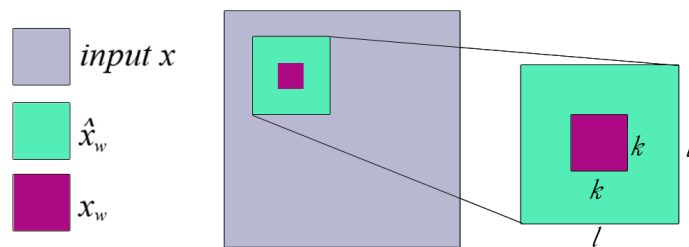
Figure 3.13: Illustration of the conditional independence assumptions used by Zintgraf et al. to make the conditioning tractable. A patch of size $k \times k$ only depends on the surrounding pixels from an $l \times l$ patch that contains the $k \times k$ patch. Figure taken from [226].

---

[11] The works considers image data. $P(x_i \mid x_{\setminus i})$: distribution of feature i given all the other features in the image.

> **Definition 3.10: Mixture of Gaussians**
>
> A Mixture of Gaussians (MoG) distribution with M components is of the form
>
> $$P(x) = \frac{1}{M} \sum_{m=1}^{M} \mathcal{N}(x; \mu_m, \Sigma_m)$$
>
> where μ_m and Σ_m are the mean vector and covariance matrix of the mth component, respectively. The MoG distribution is one of the simplest *multimodal* distributions.

Remarks for Zintgraf et al.

In Zintgraf et al., the features do not have to correspond to pixels. They correspond to image patches in this work.

The weight of evidence is a signed value, as we consider evidence *for* and *against* the prediction. When WE_i is negative for sliding window (image patch) i, it is evidence *against* the model's prediction. It is also often evidence *for* the second-highest scoring class.

To compute the attribution scores, we could use any difference/comparison between $P(c \mid x)$ and $P(c \mid x_{\setminus i})$. The authors argue that using log odds is well-founded.

It is costly to do this procedure for all features i. For each image, one needs to compute N forward operations for the main model + N inpainting computations, where

$$N = \text{number of sliding windows} \times \text{number of samples for inpainting}.$$

The authors propose two methods for estimating the true inpainting distributions $P(x_i \mid x_{\setminus i})$. The first one is to assume *independence* of feature x_i on other features $x_{\setminus i}$. If we make such an assumption, we can consider the empirical distribution of feature x_i from the dataset, i.e., we replace the feature value with a different one sampled from the dataset at random. By sampling more possible feature values from the dataset (at the same image location), we Monte Carlo estimate the expectation. As the authors also state, this is a crude approximation. The second proposal of the paper is to not assume independence but to suppose that an image patch x_i of size $k \times k$ *only depends on the surrounding pixels* $\hat{x}_i \setminus x_i$, where $\hat{x}_i$ is an image patch of size $l \times l$ that contains x_i. An illustration is given in Figure 3.13. To speed things up, the authors used a straightforward method for inpainting: a multivariate Gaussian inpainting distribution in pixel space, fit on dataset samples. In particular, the authors calculate the empirical mean μ_i and empirical covariance Σ_i of the large patch $\hat{x}_i$ on the entire dataset, using the simplifying assumption that the distribution of the large patch $\hat{x}_i$ (i.e., the *joint* distribution of the window we want to sample from and the surrounding pixels) is a Gaussian: $P(\hat{x}_i) = \mathcal{N}(\hat{x}_i; \mu_i, \Sigma_i)$. Finally, the authors use the well-known conditioning formula for Gaussians to obtain $P(x_i \mid \hat{x}_i \setminus x_i)$. Under their assumptions, we have

$$P(x_i \mid \hat{x}_i \setminus x_i) = P(x_i \mid x_{\setminus i}).$$

This is probably the simplest form of inpainting one could think of.

Other possibilities for the inpainting distribution: One could use a Mixture of Gaussians (MoG) or diffusion models [80] for inpainting. However, then it would take even longer to compute the explanation for a single image. There is always a trade-off between complexity and quality.

The method of Zintgraf et al. is a *local* explanation method (as it only gives an explanation for a single image) but a *global* counterfactual method (because the inpainter is allowed to predict anything, not just very small perturbations compared to the original image features). Note, however, that the inpainter is only used to replace small patches – it is still spatially local.

Results of Zintgraf et al.
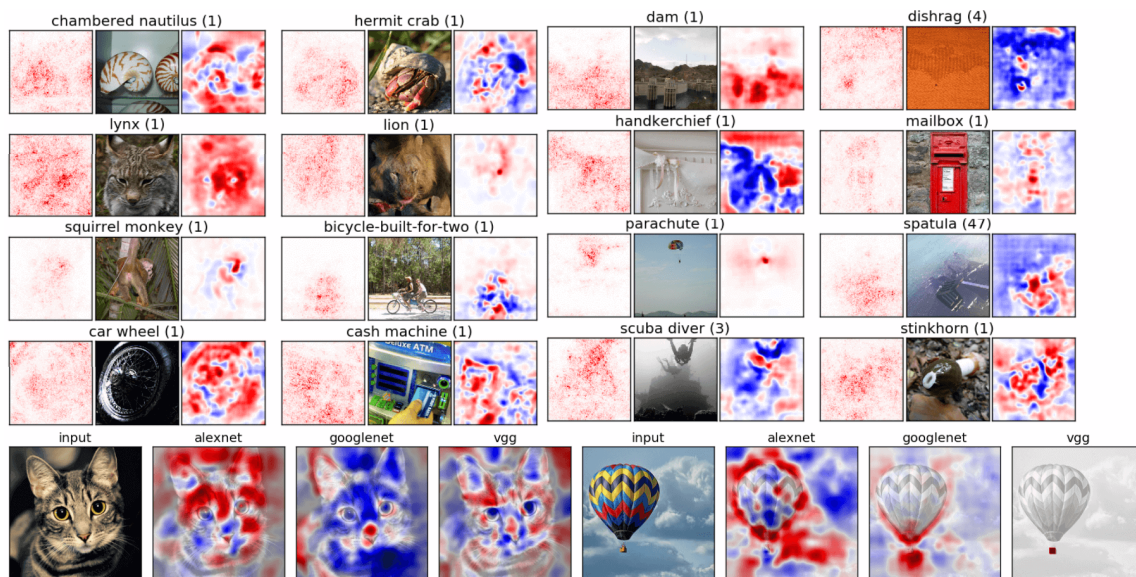
Figure 3.14: Results of Zintgraf et al. (2017) [226], taken from the paper. The attribution maps look surprisingly hard to interpret. Different architectures seem to look at notably different parts of the input image. Still, maybe it *is* the genuine contribution of each feature to the network's prediction. We should not rely too much on human intuition, as that might harm our belief about the soundness of the method. It is hard to say whether this is right or wrong without a quantitative soundness evaluation.

We show a few attribution maps in Figure 3.14. We argue that this is the most promising solution for counterfactual attribution, but it is also the most computationally heavy. Let us now give some pros and cons of the method.

- **Pro**: The method performs a global counterfactual analysis because of inpainting. It is also one of the few datatype-agnostic methods – it can be applied to image, text, and tabular data inputs as well, given that an inpainter is available.
- **Contra**: The method is way too complex to be practical. It also depends on the inpainter, which opens a new can of worms.

3.5.9   LIME: Fitting a Sparse Linear Model

LIME, introduced in the paper ""Why Should I Trust You?": Explaining the Predictions of Any Classifier" [166], has been a popular method for more than five years now that is a bit more realistic than Zintgraf et al. regarding practical use. It builds a surrogate model that is explainable by definition. Given the general formulation

$$\xi(x) = \arg\min_{g \in G} \mathcal{L}(f, g, \pi_x) + \Omega(g)$$

where f is the original model, g is the surrogate model, G is the set of possible surrogate models, π_x is a measure of distance from x used to weight loss terms, and Ω is a measure of complexity. The authors make the following choices: G should be a set of sparse linear models, and Ω should be a sparsity regularizer for the linear model g.

By optimizing the objective function, we try to make g as close to f as possible *in the vicinity* of x, the test input of interest, weighted by π_x, while also keeping it sparse.

In LIME for images, we define

- x as the original image,
- x' as the interpretable version of the original image: a binary indicator vector whether superpixel i is turned on or off (grayed out). Here, all entries are ones.
- z' as a sample around x' by drawing non-zero elements of x' uniformly at random. The number of draws is also uniformly sampled.
- z as z' transformed back to an actual image,
- $f(z)$ as the probability that z belongs to the class being explained, and
- $\mathcal{Z}$ as the dataset of (z, z') pairs.

We specify the sparse linear function g formally by

$$g(z') = w_g^\top z'$$

and the sparsity constraint by

$$\Omega(g) = \infty \mathbf{1}\left(\|w_g\|_0 > K\right),$$

i.e., f should have at most K non-zero weights. The function fitting takes place around input x. We let g follow f via the L_2 loss on the function outputs

$$\mathcal{L}(f, g, \pi_x) = \sum_{(z, z') \in \mathcal{Z}} \pi_x(z) \left(f(z) - g(z')\right)^2,$$

with π_x making sure that we focus on fitting g to f only in the vicinity of x (we only aim for local faithfulness):

$$\pi_x(z) = \exp\left(-D(x, z)^2 / \sigma^2\right).$$

Here D is the cosine distance from x to z if the input is text, or the L_2 distance for images.

An example of the fitting procedure is given in Figure 3.15. The linear model learns to respect local changes of f. This is close to taking the gradient, but we get a sparser linearization than that, which is more interpretable.

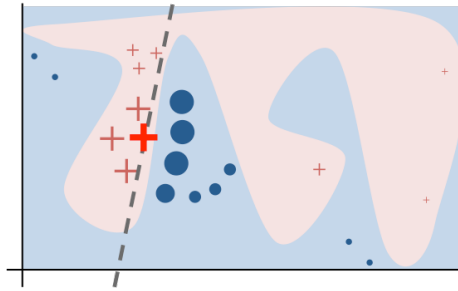The workflow with LIME for images can be explained as follows.

Figure 3.15: Toy example of LIME being fit to the bold red plus data point. The brown plus and blue circle samples are the sampled instances in the vicinity of the input being explained, $(z, z') \in \mathcal{Z}$. Their size encodes their similarity with the original input, as given by $\pi_x(z)$. The background contours encode the decision boundary of the complex model f, whereas the dashed line encodes the decision boundary of g. The surrogate model is locally faithful to the complex model. Figure taken from [166].

(a) Original Image   (b) Expl. *Electric guitar*   (c) Expl. *Acoustic guitar*   (d) Expl. *Labrador*
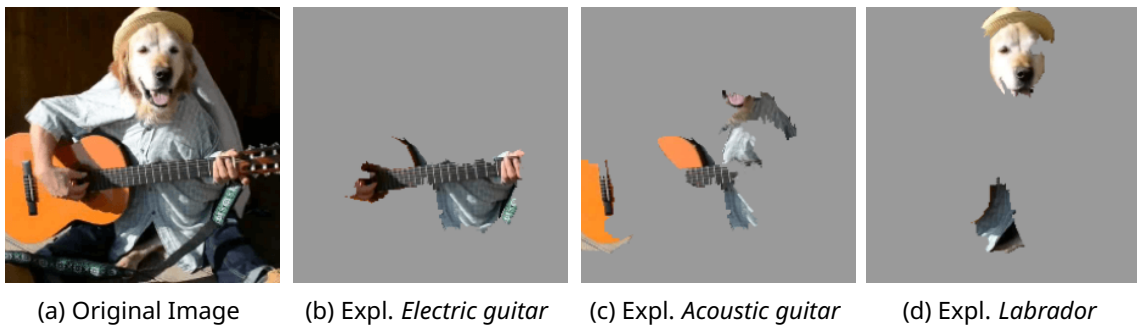
Figure 3.16: "Explaining an image classification prediction made by Google's Inception neural network. The top 3 classes predicted are 'Electric Guitar' ($p = 0.32$), 'Acoustic guitar' ($p = 0.24$) and 'Labrador' ($p = 0.21$)." [166] The results look plausible and easy to interpret, but this should not be taken as an evaluation of soundness. Figure taken from [166].

1. We pick an input x and the class to explain.
2. We train a linear model on top of the superpixel features.
3. We extract the surrogate model weights and check each superpixel's contribution.
4. The superpixel corresponding to the largest weight contributes most to the class prediction in question.

The authors do not only test the method on the actual prediction of the network. They deliberately come up with confusing images with multiple possible classes and try to explain the prediction of the network for the top $k = 3$ predictions. This is shown in Figure 3.16.

Let us discuss the pros and cons of the method.

- **Pro**: The results are interpretable by design.
- **Contra**: (1) We only have a local sparse linear approximation that can be very different from the DNN. (2) The method is expensive, as a sparse linear model has to be fit for all images we want to be explained. (3) The reference image is assumed to be a gray image, an often-used representation of missingness. We discuss in 3.5.11 that this might be suboptimal. (4) The method is not stable. The given explanations (coefficients) are not continuous in the input and are, therefore, not general. In particular, Alvarez-Melis and Jaakkola show in the paper "On the Robustness of Interpretability Methods" [8] that even explaining test instances that are very close/similar to each other leads to notably

different results.

Note 1: The reference cannot be seen in the formulation but rather in how we construct the z samples:

$$z_i' = 0 \iff \text{superpixel } i \text{ is gray}.$$

Thus, we still have an actual 0 value in "interpretable space"; the related term does not contribute to the sum in the linear model.

Note 2: When we give an input to a DNN, we typically subtract the mean of the training set to center the input. So an original image that becomes a 0 input for a DNN is usually gray (the mean of the training set samples, close to constant gray for a versatile dataset). For ImageNet (and many other vision datasets), the standard practice is to subtract the mean.

Additional Information: Surrogate Model

The LIME paper uses $f(x)$ to denote the probability that x belongs to the class being explained. The surrogate g is, however, defined to be a *linear* model that can, in principle, predict any real number and not just probabilities. We could have two other options for defining the surrogate model.

1. Use the *logit* values of the classifier as the targets for the surrogate model. This way, we are matching a real number to another (unconstrained) real number, which seems more natural. However, the coefficients of the surrogate model do not correspond to the changes in the model *output* anymore, but rather to the changes in the logits that are more disconnected from the model's final decision than its predicted probabilities.

2. Constrain the surrogate model's outputs to the $(0, 1)$ range, e.g., by using a logistic sigmoid activation function. This way, we could use any classification loss to train the surrogate model – we are matching probabilities to probabilities. The downside is that the surrogate model outputs are not *linearly* related to the outputs anymore, and the attribution scores become less interpretable.

3.5.10 SHAP (SHapley Additive exPlanations)

The setup of the SHAP method, introduced in the paper "A Unified Approach to Interpreting Model Predictions" [127], is very similar to that of the LIME method in terms of the knowledge about the system and the input/output format. In particular, we assume a black-box system with a binary input vector $x \in \{0, 1\}^N$ that gives us scores $f(x) \in \mathbb{R}$ for a particular class c. We want to assign the contribution of each feature i to the prediction.

The input is represented by a given set of features. The binary membership indicator x is a constant one vector: in the original input, all features are present. For perturbed inputs $z \subseteq x$, zeros and ones indicate whether the corresponding feature is present or turned off in the perturbed image. As we have binary input features, we have a clear interpretation of turning on (1) and turning off (0) features. For images, this is usually a superpixel representation, where the constant one vector is the full image, and the subsets specify which superpixels we switch off (i.e., replace with some base value) and on.

> **Definition 3.11: Combination**
>
> The number of possible ways to choose k objects from n objects is
>
> $$\binom{n}{k} = \frac{n!}{k!(n-k)!}.$$

SHAP determines the individual contribution of each feature i to the prediction $f(x)$ using the notion of Shapley values [177]. The value is defined as

$$\phi_{f,x}(i) = \mathbb{E}_{z \subseteq x : i \in z}\left[f(z) - f(z-i)\right].$$

This value gives the average contribution of feature i in all subset cases to the output of network f. z is a subset of x that must include i. For every subset, we analyze the effect of discarding feature i. This can be thought of as a set function version of the gradient of f at x w.r.t. feature i. The original input x is always treated as $[1, 1, \ldots, 1]$ (all features are turned on), and an example of a valid sample z is $[0, 1, 0, 0, 1, 0]$ for index $i = 2$ if $x \in \{0, 1\}^6$. (The indexing starts from 1.) The possible subsets z are thus any binary vector of the same dimensionality as x. It also follows from the formulation that Shapley values are *signed*, unlike, e.g., saliency maps. Similarly to LIME, we give an attribution score to each feature (e.g., superpixel) w.r.t. turning them on/off (global counterfactual explanation).

Note: The expectation in the SHAP attribution values is *not* uniform across all possible zs that are subsets of x. The expectation follows the procedure below:

1. Sample subset size m from $\mathrm{Unif}\{1, \ldots, |x|\}$.[12]
2. Sample a subset z of size m containing feature i with equal probabilities.

Not every subset across all subset sizes has the same probability of being picked because of sample size differences. If $|x| = 10$, then $\binom{9}{4} \gg \binom{9}{9}$, meaning particular small or large subsets are much more likely than particular medium-sized ones.

Example: Let us consider features as image patches. Suppose that feature i indicates the face region of the cat. To calculate the Shapley value corresponding to feature i, we average the function output for all possible inputs with i switched on (other parts are free to vary), then we *subtract* the average function output for all possible inputs with i switched off. The example is illustrated in Figure 3.17.

We rewrite the expectation as

$$
\begin{aligned}
\phi_{f,x}(i) &= \mathbb{E}_{z \subseteq x : i \in z}\left[f(z) - f(z-i)\right] \\
&= \frac{1}{|x|} \sum_{z \subseteq x : i \in z} \binom{|x|-1}{|z|-1}^{-1} \left[f(z) - f(z-i)\right] \\
&= \sum_{z \subseteq x : i \in z} \frac{(|z|-1)!(|x|-|z|)!}{|x|!} \left[f(z) - f(z-i)\right]
\end{aligned}
$$

by leveraging that the probability of sampling z is equal to the probability of subset size $|z|$

---

[12]Subset z must include i, so the minimal size of z is 1.

Average function output for inputs with feature 🐱   **MINUS**   Average function output for inputs without feature 🐱

Figure 3.17: Illustration of the computation of Shapley values. This is equivalent to the formulation above because the expectation is linear.

times the probability of choosing a particular subset of size $|z|$.

SHAP also satisfies the completeness axiom.

Claim 3.1. *If we sum over the Shapley values for all features i, then we get the difference of the function value for the input of interest x and the prediction for the baseline 0:*

$$\sum_i \phi_{f,x}(i) = f(x) - f(0).$$

Proof.

$$\sum_i \phi_{f,x}(i) = \sum_i \sum_{z:i \in z \subseteq x} \frac{(|z|-1)!(|x|-|z|)!}{|x|!} \left[f(z) - f(z-i) \right].$$

Here, '$\cdot f(z)$' appears $|z|$ times ($|z| \in \{1, \ldots, |x|\}$) with a *positive* sign, once for each feature i in z. Its coefficient is always

$$\frac{(|z|-1)!(|x|-|z|)!}{|x|!},$$

thus $|z|$ times the coefficient gives

$$\binom{|x|}{|z|}^{-1}.$$

Similarly, '$\cdot f(z)$' appears $|x| - |z|$ times ($|z| \in \{0, \ldots, |x|-1\}$) with a *negative* sign, once for each feature i *not* in z. Its coefficient is always

$$\frac{|z|!(|x|-|z|+1)!}{|x|!}$$

as we consider $|z| \leftarrow |z| + 1$ in the formula, thus $|x| - |z|$ times the coefficient gives

$$\binom{|x|}{|z|}^{-1}.$$

The terms of the previous two paragraphs obviously cancel whenever $z \notin \{0, x\}$.

For $z = 0$, $f(z)$ appears $|x|$ times with a *negative* sign. Its coefficient is always

$$\frac{0!(|x| - 1)!}{|x|!} = \frac{1}{|x|},$$

thus, $|x|$ times the coefficient gives 1. Therefore, the term gives $-f(0)$ in the sum.

For $z = x$, $f(z)$ appears $|x|$ times with a *positive* sign. Its coefficient is always

$$\frac{(|x| - 1)!0!}{|x|!} = \frac{1}{|x|},$$

thus, $|x|$ times the coefficient gives 1. Therefore, the term gives $+f(x)$ in the sum.

Finally, by summing all terms up, we indeed obtain $\sum_i \phi_{f,x}(i) = f(x) - f(0)$. $\qquad\square$

Note: The 0 vector can mean arbitrary missingness in the pixel space, just like in LIME. For integrated gradients, we had a very similar result: When we sum over all contributions from every pixel, we obtain $f(x) - f(x^0)$. The difference is that we are not in the pixel space with SHAP.

SHAP satisfies the strong monotonicity property.

> **Definition 3.12: Strong Monotonicity**
>
> Attribution values ϕ satisfy the strong monotonicity property if, for every function f and f', binary input x and feature i, the following holds:
>
> $$f(z) - f(z - i) \leq f'(z) - f'(z - i) \quad \forall z \subseteq x \text{ s.t. } i \in z \quad \implies \quad \phi_{f,x}(i) \leq \phi_{f',x}(i).$$
>
> In words, if the impact of deleting feature i is more significant for f' for all subsets of x containing i, then the attribution value for f' on feature i must be greater than that for f.

The fact that SHAP satisfies the strong monotonicity property follows trivially from its formulation. This seems to be a very reasonable property[13] but should not be deemed crucial. Below, we will see that Shapley values are special for measuring contribution.

Uniqueness: The attribution values ϕ of SHAP are the only ones that satisfy both the strong monotonicity and the completeness axiom [219]. The theorem is well-known in the game theory literature. This roughly translates to: "If we want these nice properties, we must use SHAP." Thus, *SHAP is sufficient and necessary for these two properties to hold jointly.* The coefficients for Shapley values are, therefore, significant to be exactly these.

Why do we want these properties?

Why are strong monotonicity and completeness useful from an applicability point of view? We do not have a strong argument for why this should be the "holy grail" for attribution. The

---

[13]If a function values a feature a lot, then that is also reflected in the Shapley value.

paper also does not give a strong reason why these properties should be strongly connected to any real-world properties. Such works that are built upon axiomatic foundations that introduce some intuitive requirements (e.g., strong monotonicity or completeness axioms) usually conclude that the only method that satisfies all the axioms is theirs. But they usually take *different axioms*, which results in different formulations. The integrated gradients method is also a unique formulation that satisfies a different set of axioms [191]. Everything depends on how we choose these axioms. We do not think that any of the axioms are *absolute necessities*. They are just one way to connect possible real-world needs to an actual explanation method we wish to have.

Using SHAP in practice

We approximate the Shapley values by sampling the expectation at random, according to the coefficients (choose size uniformly, choose a set of that size uniformly). This avoids traversing through the combinatorial number of subsets but introduces large variance in the Monte Carlo approximation, leading to a decreased trustworthiness of the attribution scores.

Let us consider the pros and cons of SHAP.

- **Pro**: Similarly to LIME, the results are interpretable by design. The method also gives global counterfactual analysis.
- **Contra**: (1) We have to use efficient approximations of the Shapley values to keep tractability. Depending on the variance of our approximations, the results we obtain this way might not be faithful to the true Shapley values. (2) Alvarez-Melis and Jaakkola [8] show also for SHAP that the attribution scores can change significantly in small input neighborhoods. (3) Just like in LIME, the reference image is assumed to be a gray image (the mean of the training distribution) in the paper. This might have unfavorable implications, which we will further discuss in Section 3.5.11.

3.5.11   Defining a Missing Feature

We needed a good definition of "no information" for the methods discussed previously.

In *integrated gradients*, we use black pixels as missing features, which is empirically justified in [191]. This gradually kills information by dimming and considers the effect for each pixel integrated through the procedure.

Zintgraf et al. use inpainting as missing features, which is, perhaps, a more sensible choice to encode missingness than any fixed color.

LIME takes the mean pixel values to indicate missingness (which corresponds to gray pixels for most datasets of natural images).

In *SHAP*, missingness is indicated the same way as in LIME. **Note**: Completeness holds when we consider a 0 *vector*. It can correspond to *any* image. The authors equate that to a gray image, but one could make different choices, such as black/white images or Gaussian noise. The choice of what the 0 vector encodes could also be made arbitrarily for LIME. The integrated gradients method also gives a freedom of choice in designing the baseline image. Usually, the choice is made based on results from cross-validation or qualitative analysis (the latter

often being flawed). It is also important to remark that neither of the methods is restricted to images, and we have to reason about the definition of "missingness" for other kinds of data in the same way. For example, for tabular data, both a zero value and the mean value of the dataset make intuitive sense, but they might give different results.

As discussed previously, we consider inpainting to be the most promising approach to defining missingness. The problem with fixed missing feature values is that they can also carry information (e.g., black pixels on a car or gray pixels on a house, illustrated in Figure 3.18), might matter a lot for the prediction but might not be attributed at all. Such pixel values can appear in natural images, yet they will automatically have a zero value in integrated gradients. This is, of course, problematic. The problem can even arise in CALM or SHAP, though perhaps not as severely as in integrated gradients: If a particular superpixel has the same constant value as the mean pixel, turning it on or off does not have any effect, so the attribution value is necessarily zero.

Using constant pixel values to encode missingness also causes problems when considering soundness evaluation methods such as remove-and-classify, introduced in Section 3.7.7, as it can introduce missingness bias, discussed in Section 3.7.8.

Figure 3.18: Example of black and gray colors – popular choices for encoding missingness – conveying information in images. Choosing *any* fixed color to encode missingness is questionable. The images were generated by Stable Diffusion [167].

> **Additional Information: Inpainting Models**
>
> Language models are also often inpainting models (context prediction self-supervised learning (SSL) objective). To get performant solutions, one needs a huge model. The same goes for diffusion-based inpainting models. They are also huge pre-trained models that can synthesize more realistic images. Inpainting is not as easy as it sounds.

3.5.12 Meaningful Perturbations

Now, we discuss the "Interpretable Explanations of Black Boxes by Meaningful Perturbation" [53] paper that introduces *meaningful perturbations*. Instead of different colors encoding missing features, one can also use *learned blurring*. Image blurring can erase information without potentially introducing some. (However, for humans, it might not be enough. Con-

sidering an image of a person playing the flute, even if we blur the flute out, a human still knows what is in their hands. However, in this paper, the authors demonstrated that DNNs do not work like this, as shown in Figure 3.19.)
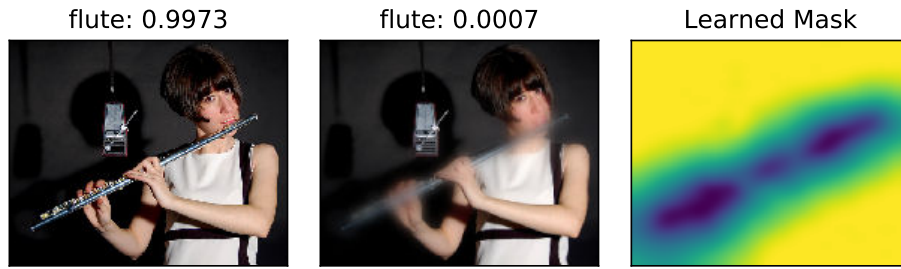
Figure 3.19: Example of a learned blur that results in diminished predictive performance, taken from [53].

The authors are optimizing for the *blur mask*. After the optimization, the final blurred region is ideally the most important region for predicting the corresponding label. The optimization problem is

$$m^* = \underset{m \in [0,1]^\Lambda}{\arg\min} \, \lambda \|\mathbf{1} - m\|_1 + f_c(\Phi(x_0; m))$$

where

- m: A continuous relaxation of a binary mask that associates each pixel $u \in \Lambda$ with a scalar value $m(u) \in [0,1]$.
    - $m(u) = 1$: We do not perturb the pixel at all.
    - $m(u) = 0$: We perturb the pixel (region) as much as possible.
- m^*: Mask that erases most information from the image while also being sparse.
- $\|\mathbf{1} - m\|_1$: Measures the area of the erased region. As m is continuous (smooth), the magnitude matters. L_1 regularization encourages the mask to be sparse. This can be considered as a relaxation of the NP-hard problem using $\lambda \|\mathbf{1} - m\|_0$ plus $m \in \{0,1\}^\Lambda$.
- f_c: Classifier score for class c. We want to minimize this in a regularized fashion.
- $\Phi(x_0; m)$: The perturbation operator, e.g., blurring of original image x_0 according to the mask m:

$$[\Phi(x_0; m)](u) = \int g_{\sigma_0 \cdot (1 - m(u))}(v - u) \cdot x_0(v) \, dv$$

    where $\sigma_0 = 10$ is the maximum isotropic standard deviation of the Gaussian blur kernel.

The objective is fully differentiable w.r.t. m; one can train end-to-end with Gradient Descent (GD).

Use cases of meaningful perturbations

After optimization, we can look at the learned continuous mask to see what region(s) have a large effect. This can unveil very interesting properties of our model. For example, to determine whether chocolate sauce is in the image, our model might be looking more at the spoon than the actual sauce (meaning the score decreases more for blurring this region), as depicted in Figure 3.20. Thus, we can even detect spurious correlations with the method. ("Did my model learn the wrong association?") After detection, we can fix them. This is much more direct than the counterfactual evaluation introduced in Section 2.10.
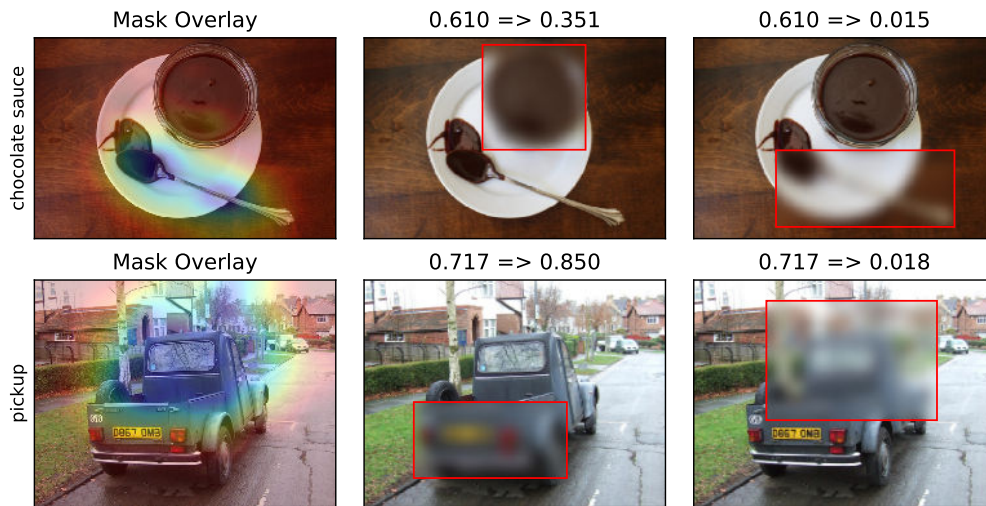
Figure 3.20: One can use the Meaningful Perturbations method to unveil spurious correlations. Figure taken from [53].

Considering the inherent linearity of various XAI methods (3.6), this method does not explicitly give rise to a linear approximation of $f(x)$, but it might be possible to obtain a linear formula in the *transformed* attributions $T(m)$ by embedding them in a non-linear fashion and still keeping them interpretable. Another possibility is that the method linearizes the model's prediction, just not in the attributions but in another property.

3.5.13   Testing with Concept Activation Vectors (TCAV)

Let us go beyond the previous low-level features. We look into higher-level and human-understandable ones because interpretable features are more relevant for most real-life applications. Saliency maps use the gradient directly to attribute to individual pixels. If we look at saliency maps, we usually gain no information about where the important object/region is for a particular label. They are simply too noisy to read and trust and to understand a network's prediction. Even if we choose other pixel attribution methods, these are not interpretable features and do not allow us to relate to more abstract *concepts*. What we really want to ask [98]:

- "Was the model looking at the cash machine or the person to make the prediction?"
- "Did the 'human' concept matter?"
- "Did the 'glass' or 'paper' concept matter?"
- "Which concept mattered more?"
- "Is this true for all other predictions of the same class?"

These are much more semantic questions than the previous methods can handle. This is because while most concepts can be expressed through examples/natural language, they are often impossible to explain in terms of input gradients or more sophisticated scores at the pixel/pixel aggregation level.

TCAV, introduced in the paper "Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)" [98], is a method that allows us to ask whether an abstract concept mattered in the prediction. Figure 3.21 gives an overview of the method

through an intuitive example. We have a classifier with one of the classes being "doctor". We want to know whether some abstract concept was important in predicting $P(z)$, the "doctor-ness". A concept does not have to be an explicit part of training: It can be implicitly globally encoded into the whole image. Instead of relying on gradients/pixel-wise or superpixel-wise attributions, we directly attribute to the human-understandable concept, e.g., woman/not woman.
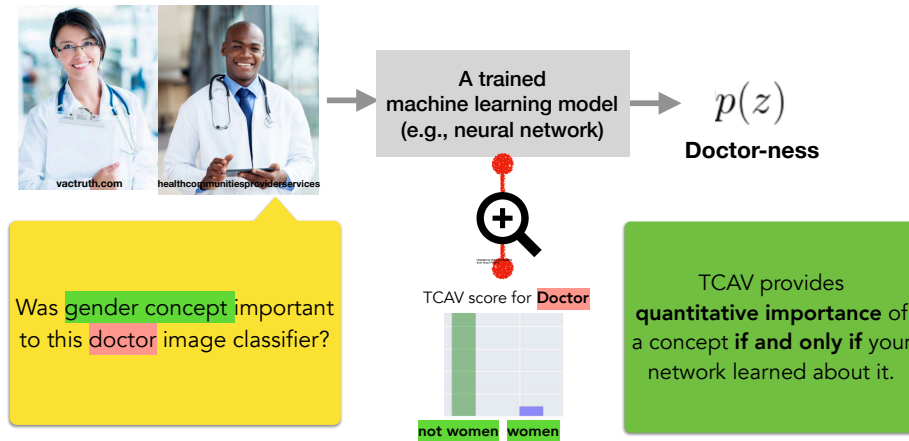
Figure 3.21: Overview of the TCAV method that attributes to human-interpretable concepts. Figure taken from the ICML presentation slides of [98].

Attributing to high-level concepts

Let us first introduce the notation used in the paper:

- C: concept;
- l: layer index;
- k: class index;
- X_k: all inputs with label k (e.g., in the training set).
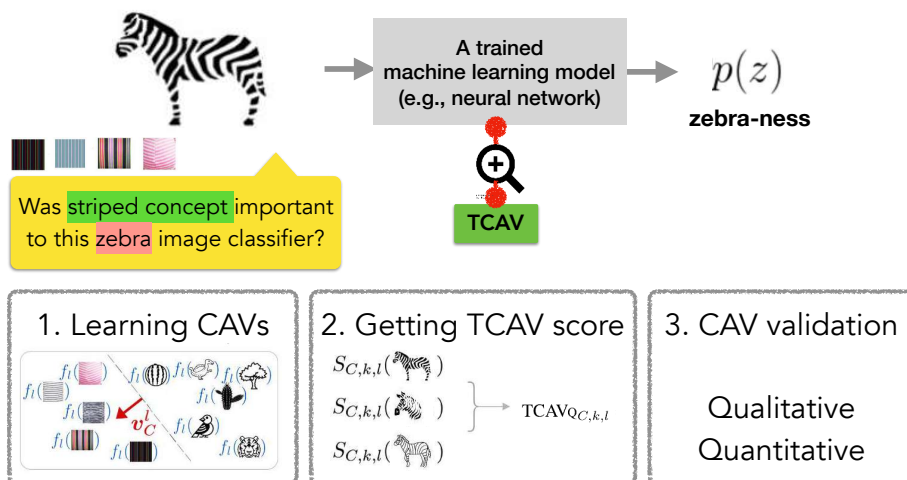
Figure 3.22: Individual stages of the TCAV pipeline, taken from the ICML presentation slides of [98]. Quantitative CAV validation can be performed using statistical testing w.r.t. the set of random samples by validating that the distribution of the obtained TCAV scores is statistically different from that of random TCAV scores. For example, one can use a t-test.

Consider the (already trained) sub-network $f_l : \mathbb{R}^n \to \mathbb{R}^m$ whose output is an intermediate representation of dimension m, corresponding to layer l. We denote the "remaining net" that gives the score to class k by $h_{l,k} : \mathbb{R}^m \to \mathbb{R}$. The method can be summarized as follows (Figure 3.22). We prepare a set of positive and negative samples for the concept (e.g., images containing stripes and other random images). We also prepare images for the studied class (e.g., from the training set). We train a linear classifier to separate the activations of the intermediate layer l between the positive and negative samples for the concept. The Concept Activation Vector (CAV) v_C^l is the vector *orthogonal to the decision boundary of the linear classifier*. This is cheap to obtain: the normal of the decision boundary is the weight vector that points into the positive class. For a particular input x, we consider the *directional derivative of the prediction $h_{l,k}(f_l(x))$ w.r.t. the intermediate feature representation of x, $f_l(x)$, in the direction of the CAV*:

$$
\begin{aligned}
S_{C,k,l}(x) &= \lim_{\epsilon \to 0} \frac{h_{l,k}(f_l(x) + \epsilon v_C^l) - h_{l,k}(f_l(x))}{\epsilon} \\
&= \nabla_{f_l(x)} h_{l,k}(f_l(x))^\top v_C^l.
\end{aligned}
$$

We treat this as the *score* of how much the concept contributed to the class prediction for this particular example. (How would it influence our predictions if we moved a tiny bit in the direction of the concept vector in the feature space?) If the directional derivative is positive, the concept positively impacts classifying the input as the class. Otherwise, the concept has a negative impact.

Finally, the TCAV score for a set of inputs with label k, X_k, is calculated as

$$
\text{TCAV}_{Q_{C,k,l}} := \frac{|\{x \in X_k : S_{C,k,l}(x) > 0\}|}{|X_k|} \in [0, 1].
$$

In words: $\text{TCAV}_{Q_{C,k,l}}$ is the fraction of samples in the dataset with label k where the contribution of the concept was positive for the prediction of the class. This metric only depends on the sign of the scores $S_{C,k,l}$; one could also consider the magnitude of conceptual sensitivities. The TCAV score turns the *instance-specific* analysis ($S_{C,k,l}$, local explanation method) into a more *global* one, for a particular class in general ($\text{TCAV}_{Q_{C,k,l}}$, more global explanation method). It tells us whether the *presence* of the concept is important for a class in general.

TCAV Results

Qualitative results of TCAV are shown in Figure 3.23. TCAV can also shine in medical image analysis, as shown in Figure 3.24. TCAV can streamline the interaction between humans and computers for making predictions.

Let us discuss some pros and cons of the method [142].

- **Pro:** TCAV produces global explanations and can therefore provide insights into how the model works as a whole. It allows users to investigate any concept they define and is, therefore, flexible.
- **Contra:** While the flexibility to investigate user-defined concepts is an advantage, it also has its downside: TCAV may require additional annotation/efforts to construct a concept dataset. Depending on the user's needs, TCAV may not easily scale to many
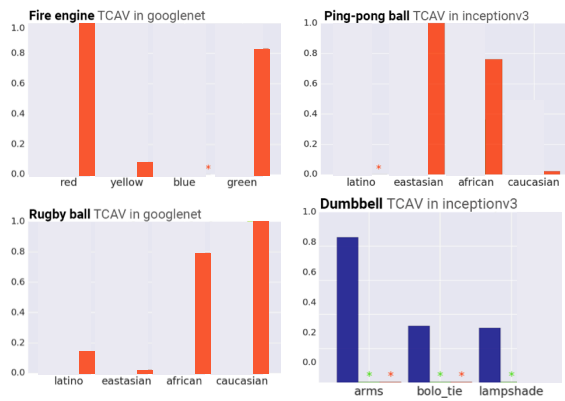
Figure 3.23: Qualitative results of the TCAV method on GoogLeNet and Inception-v3, taken from the ICML presentation slides of [98]. Stars mark CAVs omitted after statistical testing w.r.t. different random images. One can see the concepts the model looks at to make predictions. TCAV can measure how important the presence of {red, yellow, blue, green} color is for the prediction of 'fire engine'. The experiment results show that the red and green colors are important. This signals a strong geographical bias towards countries in the dataset with red and green fire engines. TCAV can also measure how important the presence of different ethnicities is for the prediction of 'ping-pong ball'. The result of the experiments is that the East Asian and African concepts are important. This signals a strong bias towards the ethnicity of players. Agreeing with human intuition, the 'arms' concept is more important for the prediction of 'dumbbell' than the 'bolo tie' or 'lamp shape' ones.
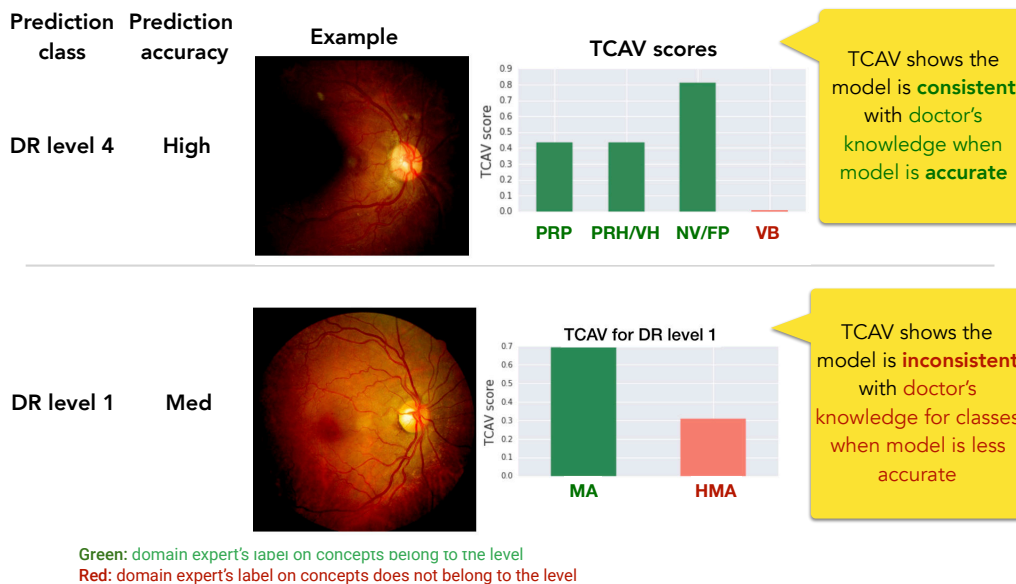
Figure 3.24: Results of using the TCAV method for Diabetic Rethinopathy, taken from the ICML presentation slides of [98]. When the model is accurate, TCAV also shows that it is consistent with the doctor's knowledge: It gives high scores to features deemed by doctors as a precise cause for the prediction. When the model is less accurate, TCAV shows that the model is inconsistent with the doctor's knowledge: It gives a high score to a concept that the doctors deem not helpful to look at.

concepts. Furthermore, TCAV requires a good separation of concepts in the latent space. If a model does not learn such a latent space, TCAV struggles and may not be applicable, as e.g. in shallow networks.

3.5.14 Class Activation Maps (CAM)

CAM, introduced in the paper "Learning Deep Features for Discriminative Localization" [223], is a method that attributes to interpretable intermediate features. A high-level overview of the method is shown in Figure 3.4. CAM employs a typical CNN-based architecture with only a linear operation after calculating the intermediate score map. Up to the score map, the network is very complicated. Afterward, it is just a linear model using Global Average Pooling (GAP) and an intrinsically interpretable linear layer. The key assumption of CAM is that the attribution to pixels in the score map "kind of" corresponds to the attribution to original pixels. This is a huge leap of trust, but CNNs preserve localized information throughout the network (as given by the receptive field of individual neurons). Thus, the explanation w.r.t. the score map also roughly corresponds to the original image. Because of this, we do not have to do linearization for the earlier part of the network to attribute to pixels. We can easily find the pixel in the score map that contributes most to the final prediction. We can also do thresholding w.r.t. the label of choice, and then we obtain a foreground/background mask as an explanation.

Original CAM Formulation

Our training likelihood (or prediction) is

$$P(y \mid x) = \mathrm{softmax}\left(\sum_l W_{yl}\left(\frac{1}{HW}\sum_{hw}\bar{f}_{lhw}(x)\right)\right).$$

(We use NLL to train the model.) We obtain our explanation score map at test time w.r.t. label $\hat{y}$ by using the formula

$$f_{\hat{y}hw} = \sum_l W_{\hat{y}l}\bar{f}_{lhw}(x).$$

That is, we weight each channel of our convolutional feature map $\bar{f}$ w.r.t. the weights between channels l and class $\hat{y}$.

The used shapes of the tensors in the above formulation are $\bar{f}(x) \in \mathbb{R}^{L \times H \times W} = \mathbb{R}^{2048 \times 7 \times 7}$ for the ResNet-50 CAM uses[14] and $W \in \mathbb{R}^{C \times L} = \mathbb{R}^{1000 \times 2048}$ where $C = 1000$ is the number of classes (using ImageNet-1K).

---

[14]Note the low resolution. To overlay the score map on images, further upscaling is needed.

Simplified CAM Formulation

We rewrite our training likelihood as

$$P(y \mid x) = \text{softmax}\left(\sum_l W_{yl} \left(\frac{1}{HW} \sum_{hw} \bar{f}_{lhw}(x) \right) \right)$$

$$= \text{softmax}\left(\frac{1}{HW} \sum_{hw} \underbrace{\sum_l W_{yl} \bar{f}_{lhw}(x)}_{f_{yhw}(x):=} \right)$$

$$= \text{softmax}\left(\frac{1}{HW} \sum_{hw} f_{yhw}(x) \right)$$

where $f(x) \in \mathbb{R}^{C \times H \times W} = \mathbb{R}^{1000 \times 7 \times 7}$. After this, we trivially simplify our explanation algorithm by indexing into our last-layer feature map f that was already calculated in the forward propagation:

$$f_{\hat{y}hw} = \sum_l W_{\hat{y}l} \bar{f}_{lhw}(x).$$

We do not have to do an additional matrix multiplication to generate the score map, i.e., we do not have to do the linear computation twice. We calculate it once for the forward propagation and then reuse the intermediate result for the class of interest by taking the last-layer feature map with channel index = class of interest. This gives us the score map directly. In the original formulation, we first perform GAP, and then use the FC layer during forward propagation. In the new formulation, we have to modify our original model a bit: We exchange the GAP and FC layers and turn the FC layer into a 1×1 convolutional layer.

The FC operation is identical to the 1×1 convolution operation, except that FC operates on non-spatial 1-dimensional features, but 1×1 convolution operates on spatial 3-dimensional features. We apply the same matrix multiplication for every "pixel" ($\in \mathbb{R}^L$) in the spatial dimensions of the feature map. Shape of weights: $\mathbb{R}^{1000 \times 2048 \times 1 \times 1}$. This ResNet [73] variant is fully convolutional. These have been used extensively in the era of CNNs for semantic segmentation. In this case, we are training for a pixel-wise prediction of the class; thus, we need a tensor output. Usually, the exact spatial dimensionality is not retained; we have an hourglass architecture and upscaling at the end [168]. Mask R-CNN [74] also predicts a binary *instance* mask for each detection, and it also has to upscale to the original window size.[15]

We compare the implementation of both approaches. With the simplified formulation, extracting the score map becomes much more straightforward. The two approaches are visualized in Figure 3.25. Python code for the original CAM formulation using PyTorch is shown in Listing 3.1. Similarly, Python code for the simplified CAM formulation is shown in Listing 3.2.

3.5.15 Comparison of the two CAM implementations

Let us consider the pros and cons of the simplified CAM implementation.

---

[15]Nowadays, many people are using Transformer-based [201] baselines for doing semantic segmentation. The convolutional baselines are a bit old-fashioned but are still widely used.

```python
class ResNet(nn.Module):
    def __init__(self, block, layers, num_classes):
        super().__init__()
        self.conv1 = DNN.Conv2d(
            3, 64, kernel_size=7, stride=2, padding=3, bias=False
        )
        self.bn1 = DNN.BatchNorm2d(64)
        self.relu = DNN.ReLU(inplace=True)
        self.maxpool = DNN.MaxPool2d(kernel_size=3, stride=2, padding=1)
        self.layer1 = self._make_layer(block, 64, layers[0])
        self.layer2 = self._make_layer(block, 128, layers[1], stride=2)
        self.layer3 = self._make_layer(block, 256, layers[2], stride=2)
        self.layer4 = self._make_layer(block, 512, layers[3], stride=2)

        self.avgpool = DNN.AdaptiveAvgPool2d((1, 1))
        self.fc = DNN.Linear(512 * block.expansion, num_classes)

    def forward(self, x):
        x = self.conv1(x)
        x = self.bn1(x)
        x = self.relu(x)
        x = self.maxpool(x)

        x = self.layer1(x)
        x = self.layer2(x)
        x = self.layer3(x)
        x = self.layer4(x)

        x = self.avgpool(x)
        x = torch.flatten(x, 1)
        x = self.fc(x)
        return x

    def compute_explanation(self, x, y):
        x = self.conv1(x)
        x = self.bn1(x)
        x = self.relu(x)
        x = self.maxpool(x)

        x = self.layer1(x)
        x = self.layer2(x)
        x = self.layer3(x)
        x = self.layer4(x)  # (1, 512 * block.expansion, w, h)

        weights = self.named_modules(
        )["fc"].weight.data[y, :].unsqueeze(0).unsqueeze(2).unsqueeze(3)
        # (1, 512 * block.expansion, 1, 1)

        return torch.nansum(weights * x, dim=1)  # (1, w, h)
```

Listing 3.1: Naive approach of computing CAM. Obtaining the score map requires complicated indexing.

```python
class ResNet(nn.Module):
    def __init__(self, block, layers, num_classes):
        super().__init__()
        self.conv1 = DNN.Conv2d(
            3, 64, kernel_size=7, stride=2, padding=3, bias=False
        )
        self.bn1 = DNN.BatchNorm2d(64)
        self.relu = DNN.ReLU(inplace=True)
        self.maxpool = DNN.MaxPool2d(kernel_size=3, stride=2, padding=1)
        self.layer1 = self._make_layer(block, 64, layers[0])
        self.layer2 = self._make_layer(block, 128, layers[1], stride=2)
        self.layer3 = self._make_layer(block, 256, layers[2], stride=2)
        self.layer4 = self._make_layer(block, 512, layers[3], stride=2)
        self.conv_last = DNN.Conv2d(
            512 * block.expansion, num_classes, kernel_size=1
        )
        self.avgpool = DNN.AdaptiveAvgPool2d((1, 1))

    def forward(self, x):
        x = self.conv1(x)
        x = self.bn1(x)
        x = self.relu(x)
        x = self.maxpool(x)

        x = self.layer1(x)
        x = self.layer2(x)
        x = self.layer3(x)
        x = self.layer4(x)

        x = self.conv_last(x)
        x = self.avgpool(x)
        return x

    def compute_explanation(self, x, y):
        x = self.conv1(x)
        x = self.bn1(x)
        x = self.relu(x)
        x = self.maxpool(x)

        x = self.layer1(x)
        x = self.layer2(x)
        x = self.layer3(x)
        x = self.layer4(x)  # (1, 512 * block.expansion, w, h)

        x = self.conv_last(x)  # (1, num_classes, w, h)
        return x[:, y]  # (1, w, h)
```

Listing 3.2: Simpler approach of computing CAM. Obtaining the score map becomes trivial.
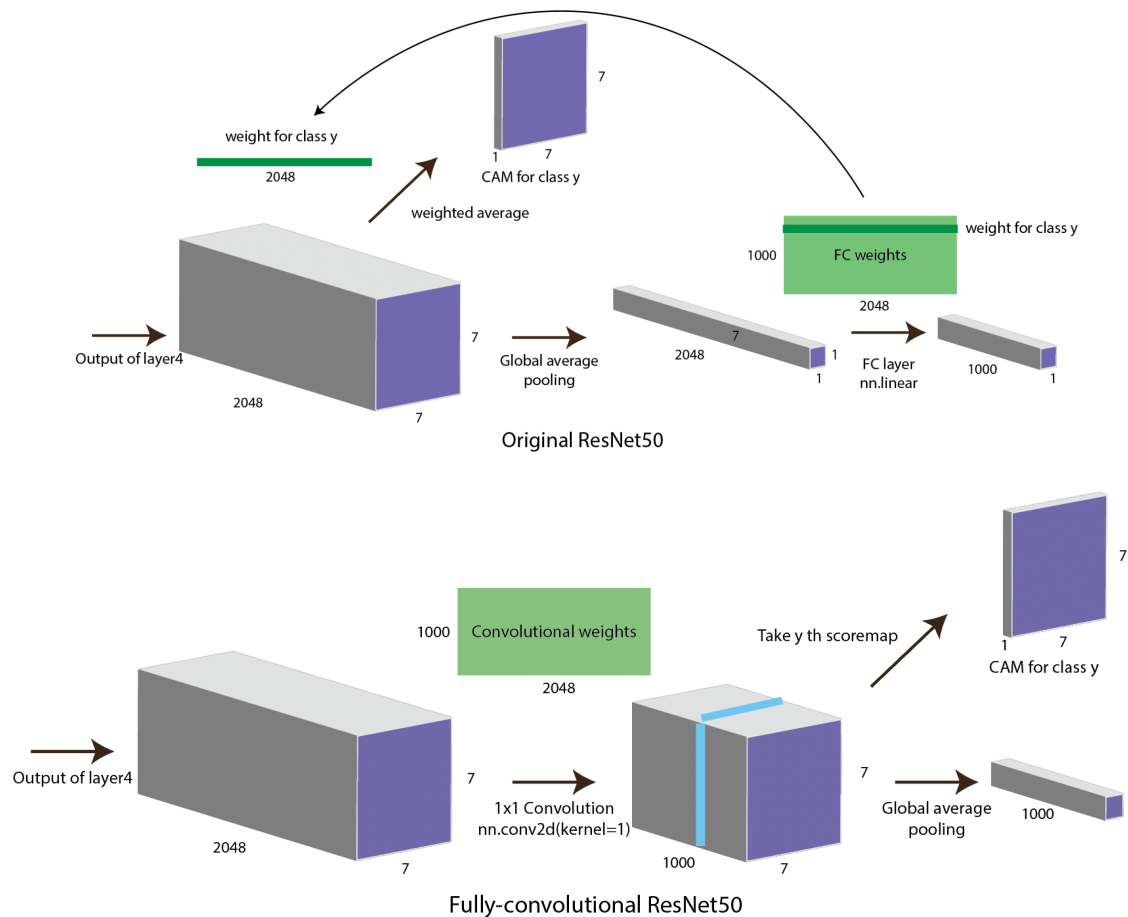
Figure 3.25: Comparison of the two formulations of the CAM method. Extracting the score map corresponding to any of the classes becomes significantly easier when using the bottom formulation. The overhead of having to store the tensor of shape $1000 \times 7 \times 7$ in memory is negligible.

- **Pro**: Simpler implementation (especially for CAM computation) without changing the model performance (confirmed through numerous experiments).
- **Contra**: More memory usage (but negligible). In the original formulation, after we perform GAP, we are left with a 2048D vector. In the simplified formulation we first perform the 1×1 convolution, which results in a tensor of shape $\mathbb{R}^{1000\times7\times7}$. We need to store more floating point values for backprop (and for the CAM computation), but this is negligible compared to the total memory usage of a deep net.

3.5.16   Assumptions to Make CAM Work

As we have seen, CAM assumes an architecture in which we only perform linear operations after computing the score map. There should be a linear mapping from the feature map to the final score. (This does not hold if we also consider the softmax activation, but that is generally considered an interpretable operation, and we usually attribute to the logits.) The GAP operation is just a linear sum of 7×7 features channel-wise, which is very interpretable. (Final prediction can be split into predictions from each of the features of the feature map.

For sums, people have an excellent intuition about what is contributing by how much.)

As discussed previously, another crucial assumption of CAM is that the feature map pixels contain information specific to the corresponding input pixels. We treat each "pixel" of the score map as the feature corresponding to the input pixels at the same spatial location. This was empirically found to be true for CNNs because of translational equivariance. We can trace every feature pixel in the score map back to the possible range of pixels in the input that influenced that feature pixel, called the *receptive field*. This tends to be huge, but the corresponding pixels "move" with the feature pixels via translational equivariance.

Notably, the assumptions only work to some extent. We get *coarse* attribution scores upscaled to fit the input shape. This upscaling (or overlaying) is not theoretically justified, but we still get pretty sound attributions, as measured by soundness evaluation techniques. We can find worst-case examples where (because the receptive field can be much larger than the region one upscaled attribution "pixel" covers) the attribution map does not attribute the pixel responsible for the prediction at all. However, these are pretty artificial examples.

A spatial location of the score map usually has a very large receptive field, especially for a deep architecture like the ResNet-50. When we make use of CAM, we simply upscale it to match the input dimensions. By doing so, the input pixels that correspond to the score map "pixels" according to CAM can be much fewer than the number of pixels in the receptive field of a particular score map "pixel", i.e., the number of pixels that actually influence this score map value. CAM might be localizing more than it should. There is no guarantee that there is a nice straight mapping between the feature map pixels and the raw pixels. Researchers might have gained such insight through semantic segmentation models using a fully convolutional architecture (e.g., DeepLab [34]) where the pixel-wise prediction is directly generated from the feature map (after upsampling). However, such models are trained with pixel-wise supervision, and so they are explicitly instructed that each feature pixel should mostly encode the content of the input area around that feature pixel (not the entire receptive field). CAM is not trained with such a signal – only the aggregation of the pixel-wise predictions is supervised – so there is even less guarantee for the correspondence. In fact, the CAM activation pattern tends to reflect the shape of the receptive field. There exist architectures with non-unimodal receptive fields; for example, DeepLab has hierarchical, checkerboard-like receptive field patterns due to the dilated convolutions and increased convolution strides. As a result, CAM applied to DeepLab shows checkerboard-like activation maps even if the object is at one location of the given image.

The upsampling also introduces various possible problems with interpretability. The upsampling method also influences the attributions (e.g., nearest vs. bilinear vs. bicubic) and detaches the explanation from the extracted feature map values. This is similar to how the normalization of the CAM attributions is detached from the intuitive feature map values and how they relate to the output of the network.

We will see that more sophisticated methods (e.g., CALM [99] in Section 3.5.19) still suffer from the "handwaviness" of the upscaling.[16] In summary, the foundation of CAM is questionable, and the attribution maps should be taken with a grain of salt.

---

[16]In CALM, the intermediate feature map elements also do not have a one-to-one correspondence to the input pixels. As we will see in Section 3.5.19, however, CALM resolves one of the many problems CAM has (namely, the unintuitive normalization of the attribution map).

The validity of the assumption that the feature map pixels correspond to the respective input pixels is unclear for Transformer-based models, as they do not have translational equivariance. It is unclear if the token-wise features after all self-attention layers actually convey the same semantic content as the corresponding tokens at the beginning. Most likely, CAM would not work well for Visual Transformers (ViTs) [48].

> **Additional Information: Attribution Methods for Transformers**
>
> For Transformers, various attribution methods are discussed in the paper ""XAI for Transformers: Better Explanations through Conservative Propagation" [7]. It discusses Generic Attention Explainability (GAE), input $\times$ gradient methods, and several other methods for Transformers.

3.5.17   Grad-CAM – Generalizing CAM to non-linear h
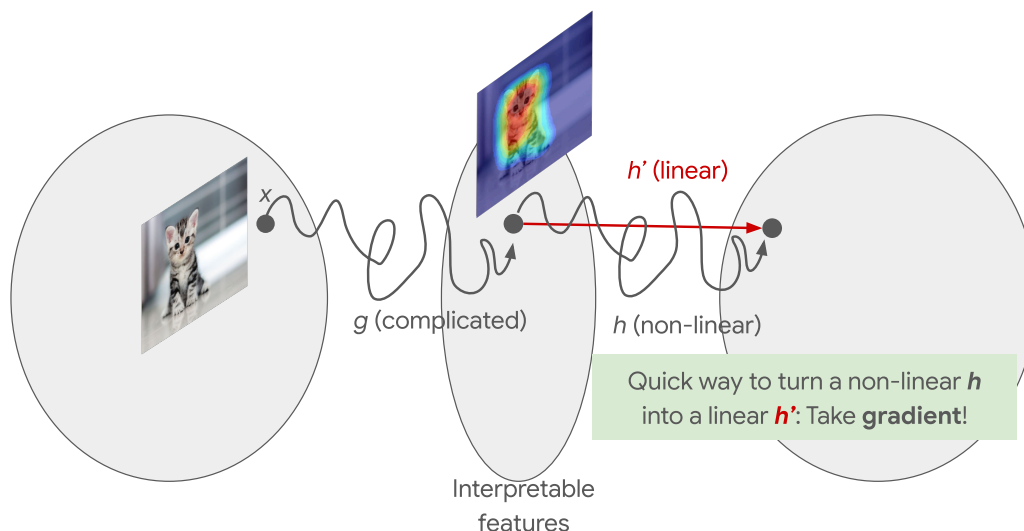
Figure 3.26: High-level motivation of the Grad-CAM method. In Grad-CAM, irrespective of the particular task-specific network we have on top of the convolutional feature representation, as long as it is differentiable, we can linearize it around the feature space point of interest $g(x)$.

What happens if the part of our network between the feature map and the final scores is not linear? This assumption was one of the reasons why we could take an intermediate layer as a feature attribution map in CAM – the remaining layers were linear and, therefore, intrinsically interpretable. Had they not been linear, we could have not applied CAM directly.

Conveniently, we can extend CAM using linearization techniques we have seen before. Grad-CAM [174] is a follow-up method on CAM that extends it to non-linear parts between the feature map and the final scores. A high-level overview and description of the method is given in Figure 3.26. A detailed overview of Grad-CAM is shown in Figure 3.27.

Remarks: It is unclear why the authors apply ReLU on the weighted sum. CAM already performs max-normalization, i.e., they drop all negative values and normalize with the max value anyway. Using the notation of Figure 3.26, instead of having a linear h part, the authors lin-
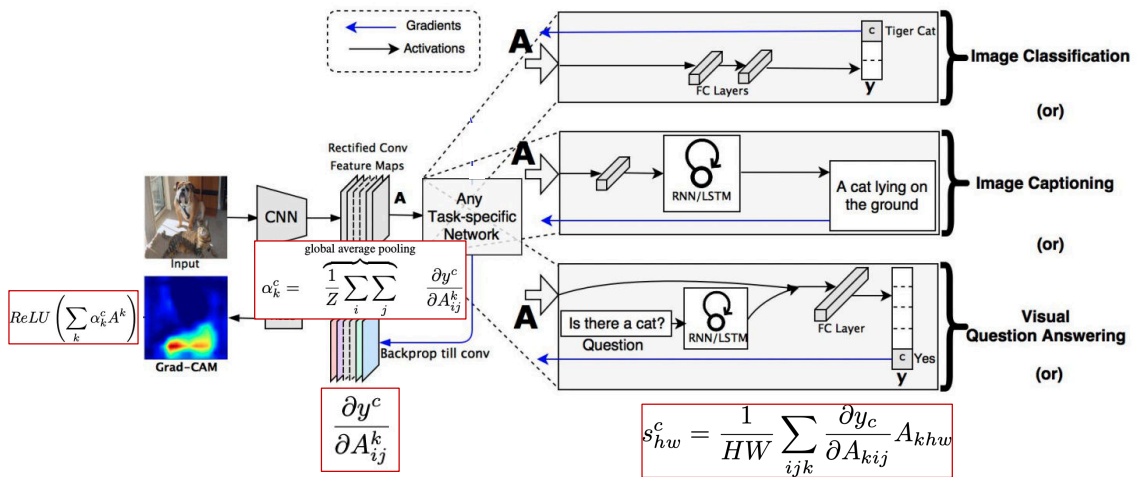
Figure 3.27: Detailed overview of the Grad-CAM method, a generalization of the CAM method that employs linearization. CNN denotes a fully convolutional network like the one we have seen in CAM. Until the *Rectified Conv Feature Maps*, the architecture considered is the same as the one in CAM – $A \in \mathbb{R}^{2048 \times 7 \times 7}$ is the 3D convolutional feature map we previously denoted by $\bar{f}$. Note: A can have other shapes, too – Grad-CAM is not restricted to fixed shapes, and neither is CAM. However, the later layers require some modifications to the CAM method. First, we consider image classification. In VGG, after the convolutional layers, the authors employ three more FC layers with activations. This results in a non-linear second part – vanilla CAM no longer works. In 2017, LSTMs for NLP tasks (including image captioning) were the gold standards. (A lot has changed since, Transformers have taken over the field.) An LSTM also contains many non-linearities – vanilla CAM does not work here, either. To solve this, we linearize the second part of the network architecture for each input x. In particular, we calculate the 3D gradient map of the logit for class c w.r.t. the intermediate representation, $\frac{\partial y^c}{\partial A} \in \mathbb{R}^{2048 \times 7 \times 7}$ (meaning that $\frac{\partial y^c}{\partial A_{ij}^k} \in \mathbb{R}$). Base Figure taken from [174].

earize h locally around $g(x)$ and then compute CAM w.r.t. this linearized network.

Note: People often refer to CAM as Grad-CAM in many cases because the latter is a generalization of the former. When Grad-CAM is mentioned in a paper, it could refer to CAM. It depends on the architecture it is being used on.

Additional Information: Grad-CAM Generalizes CAM

In CAM, we had

$$f_{\hat{y}hw} = \sum_l W_{\hat{y}l} \bar{f}_{lhw}(x).$$

Here, we have a *generalization* of CAM (considering $A = \bar{f}$). Using the notation from Figure 3.27,

$$s_{hw}^c = \frac{1}{HW} \sum_{ijk} \frac{\partial y^c}{\partial A_{kij}} A_{khw}$$

$$= \sum_k \underbrace{\frac{1}{HW} \sum_{ij} \frac{\partial y^c}{\partial A_{kij}}}_{\alpha_k^c} A_{khw}$$

To see that this is indeed a generalization of CAM (with a twist), observe that when we have a linear second part, i.e.,

$$y^c = \frac{1}{HW} \sum_{hw} \sum_k W_{y^c k} A_{khw},$$

then

$$\begin{aligned}
\frac{\partial y^c}{\partial A_{lij}} &= \frac{\partial}{\partial A_{lij}} \frac{1}{HW} \sum_{hw} \sum_k W_{y^c k} A_{khw} \\
&= \frac{1}{HW} \sum_{hw} \sum_k W_{y^c k} \frac{\partial}{\partial A_{lij}} A_{khw} \\
&= \frac{1}{HW} \sum_{hw} \sum_k W_{y^c k} \delta_{lk} \delta_{ih} \delta_{jw} \\
&= \frac{1}{HW} W_{y^c l},
\end{aligned}$$

thus,

$$\alpha_k^c = \frac{1}{HW} \sum_{ij} \frac{1}{HW} W_{y^c k} = \frac{1}{HW} W_{y^c k},$$

which nearly gives us back the CAM formulation but has an additional scaling term $\frac{1}{HW}$. This, however, does not matter for the final activation map because it is normalized. This scaling factor just makes the computation a bit more stable by averaging. This shows that this method is a natural extension of CAM.

3.5.18 Remaining Weakness of CAM

CAM is not as interpretable as we would want. While the function on top of the feature map is linear (GAP + 1×1 convolution), that is not the end of the story. This is because when we compute CAM, we have *an additional step of normalization* of the map to be in the image value range. The unnormalized score map is taken from the pre-softmax values; thus, we have no guarantee of normalization. We do not only need normalization to be in the image value range – a fixed range for the score map is needed anyway, as otherwise, there would be no way to compare score maps consistently across many images. The model we train in CAM is

$$P(y \mid x) = \mathrm{softmax}\left(\frac{1}{HW} \sum_{hw} f_{yhw}(x) \right),$$

which is very interpretable. However, the final score map is calculated in two ways:

$$s = \begin{cases} \frac{\max(0, f^{\hat{y}})}{f^{\hat{y}}_{\max}} & \text{if max} \\ \frac{f^{\hat{y}} - f^{\hat{y}}_{\min}}{f^{\hat{y}}_{\max} - f^{\hat{y}}_{\min}} & \text{if min-max} \end{cases} \in [0, 1]^{H \times W}.$$

These non-linear transformations of our feature map are hard to interpret. In English, the max-version could be explained as

"The pixel-wise pre-GAP, pre-softmax feature value at (h, w), measured in relative scale

within the range of values $[0, A]$ where A is the maximum of the feature values in the entire image."

It is clear that we could not explain this to an end user who has no knowledge of ML. They would not understand what is being shown in the score map, which is a necessary condition of the attributions to be deemed human-understandable.

A summary of problems with CAM as an attribution method is given below:

- *The test computational graph is not a part of the training graph.* In a sense, we are making up values later, at test time, for the score map.
- *We only have an unintuitive description of the score map values in English.* It is difficult to explain the attribution values to clients. Another problem with the normalization method is that min-max or max normalization suffers from outliers without clipping. If one is not careful, whenever the pre-normalized scores contain outliers, the normalized score maps can become uninformative: when visualized, everything seems roughly equally important and the displayed map is not faithful anymore to the actual attribution scores.
- *CAM also violates widely accepted "axioms" for attribution methods.* Details are given in the CALM paper.

3.5.19 Class Activation Latent Mapping (CALM)

To fix the problems introduced in Section 3.5.18, we discuss the paper "Keep CALM and Improve Visual Feature Attribution" [99]. In CALM, we approach the problem with a fully probabilistic treatment of the last layers of CNNs.

Notation:

- X: Input image.
- Y: Class label $\in \{1, \ldots, C\}$.
- Z: Pixel index (location) $\in \{1, \ldots, M\}$ in the feature map. For example, possible values for Z are $1, \ldots, 49$ for a 7×7 feature map. Z is a discrete random variable in the spatial feature map dimensions.

Task: "Predict Y from X by looking at pixel Z." Our prediction is based on the observations at feature location Z. Z is a latent variable not observed during training. Only X and Y are observed; the training set is the same as always. In particular, we do not have GT values for Z.

We use the following decomposition of the joint distribution:

$$P(x, y, z) = P(y, z \mid x)P(x)$$
$$= P(y \mid x, z)P(z \mid x)P(x).$$

This corresponds to the probabilistic graph (directed graphical model) illustrated in 3.28. A detailed overview of CALM is provided in Figure 3.29. We discuss the individual parts of the model here.
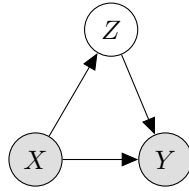
Figure 3.28: Directed graphical model representation of the relationship of variables X, Y, Z in CALM.

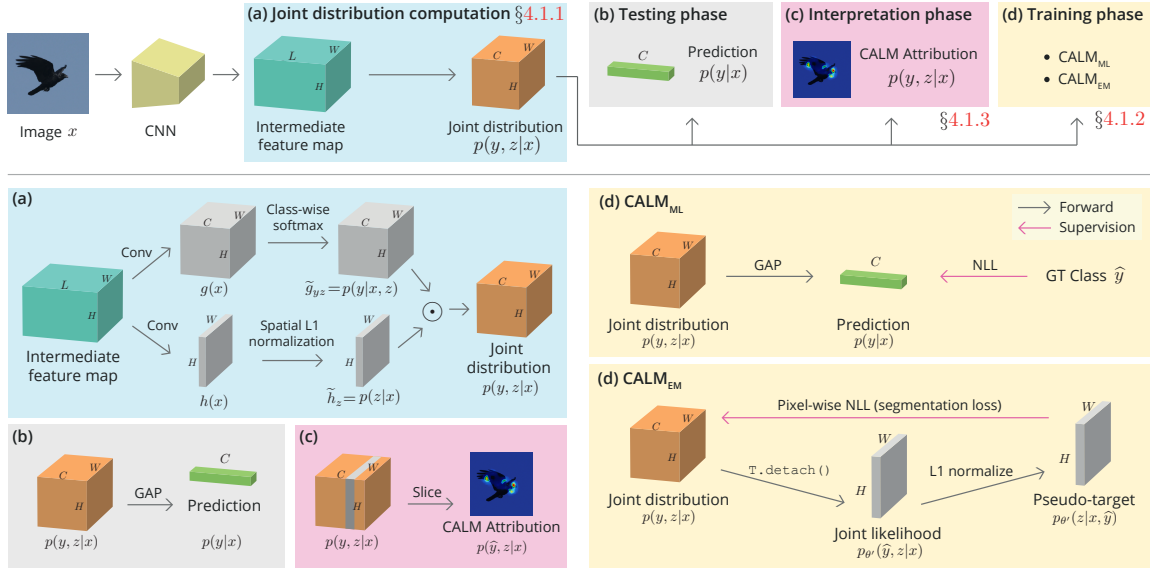**Probabilistic framework for CALM: inference and learning.**

Figure 3.29: Detailed overview of CALM, a fully probabilistic approach to feature attribution. The CNN used is an FCN, just like in CAM and Grad-CAM. Figure taken from [99].

Part (a) of Figure 3.29

We obtain the conditional joint distribution of Y and Z given X. Both Y and Z are discrete random variables; thus, we can fully represent their joint distribution by a 3D tensor, where z is a 2D spatial index and y is a 1D index. Before spatial L_1 normalization, we apply softplus. $g(x)$ and $h(x)$ are network predictions. The only requirement for the joint distribution is that the values are between 0 and 1, and they sum up to 1. This is enforced by the normalization before the element-wise multiplication. We could also just apply global softmax on the entire $g(x)$ pre-activation tensor that normalizes w.r.t. both the class and spatial dimensions, but it did not perform well in the early experiments, according to the authors. Softmax and softplus + L_1 norm are very similar: both are eventually L_1 normalization, but softmax exponentiates before normalizing and softplus + L_1 uses softplus before normalization. Exponentiation can sometimes be too harsh because it can blow up high values to infinity or push low values down to virtually 0. Softplus, on the other hand, is much better behaved – the transformation is approximately linear on the positive side. For this reason, one should always consider using softplus + L_1 norm when softmax blows up neural network training. It would be interesting to observe how turning softmax in Transformers into softplus + L_1 norm influences the behavior of these networks.

Part (b) of Figure 3.29

We obtain the test-time prediction from the network. Similarly to CAM, we do global *sum* pooling. We sum instead of averaging because the elements are probabilities, and this corresponds to marginalization.

Part (c) of Figure 3.29

We obtain the attribution map from the network for a particular input x. $\hat{y}$ is the ground truth label. For a particular location z, the *attribution score* s_z is

$$s_z := P(\hat{y}, z \mid x).$$

In English, the map is significantly simpler to explain than CAM:

> "The probability that the cue for recognition was at z and the ground truth class $\hat{y}$ was correctly predicted for image x."

A nice property of this formulation is that the attribution map is well-calibrated: it lies between 0 and 1 and has a probabilistic interpretation. One can also normalize w.r.t. z and calculate the attribution map w.r.t. the predicted class to get a similar formulation as in CAM. We have a simpler way to compute a calibrated explanation score map.

Part (d) of Figure 3.29

> **Additional Information: DeepLab**
>
> DeepLab [34] is a semantic segmentation network from 2016. It was SotA on the PASCAL VOC-2012 semantic segmentation task at the time of its publication.

We consider two ways of training CALM: Marginal Likelihood (ML) and Expectation maximization (EM). These are typical methods to train a latent variable model. Let us discuss them in this order.

Marginal likelihood. This method directly minimizes the negative log-marginal likelihood. This is the usual way to train when obtaining $P(y \mid x)$ is tractable. The NLL is simply the CE loss

$$-\log P(\hat{y} \mid x) = -\log \sum_z P(\hat{y} \mid x, z) P(z \mid x) = -\log \sum_z \tilde{g}_{\hat{y}z} \cdot \tilde{h}_z.$$

Expectation maximization. Segmentation methods using CNNs use it often. This is exactly how we train a DeepLab [34] model using pixel-wise GT masks. We optimize for the joint tensor. `detach()` is needed to not have any gradient flow from the *target*. We take the GT slice of the joint distribution. It is a likelihood because we apply our knowledge of what the true y is, and it is unnormalized in z. L_1 normalization means dividing by the sum of values in the matrix. The pseudo-target is what we want to reach with $P(\hat{y}, z \mid x)$, as we want only the $\hat{y}$ slice to have a positive probability. Then the joint becomes properly normalized in z when considering $\hat{y}$. We have an entire prediction vector for every pixel in the joint. In our

minds, we expand the pseudo-target into a one-hot vector ($\hat{y}$ dimension is the pseudo-target, all other class dimensions are zeros). Then we apply a CE loss.

CALM Addresses the Limitations of CAM

CALM addresses the limitations of CAM detailed previously:

- *The test computational graph is a part of the training graph.* The training, test, and interpretation phases are all probabilistic.
- *We have an intuitive description of the score map values.*
- *CALM respects all widely accepted "axioms" for attribution methods.* Exact details are discussed in the CALM paper. Being probabilistic, CALM has many linear components.

While this method solves many problems with CAM, it still lacks reasoning about upscaling the score map instead of taking receptive fields into account more rigorously.

Windfall features for CALM attributions
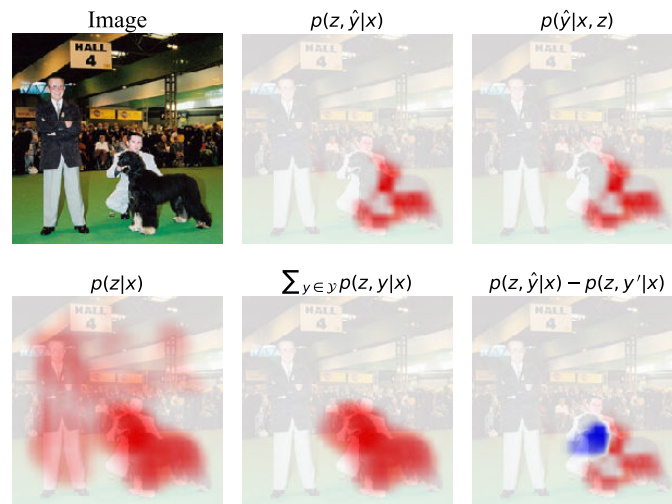
Figure 3.30: Examples of different windfall attributions we can obtain from the joint $P(y, z \mid x)$ in CALM, taken from the paper [99].

CALM comes with numerous windfall gains.[17] When the attribution map is well-calibrated and probabilistic, we can compute a lot of derivative[18] attributions on top of it, as illustrated in Figure 3.30. Score maps can be given, e.g., for

- the GT class (first row, second image);
- the likelihood of the GT class (first row, third image – the difference is the normalization factor);
- the predicted class (not shown);
- a generic class (not shown);
- all classes (second row, first image);
- multiple classes (second row, second image);

---

[17]Unexpected, large gains.
[18]What we mean by "derivative" is not mathematical derivatives but computations that are derived from the probability tensors.

- and counterfactuals (second row, third image).

We discuss some of the options in detail below.

Marginalizing out all classes allows us to gain an overview. "Where is any object that belongs to the 1k classes in ImageNet-1K?" (Only a somewhat valid interpretation when the network is void of any spurious correlation, but even then, its prediction might only depend on small object parts that are very predictive.) "What image regions does the network attribute to any of the classes?" (Valid interpretation for any network.)

We can also sum scores for a subset of classes (e.g., dog, living thing, equipment, object, edible fruit, or food). Here, we sum up score maps for all dog classes (118 in ImageNet-1K). We get better-delineated boundaries for the dog meta-class.

Subtracting different score maps gives a counterfactual explanation of why we chose a class over another. The score map still makes sense; we just use different colors for the two classes.
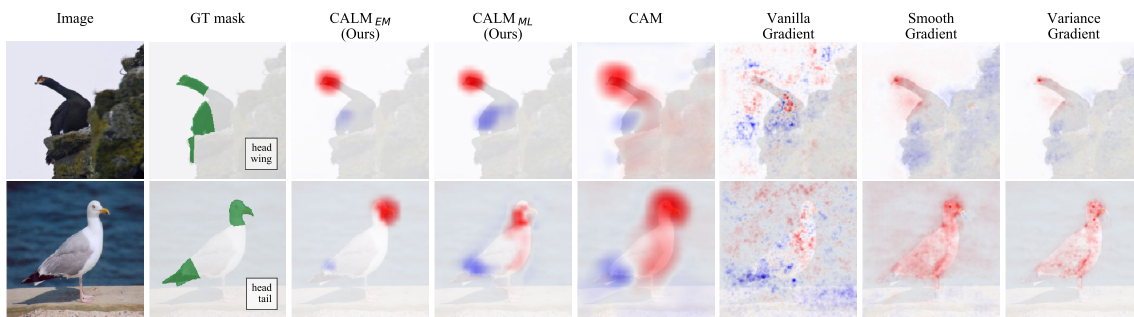
Figure 3.31: Qualitative comparison of CALM and other attribution methods against the GT CUB annotations, taken from the paper [99]. In detail, the authors select the ground truth class and one that is easily confused with it (i.e., the differences appear only on a few body parts of the bird species). They want the model to give the same attributions to the body parts where the classes' (birds') attributes are the same.

Let us now turn to Figure 3.31. One can evaluate the quality of the attribution maps on the CUB dataset as follows.

"We compare the counterfactual attributions from CALM and baseline methods against the GT attribution mask [on CUB]. The GT mask indicates the bird parts where the attributes for the class pair (A, B) differ. The counterfactual attributions denote the difference between the maps for classes A and B: $s^A - s^B$. [...]" [99]

The corresponding results are shown qualitatively in Figure 3.31 and quantitatively in Table 3.1. One possible problem with the evaluation in Figure 3.31 is that the attribution maps that are compared are $P(z, \hat{y} \mid x)$ and $P(z, \tilde{y} \mid x)$ (where $\hat{y}$ and $\tilde{y}$ are two similar classes), which are not normalized in z. In particular, they do not even sum to the same value in z, as the predicted probabilities $P(\hat{y} \mid x)$ and $P(\tilde{y} \mid x)$ are never exactly equal for NN predictions. The paper mentions that if a pixel for both classes is equally important, the difference ideally cancels out so the counterfactual attribution map ideally focuses on pixels that affect the two classes differently. But because of the two maps not being on the same scale, even if proportionally some pixel has the same relative importance, the values are not going to cancel. Thus, the plot very likely shows that the individual attributions are already *only* focusing on parts that are discriminative between the two classes. This would mean that the given reasoning for

the feature maps is slightly incorrect. Without knowing the individual attribution maps, the difference is also not very descriptive. For example, for the ground truth class in the above row, the bird's head seems to be a very distinctive factor for the prediction. However, the wing might also be a factor that the network takes into consideration *for* the GT class, but it is certainly taken with a higher attribution value for the alternative class because it is pale blue. Still, we do not know the exact attributions.

CALM gives counterfactual score maps that often coincide with the GT masks on the CUB task. CALM with EM beats CAM on the CUB benchmark, as shown in Table 3.1. Both CALM variants beat other attribution methods. (This is not an evaluation of soundness. The model has all the rights to look elsewhere, e.g., because it suffers from spurious correlations.)

The authors evaluate the soundness of their method using remove-and-classify (Figure 3.32; discussed in Section 3.7.7). CALM performs best and seems to be the most sound. [99]

Table 3.1: Quantitative comparison of CALM and other attribution methods against the GT CUB annotations, taken from the paper [99].

#part differences	1	2	3	
#class pairs	31	64	96	mean
Vanilla Gradient	10.0	13.7	15.3	13.9
Integrated Gradient	12.0	15.1	17.3	15.7
Smooth Gradient	11.8	15.5	18.6	16.5
Variance Gradient	16.7	21.1	23.1	21.4
CAM	24.1	28.3	32.2	29.6
CALM$_{ML}$	23.6	26.7	28.8	27.3
CALM$_{EM}$	**30.4**	**33.3**	**36.3**	**34.3**
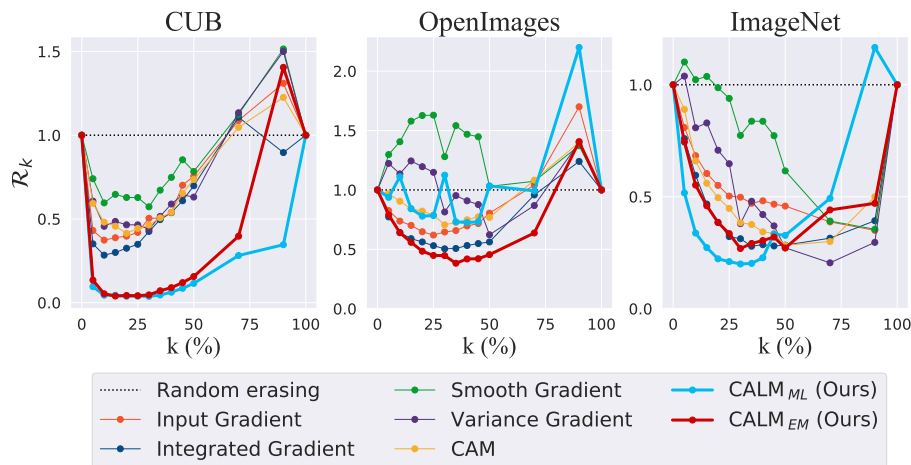
Figure 3.32: "**Remove-and-classify results.** Classification accuracies of CNNs when k% of pixels are erased according to the attribution values s_{hw}. We show the relative accuracies $\mathcal{R}_k$ against the random-erasing baseline. Lower is better." [99] CALM performs well on the remove-and-classify benchmark. Figure taken from [99].

Cost to Pay in CALM

CALM is clearly a better explainability method than CAM but is not necessarily a better classifier. CALM is changing the network structure, so it is very different from the reformulation

of CAM. There, we had an equivalent formulation: The ResNet-50 architecture is fully compatible with CAM. Here, we do not have such an equivalent formulation. Now if we change the original network structure to the CALM formulation, we are changing the mathematical structure of the model. We cannot expect the same accuracy.

As shown in Table 3.2, CALM ML sometimes gains a few points of accuracy and sometimes loses a few against CAM. CALM EM sometimes becomes much worse than CAM in accuracy, sometimes stays close to CAM, and is only behind by a few percentage points. One has to be careful about the possible accuracy loss with CALM.

There is an inborn trade-off between interpretability and accuracy. The existence of this trade-off is very curious: it also means that depending on our actual needs, we might need to choose a different model. For example, losing 4% accuracy might not be as important as gaining confidence and a better picture of how our model works. For such applications, we probably need CALM-trained models. Why must there be a trade-off between accuracy and the model's ability to explain itself? Because we limit ourselves to a smaller fraction of models if we are confined to interpretable models.

There are diverse requirements for deployment. We need to develop more diverse types of models. We should not only aim for models that perform well on a validation set but also develop slightly suboptimal models that are, e.g., interpretable or generalize very well to unseen situations. As an attribution method, there is room for improvement for CAM. CALM improves upon CAM regarding explainability. The better interpretability of CALM also contributes to better Weakly Supervised Object Localization (WSOL) [99], even though WSOL is not precisely aligned with explainability. Better interpretability, however, comes with a cost to pay (accuracy).

The human-interpretability of an XAI method does *not* mean that we wish to make the model recognize things as humans do (human alignment). Instead, we wish to present the behavior of the model in a form that is *understandable* by humans. No model will "start thinking like humans" by using human-interpretable XAI methods. There is no human alignment involved in the above reasoning. These two are also orthogonal axes of variation: the model might make decisions just like humans do, but the XAI method might fail to capture this. Vice versa, the XAI method might *show* that the model makes decisions in a very human-aligned way, but it might just be because of the poor soundness of the method. However, there *are* also occasions where we want better human alignment even at the cost of some loss in accuracy – e.g. when the model is helping experts. This is a different trade-off, namely, the alignment-accuracy trade-off.

Table 3.2: Classification accuracies of the Baseline (ResNet-50), CALM ML and CALM EM, taken from [99]. Both formulations of CALM result in decreased accuracy in most situations. These can also be quite severe: The accuracy of $CALM_{ML}$ is more than 10% less than that of the baseline on CUB. However, there are also some situations where CALM can increase accuracy: On CUB, $CALM_{EM}$ improves upon the baseline.

Methods	CUB	Open	ImNet
Baseline	70.6	72.1	74.5
$CALM_{EM}$	71.8	70.1	70.4
$CALM_{ML}$	59.6	70.9	70.6

Additional Information: Layer Norm

Layer Norm [15] is a normalization technique that normalizes w.r.t. the mean and variance calculated across the feature dimension, independently for each element in the batch. On the other hand, Batch Norm calculates the mean and variance statistics across all elements in the batch. Layer Norm is widely used in Transformer-based [201] architectures.

Additional Information: Modified Backpropagation Variants

This book does not mention some attribution methods: the class of modified backpropagation variants. There are many such methods:
- LRP – layer-wise relevance propagation [22]
- DeepLIFT [180]
- DeepSHAP
- GuidedBP [187]
- ExcitationBP [220]
- xxBP
- …

They are all based on some form of modification to backpropagation. They modify the gradients and do the backpropagation with some broken gradients. Eventually, we get the attribution in some intermediate feature layer or the input space. Vanilla backprop propagates gradients all the way back to the weights. However, gradients are (1) very local and (2) sometimes the function value is more important than the local variations. For example, when a function is constant, it is still contributing something to the next layer.

We do not deal with them in this book for the following reasons:
- It seems complicated to explain what the explanation shows.
- For new types of DNN layers, one needs to develop a new recipe for modified backpropagation. For example, for Transformers, we sometimes need to skip the layer norms with a straight-through estimator (seen in obfuscated gradients). There is no good intuition yet for how to modify backpropagation correctly across layer norms.
- Results depend on the implementation of the DNN. The attributions we obtain can differ for mathematically equivalent networks with different implementations. For example, we can consider two linear layers without non-linearity in between. For fixed, already trained weights, we (1) multiply them together beforehand, use the resulting matrix in a single linear layer, or (2) keep them separate for modified backprop. The results will differ because these methods modify backprop, and the separate modules are different in the two cases. This is a severe issue. We do not have the uniqueness of our attribution score.
- Caveat: They still show good soundness results, especially for Transformer architectures. We should have more understanding of why or how they work.

3.5.20 Summary of Test Input Attribution Methods

Linear models provide nice contrastive explanations. Therefore, we explored ways to linearize complex models (DNNs).

Local linearization around input x for the entire function f is employed by, e.g., input gradients, SmoothGrad, Integrated Gradients, LIME, and SHAP. The caveat is that it is hard to choose the right way to encode "no information". If we perform global linearization for an entire function, we just obtain a linear function that is interpretable by design but not globally sound.

We also discussed the diversity of features for contrastive explanations. One may use pixels, superpixels, instance segments, concepts (that are high-level, i.e., they cannot be represented as an aggregation of pixels), or feature-map pixels (that are aggregations of receptive field pixels with highly non-linear transformations).

We have no guarantee for Transformers that they see the same influence from the corresponding location of tokens/image patches. We cannot expect CAM to work. (Strictly speaking, we do not even have guarantees for ResNets because the receptive field does not coincide with upscaling the feature map.)

Attribution methods come with various pros and cons (depending on the method), and none of them is perfect.

3.6 Explanations Linearize Models in Some Way

As we have seen, the attributions are often based on some form of linearization of the original complex function (the DNN). This is because sparse linear models are already intuitive for humans. Let us give an overview of previously introduced methods and discuss what linearizations they employ.

3.6.1 Input Gradient

Taylor's theorem tells us that for a differentiable function $f : \mathbb{R}^d \to \mathbb{R}$,

$$f(y) = f(x) + \langle y - x, \nabla_x f(x) \rangle + o(y - x);$$

thus, for very small perturbations around the input x, our function is approximately linear. Taking the first-order Taylor approximation means finding the tangent plane of f at input x.

Note: Only the input gradient linearizes the entire model w.r.t. the input x out of the methods in this overview. In the subsequent cases, we will observe linearization in either the *attributions*, the *discretized versions* of the input, or linearization of parts of the network. The only commonality is that the models are analyzed with some form of a linear model, but not necessarily a linear model *w.r.t.* x. The point we make here is that, despite clever ways to formulate the attributions, none of the discussed methods could eventually avoid borrowing the immediate intuitiveness and interpretability of linear models. It would be an interesting research objective to try to formalize this intuition and show that any reasonable XAI method is inher-

ently linear. (For an example of a weaker result, any method that satisfies the completeness axiom is inherently linearizing the predictions in the attributions.)

3.6.2   Integrated Gradients

This method also turns our model into a linear model for an input x around the baseline x^0. Linearity is not w.r.t. the input x rather w.r.t. the attributions:

$$f(x) = f(x^0) + \sum_i a_i(f, x, x^0)$$

where

$$a_i(f, x, x^0) = (x_i - x_i^0) \left\langle e_i, \int_0^1 \nabla_x f(x^0 + \alpha(x - x^0)) d\alpha \right\rangle.$$

The prediction for the input x equals the prediction for the baseline image x^0 plus the sum of the contribution of each pixel.

3.6.3   LIME

LIME makes an obvious, explicit linearization by approximating our possibly highly non-linear model $f : \mathbb{R}^d \to \mathbb{R}$ by a sparse linear model locally:

$$g(z') = w_g^\top z'.$$

3.6.4   SHAP

SHAP also turns our model into a linear model for a binary input x around the baseline 0 w.r.t. the attribution values:

$$f(x) = f(0) + \sum_i \phi_{f,x}(i).$$

The prediction for the input x equals the prediction for the baseline plus the sum of the contribution of each feature (e.g., superpixel). By turning on/off our features, we regulate whether we include a contribution in the final prediction, so in this sense, this is a linear approximation of our model.

3.6.5   TCAV

In TCAV, we have

$$f(x) = h(z) = h(g(x)),$$

which is illustrated in Figure 3.33.

We take the gradient of the output w.r.t. the intermediate layer

$$\nabla_z h(z)\big|_{z=g(x)},$$

thereby linearizing the second part of our network locally, around $g(x)$.

175
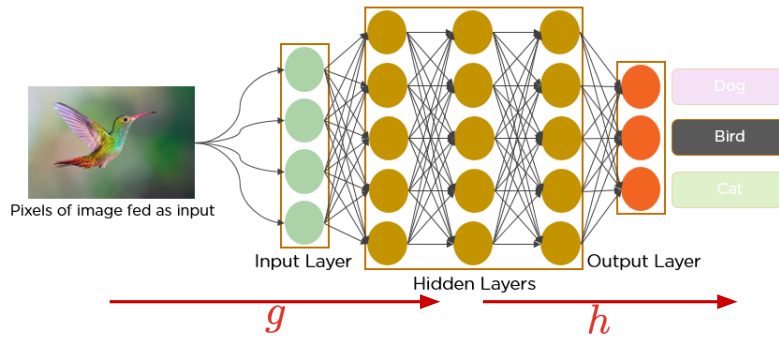
Figure 3.33: High-level overview of TCAV. We linearize the second part, $h(z)$, w.r.t. the intermediate representation, $z = g(x)$, making it a *local* (linear) approximation of the second part of the network.

3.6.6 Different Types of Linearization

We enumerated methods using linearization for explanations. Let us now see what categories of linearization we can establish.

(1) Locally linear around x, completely linear for f.

This scenario is illustrated in Figure 3.34. Examples include Input Gradient, LIME, Integrated Gradients, and SHAP (*even though some of them employ global perturbations*). Here, the entire model is linearized, but only locally. It is quite simple to achieve: Around x, we can explain everything nicely and interpretably.
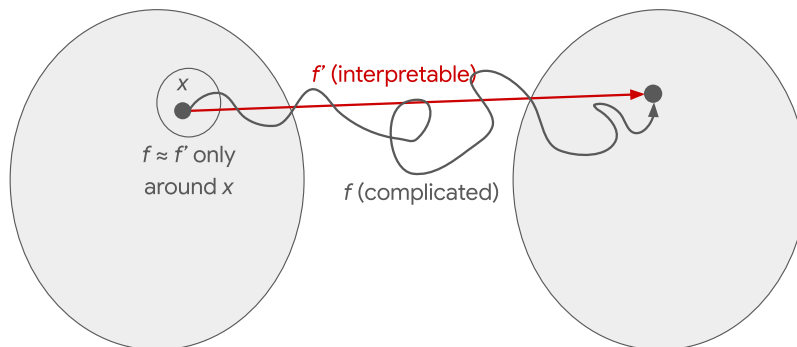
Figure 3.34: Local linearization around x and global linearization in f. This approach is used in the Input Gradient, LIME, Integrated Gradients, and SHAP methods.

(2) Globally linear over g-space, partially linear for f.

This scenario is shown in Figure 3.35. Examples include CAM and CALM. CAM's/CALM's second part is already linear, therefore, there is no need to linearize it. Instead of explaining everything in terms of the input features, we are getting help from the interpretable intermediate features. The second part of the network is naturally interpretable. Therefore, we explain in terms of interpretable features without approximations, under some assumptions.
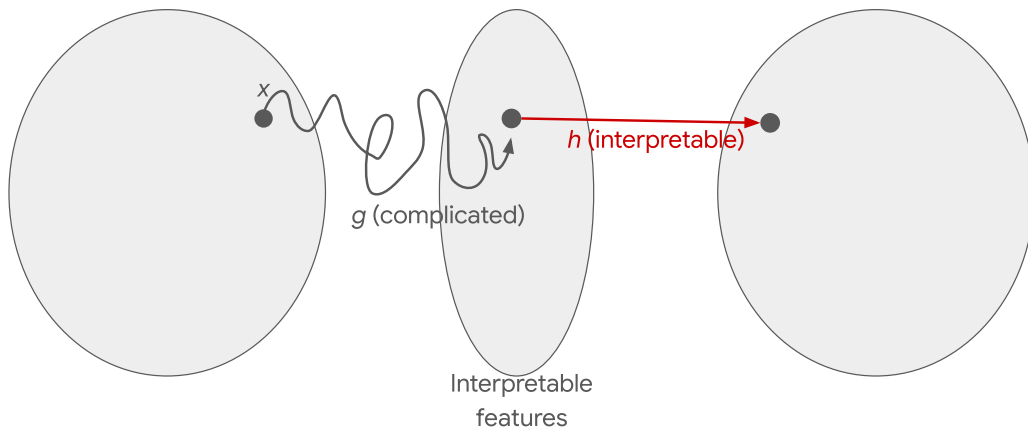
Figure 3.35: Global linearization over g-space, partial linearization for f. Examples of methods employing this strategy include CAM and CALM.

(3) Locally linear around $g(x)$, partially linear for f.

This category is illustrated in Figure 3.36. Examples include TCAV ($S_{C,k,l}$) and Grad-CAM. TCAV takes all x into account that correspond to some label in the TCAV score (quite global explanations), but it linearizes the second part of the network (partial linearization) for each x separately. Whether we perform partial or total *linearization* does *not* depend on whether the method gives local or global *explanations*. In Grad-CAM, we also only linearize the second part of the network (for a single input x).

Instead of explaining everything in terms of the input features, we are getting help from the interpretable features. We are only approximating the second part of our network through gradients. Therefore, we explain in terms of interpretable features but *with approximations*.
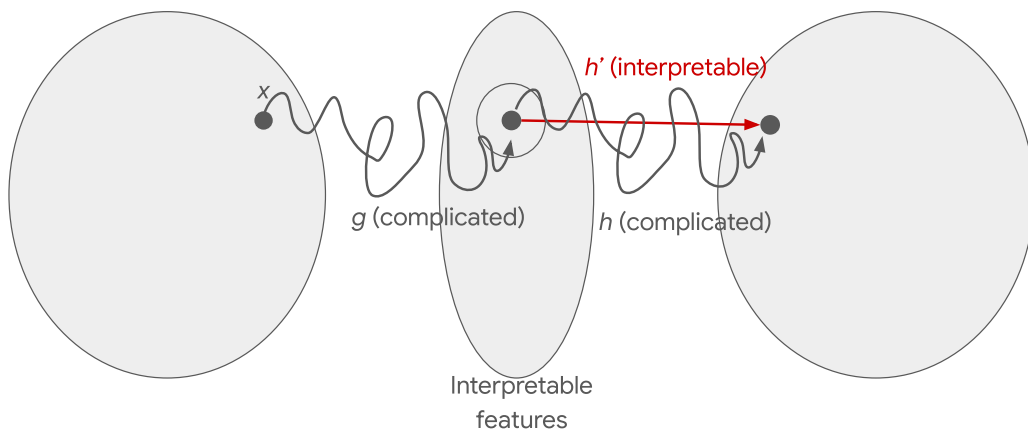
Figure 3.36: Local linearization around $g(x)$ and partial linearization for f. TCAV and Grad-CAM follow this approach.

3.7 Evaluation of Explainability Methods

First, let us discuss why we even need empirical evaluation.

3.7.1 Why do we need empirical evaluation?

The fundamental limitation of explainability is the soundness-explainability trade-off: Our explanation cannot be fully sound and fully explainable. We consider two extremes. The original model is too complex for humans to understand. This is the reason why we needed a separate explanation in the first place. We need simplifications to make humans understand. A global linear approximation makes our model interpretable again, but the model is not the same as before. The soundness of the explanations suffers a lot.

If we look at different XAI methods, they are in the trade-off frontier between soundness and explainability. One cannot say that some explanation method is conceptually perfect by design. Eventually, what matters is *whether the method is serving our need* (the end goal). For that, we need empirical evaluation. We need ways to quantify different aspects of explanations in numbers.

3.7.2 Types of Empirical Evaluation

Doshi-Velez and Kim [47] distinguish three types of empirical evaluation: *functionally-grounded*, *human-grounded*, and *application-grounded* evaluation (Figure 3.37). Let us briefly discuss each of these standard evaluation practices below.
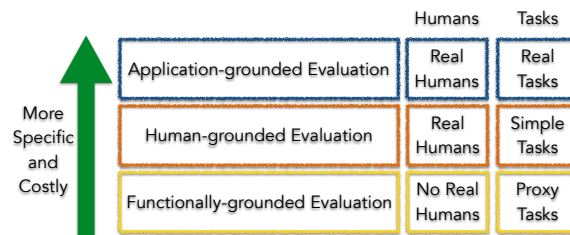
Figure 3.37: Comparison of three types of evaluation methods. As we go from bottom to top, the methods become more aligned with human needs but also become more expensive to carry out. Figure taken from [47].

Functionally-Grounded Evaluation

Functionally-grounded evaluation uses *proxy tasks* to evaluate explanations. Here, no human subjects are required for the evaluation, making this type of evaluation appealing from a time and cost point of view. However, as explainability is necessarily human-grounded, such evaluations should only be considered in addition to human-grounded studies. **Example**: One linear model might be more sparse than another, signaling better human-interpretability [47]. Sparsity can be evaluated without the involvement of humans.

Human-Grounded Evaluation

Human-grounded evaluation considers human subjects but conducts *simple* experiments. This is desired when one wants to evaluate general aspects of the explanation that do not require domain expertise. No specific end goal is considered in such evaluation tasks, but they can still be used to judge general characteristics of explanations. **Example**: Human subjects are presented with explanation pairs and are asked to choose the 'better' one [47].

Application-Grounded Evaluation

In application-grounded evaluation, it is measured how well an explanation method helps *humans* when considering *real* applications/problems. The helpfulness of an explanation method can be quantified by how much it increases human performance on a certain real task. This is the evaluation type that is most aligned with the *human aspect* of explanations, but it is also the most expensive to carry out. **Example**: A computer programmer is evaluated based on how well they can fix their code after being given an explanation.

3.7.3 Soundness Evaluation Techniques

As discussed in Section 3.3.1, soundness (also referred to as faithfulness or correctness of our explanation) is arguably one of the most important and possibly the most widely used criterion. A sound explanation must identify the true cause(s) for an event. Currently, this seems to be the primary focus of XAI evaluation, but it is also *not* the only criterion for a good explanation. This is crucial to keep in mind.

> **Definition 3.13: Confirmation Bias**
>
> Confirmation bias is confirming the performance of our explanation method against what humans think would be the proper attribution instead of investigating further whether the model was actually basing its prediction on these causes.

For measuring soundness, much previous research relied on qualitative evaluation of (potentially cherry-picked) examples. Consider the integrated gradients paper referring to the attribution maps shown in Figure 3.11:

"Notice that integrated gradients are better [than input gradients] at reflecting distinctive features of the input image [for the prediction]." [191]

Can we really conclude that for the images provided? Maybe the integrated gradients method *delineates* the objects better than input gradients, but does that mean they reflect distinctive features for the model's predictions (i.e., what the model is looking at) better? That is an entirely different question.[19] Another claim from the paper:

"We observed that the results make intuitive sense. E.g., 'und' is mostly attributed to 'and', and 'morgen' is mostly attributed to 'morning'." [191]

To humans, this makes perfect sense. However, what if the model looked at a different feature for predicting these words? We argue that this is a case of confirmation bias. If we keep relying on human intuition to measure/evaluate explainability, how could we detect models that rely on new knowledge humans have not learned before? This point of view prohibits us from *learning* from models. Another example from CAM (referring to a bunch of visualizations of different methods, shown in Figure 3.38):

"We observe that our CAM approach significantly outperforms the backpropagation approach [...]" [223]

---

[19]The paper back in 2017 was not rejected for just making qualitative evaluations. The field has grown and matured a lot since then – today, it is always a requirement to provide proper quantitative evaluations.

What do they exactly mean by outperforming? When is an explanation method doing better? Can we really conclude this?[20] This is likely another case of confirmation bias. We can see that qualitative evaluations of soundness are susceptible to confirmation bias. This is made even more severe by the fact that *no GT explanation exists in general*.
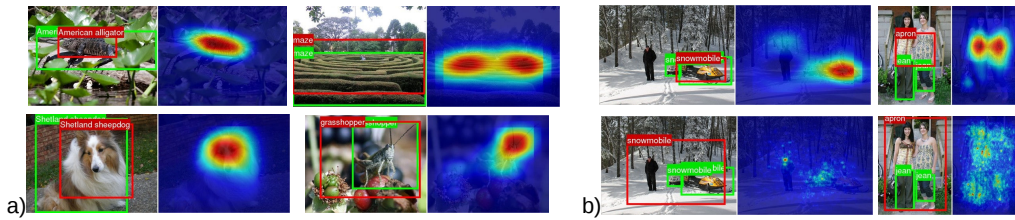

Figure 3.38: "a) Examples of localization from GoogleNet-GAP. b) Comparison of the localization from GooleNet-GAP (upper two) and the backpropagation using AlexNet (lower two). The ground-truth boxes are in green, and the predicted bounding boxes from the class activation map are in red." [223] The authors conclude that "our CAM approach significantly outperforms the backpropagation approach." What they exactly mean by outperforming is not disclosed. In particular, it is questionable if any form of 'outperforming' can be concluded by observing these results. Figure taken from [223].

Does localization evaluation make sense for soundness?

For the quantitative evaluation of CAM, the authors measure the number of times their attribution score map corresponds to the object bounding box. They segment regions whose CAM value is above 20% of the maximum CAM value. Then they take the tightest bounding box that covers the largest connected component in the segmentation map. Finally, they measure the IoU between this box and the GT object box of the class of choice. When IoU $\geq 50\%$, they consider it a success. They measure the success rate on "the" ImageNet validation set. They find that the CAM variants perform better than the backprop variants.

Setting: Explanation method A finds GT object boxes better than explanation method B. Does this mean that explanation method A is working better than B? *We do not think so.* The model may have been looking at a non-object region to make the prediction. If that were the case, the explanation method with a lower localization score might explain the model better.

Takeaway: We should not evaluate according to our expectations when evaluating explanation methods.

How to interpret unintuitive explanations?

Suppose we have a case when the provided explanation differs greatly from our expectations. Does that mean that the explanation method failed while the model was working fine (it was looking at the right thing), or did the explanation method correctly expose a bug in our model (or in the data), like spurious correlation? There is no way to tell these two scenarios apart from a single visual inspection.

---

[20]The authors likely refer to WSOL performance. However, that is just a coarse proxy for explainability methods and does not directly measure the quality of explanations in any way.

Typical pitfalls of soundness evaluation

Soundness aims to evaluate the following: Does the score map $s(f, x)$ represent the true causes for f to predict $f(x)$? The true explanation depends on both the input x and the model f. We can also calculate the attributions for the GT class y or any other class. This is usually done less in practice. Explaining something that has already happened (e.g., f predicted $f(x)$ for x) makes more sense than explaining hypothetical situations. In this case, the true explanation depends on x, f, and y. The problem with qualitative evaluation is that humans also cannot tell what the cues were that f looked at to predict a certain class. We are only looking at x and y to make the evaluation, not f. This seems wrong *by design*.

The problem with localization evaluation is that if we compare to GT localization, we also do not take the model f into account, only x and y. This also seems wrong by design. We know already that models do not always look at foreground cues to predict classes.

The fundamental issue with evaluating soundness is that there is no GT explanation in general.[21] Humans cannot provide GT explanations. This is precisely the reason we are developing an explanation technique in the first place. If there were a GT explanation for a model, then that itself would be a good explanation, and there would be no need to study what a good explanation is and evaluate explanations. We should start from somewhere, but it is hard. We are facing a chicken-egg problem.

3.7.4   Evaluation of Soundness of Explanations based on Necessary Conditions

There is a trick that people consider to test the soundness of explanation methods. We define a few criteria that a successful explanation method must satisfy.[22]

Example: The explanation $s(f, x)$ must not contain *any* information if f is not a trained model (i.e., it is randomly initialized). The intuition is, "How could any explanation contain any interesting information for an untrained model?" Otherwise, our explanation is rather trying to please human qualitative evaluations by producing plausible explanations. Interestingly, a randomly initialized CNN achieves a better score than random guessing with a trained linear layer on top because of inductive biases. On ImageNet-1K, one can achieve 4% accuracy [20].[23] It seems that this is probably a way too strong necessary condition. There can be some information in the score map (Why not? We do not fully know the behavior of a randomly initialized model.), but the main point is that the score map should strongly depend on function f. The explanation does not have to be *informationless* when a model is randomly initialized. If it turns out that the score map is independent of the model altogether, then something is wrong.

A relaxed version of the above is that when the model changes (becomes gradually randomly initialized from a trained model), we should also see notable changes in the attribution map.
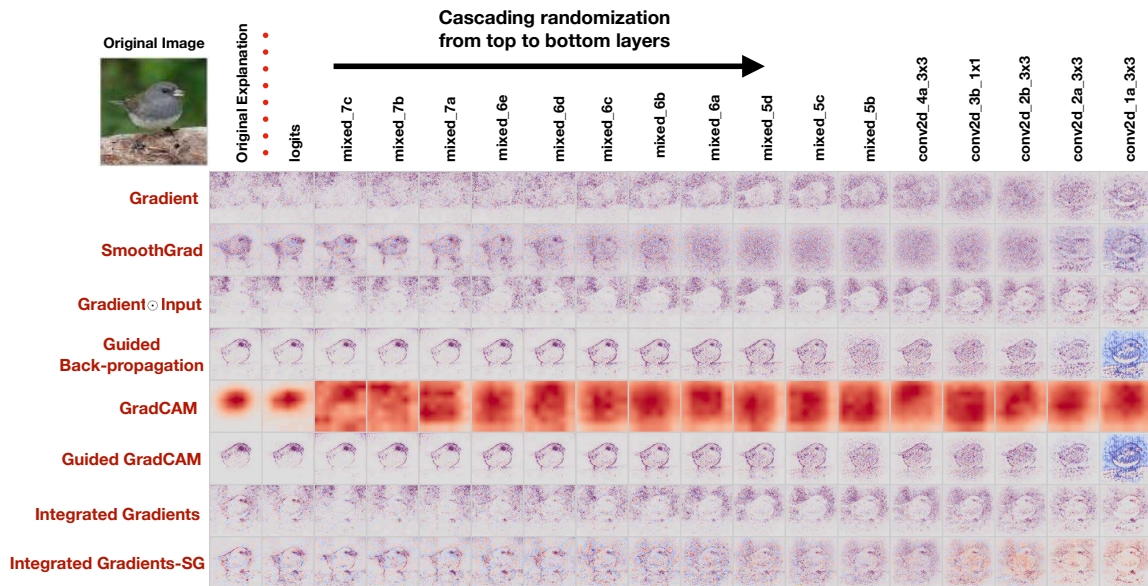
Figure 3.39: Results of various explainability methods on the cascading model parameter randomization sanity check. This sanity check is passed by saliency maps, SmoothGrad, and Grad-CAM. Details are discussed in the text. Figure taken from [2].

3.7.5 Sanity Checks for Saliency Maps

Let us discuss the paper "Sanity Checks for Saliency Maps" [2] where the authors benchmarked various explainability methods on sanity check tasks. We highlight two of these:

1. **Cascading randomization.** As we randomize the network's weights (starting from the latest layers and going toward the input layer), we should see notable changes in the explanations XAI methods give.
2. **Label randomization.** For models trained with randomized labels (that should not learn anything meaningful), XAI methods should not highlight parts of the input that are discriminative for the *original task* (without label randomization). They should return irrelevant attribution maps.

Cascading randomization

Let us discuss Figure 3.39 showcasing *cascading normalization*. Saliency maps (called 'gradient' in the Figure) exhibit large changes in the attribution map. SmoothGrad is also heavily influenced by randomization: it "passes the check." Curiously, the image does not become complete noise from the initially clear attribution map, rather, it becomes a noisy edge detector.) For Gradient ⊙ Input, the outline of the bird is always visible: the changes are not so large.

Guided Backpropagation [187] shows a similar attribution map all the way, even after a global change of the model. It only becomes noisier, the edges are clear all the way. It seems like it does not take model f into account that much. For Guided Backpropagation, the authors

---

[21] The model is, of course, sound to its own behavior. However, we cannot treat a system as its own explanation. That kills the purpose.

[22] Without question. Full stop.

[23] A simple linear classifier cannot perform better than random guessing.

were selling the fact that they get very nice visualizations of objects [187]. This is true, but it does not reflect well what the model is doing. It is close to being an edge detector, but at the time when they published it, it looked like a ground-breaking technology. No one has done it before, and the results looked like the classifier had all the knowledge about where objects are. But even though it looked promising at the time, people have since then realized it does not work. It does not contain enough information about the model. Let us give a rough outline of the Guided Backpropagation method. If we use ReLU, then during backpropagation, when the pre-activation is negative, we do not backpropagate gradients. When it is positive, we just let the gradient through. In Guided Backpropagation, we do not let the gradient through when it is negative. (Like a ReLU on gradients.) This results in an AND condition: if the gradient was positive and the pre-activation was positive, then we let the gradient through. There is no justification for why this should work. And it does not, apparently.

Continuing with previously discussed methods, Grad-CAM showcases large changes in the attribution map as well. Regarding Guided Grad-CAM, we could give the same remarks as for Guided Backpropagation. This is a multiplication of the guided backpropagation score map and the CAM score map. It is natural that the method inherits lots of issues from guidance. Integrated Gradients gives very similar results to Gradient $\odot$ Input. Attributions change only slightly – definitely not as radically as for, e.g., SmoothGrad. Integrated Gradients-SG is very similar to Integrated Gradients, maybe even a bit worse.

Note: Earlier versions of [2] give significantly different results (even more extreme).

This sanity check was set up as a necessary condition: Any explanation method (even the simplest, most naive ones) should satisfy it. If they do not, the method is unusable. It is the bare minimum requirement an explanation method has to satisfy. Nevertheless, some methods already fail to pass this simple test. Namely, Guided BP and Guided Grad-CAM are essentially edge detectors. Gradient $\odot$ Input and Integrated Gradients are also not so convincing.
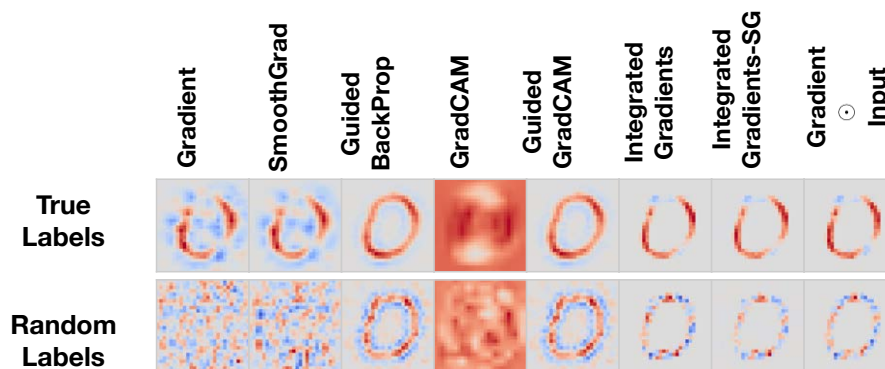
Figure 3.40: Results of various explainability methods on the data randomization (randomized labels) sanity check. This check is also passed by saliency maps, SmoothGrad, and Grad-CAM. Details are discussed in the text. Figure taken from [2].

Label randomization

We discuss the other aforementioned sanity check the paper considered, which uses *random labels* named 'data randomization test'. The results can be seen in Figure 3.40. One can also compare explanations for two models trained on MNIST with true (original) labels or random labels (control group). Random-label-trained models should return explanations without information. With Guided Backpropagation, we can clearly see the shape of 0 for random labels; it seems to give edge detection regardless of the label used for training. Guided Grad-CAM also fails the test again. Methods depending on pixel values tend to show a "0" shape even for random label models: Integrated Gradients[24], Integrated Gradients-SG, and Gradient ⊙ Input all showcase the same problem. Gradient, SmoothGrad, and Grad-CAM look more random: We say they pass the test.

This shows that any method trying to multiply the input onto the score map is strange. Even in such cases where we should not attribute to any meaningful pixels, we see patterns in the map dependent on just the raw input image. Notice how this seemingly simple sanity check already conflicts with the theoretically justified completeness axioms.
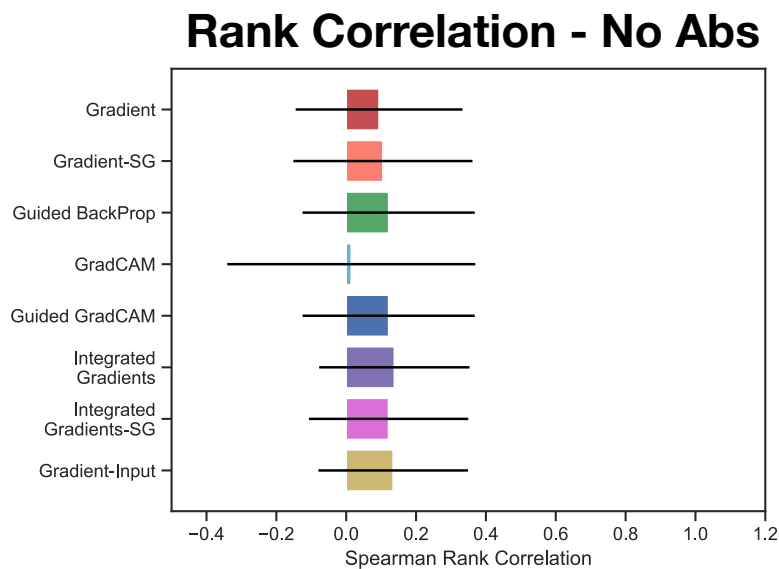
Figure 3.41: Spearman rank correlation barplot (without absolute values) of various explainability methods for an MLP. Grad-CAM gives convincing results. Details are discussed in the text. Figure taken from [2].

Quantitative results: rank correlation

We also briefly discuss a correlation plot in the paper, shown in Figure 3.41. How much correlation can we see between the upper and bottom rows for each method in Figure 3.40? Grad-CAM has a rank correlation of almost 0 for pixels on average. It satisfies the overall necessary condition the best – no correlation in attribution ranking for true/random labels. It does not seem to show any correlation between the explanation for the model trained with true labels vs. the model trained with random labels.

---

[24]This is understandable – it is multiplying gradient-based attribution with the pixel value differences between the image and baseline (a black image – MNIST). If we multiply the gradient-based attribution with the image of this 0 number, we will see a 0 in the attribution map.

Grad-CAM and SmoothGrad are still generally perceived as one of the best explanation methods.

We have seen that there is no GT explanation in general, and we are facing a chicken-egg problem. However, if we are a bit creative, we can simulate samples where the GT explanation actually exists (to some high extent). We know with very high confidence where the model should be looking for these images. According to the attribution map, we then check whether it is actually looking at the "GT part" of the image.

3.7.6 Simulation of Inputs with GT Explanations

We discuss a possible way to simulate inputs with GT explanations from the paper "Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV) " [98]. We define three classes: zebra, cab, and cucumber.

cab image   cab image with caption   cucumber image   cucumber with caption
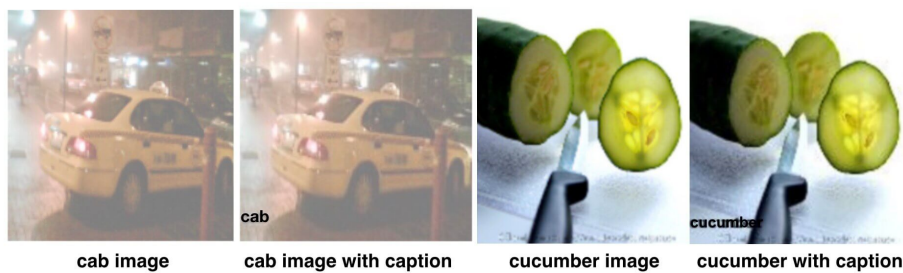
Figure 3.42: Samples from the dataset with "GT attributions" introduced in [98]. The image label is included in the image with a noise parameter that controls the probability that the label is correct. Figure taken from [98].

We provide potentially noisy captions written in the bottom left corner of the image. We have a controllable noise parameter $p \in [0, 1]$ to control the impact of the captions. In detail, p is the probability that the caption disagrees with the image content. $p = 0$ means there is no disagreement: an image of a cucumber would always have the caption 'cucumber'. For $p = 0.5$, each image has a 50% chance of the caption and the image content disagreeing. Examples are given in Figure 3.42. We have a feature selection problem (look at caption vs. image), but for low noise levels, the caption is a very prominent feature the model cannot resist looking at (refer to *simplicity bias* in 2.9). We will measure if the attribution methods are correctly picking that up. When the noise level is high, the model cannot rely on the captions at all. Thus, we will measure if the attribution methods are correctly **not** attributing the predictions to the label.

GT Attribution Results

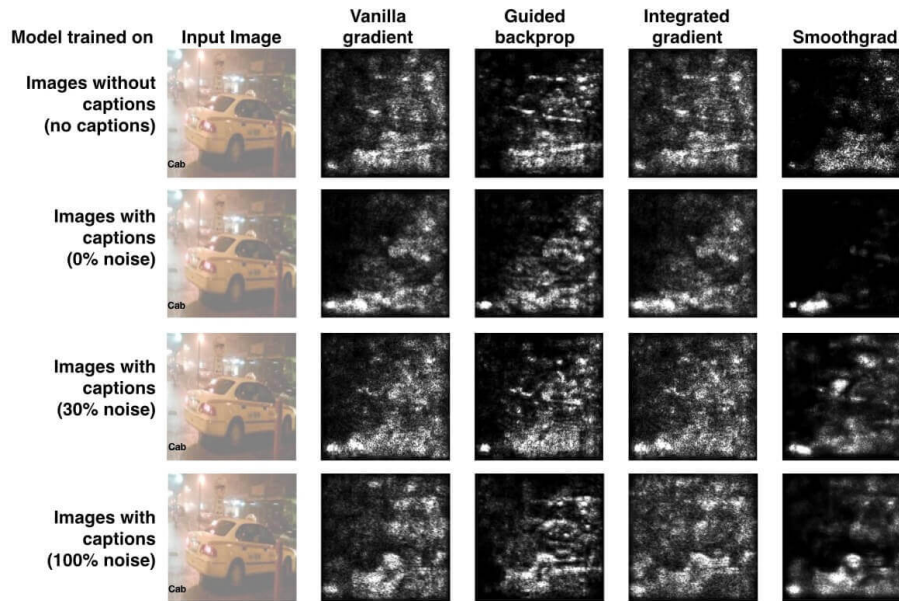Results are shown in Figure 3.43. Based on these results, SmoothGrad seems to be a great explainability method.

Figure 3.43: Results of various XAI methods on the dataset with "GT attributions" introduced in [98]. The 'cab' caption has to be a strong cue for recognition if $p = 0$. Conversely, it has to be a weak cue for recognition if $p = 1$. The test input image contains the *correct* caption. For the model trained on images with captions and 0% noise, we expect the attribution to be wholly focused on the caption. It seems like (without quantification) SmoothGrad is doing that most prominently, at least from how they show it. However, the gradient-based explanations are not well-calibrated (not ideal). Depending on how we renormalize the map, we may also get such a strong attribution to the caption for the other gradient maps. We cannot fully trust these kinds of score maps. Figure taken from [98].

3.7.7   Remove-and-Classify/Remove-and-Predict

Definition 3.14: Remove-and-Classify/Remove-and-Predict

The Remove-and-Classify algorithm is a prevalent soundness evaluation method for feature attribution scores. Attribution scores define a ranking over features: the feature attribution explanation $s(f, x) \in [0, 1]^{H \times W}$ ranks each feature in the input x. Remove-and-Classify removes features from the test input(s) iteratively, according to the attribution ranking of the explainability method. In the most popular variant, where the feature with the highest attribution score (most important) is removed first, the explainability method with the steepest drop in classification accuracy performs best. One usually calculates the Area under the Curve (AUC) to compare explanation methods. The features might be removed one by one or in a batched manner.

There are several variants of the Remove-and-Classify method. Compared to the variant introduced above, one might...

1. ...remove the features with the lowest attribution score (least important) first. In this case, the explainability method with the shallowest drop in accuracy performs best.

2. ...start from the base image and introduce features one by one (or in a batched way) according to the ranking of the explainability method – either w.r.t. increasing

> or decreasing attribution score.
>
> Sometimes, people take the average performance on these four possible benchmark combinations.
>
> There is no "correct" choice of encoding missingness. One must be particularly careful not to introduce *missingness bias* (3.7.8). The most prevalent removal technique for natural images and pixels as features is replacing the pixel with the mean pixel value(s) in the dataset, which is usually gray for natural images.
>
> **Note**: Just like in counterfactual explanation methods, grayness can still convey information – it can be problematic to consider this the base value. The quality of this choice also depends on our task – e.g., what if our task is to detect all gray boxes? Even though the signed distance of all data points to the mean (usually gray) image is zero on average, and the images are scattered around the mean image, it does not mean that individual gray pixels cannot contribute to a model's decision. They can be grouped into arbitrary shapes that have semantic meaning, even though a completely gray image might not convey much semantic information. Sticking to *any* color has a potential pitfall.

A feature attribution explanation gives a "heat map" of the given input. Suppose $s(f, x)$ is sound and correctly cites the causes for the prediction (in the correct order of importance). In that case, removing the most critical feature $i^* = \arg\max_i s_i(f, x)$ will significantly decrease the score $f(x)$ for the class in question. We measure the speed of decrease in classification accuracy as we remove pixels in the order dictated by $s(f, x)$.

We now discuss the result of remove-and-classify shown in Figure 3.32 that was reported in the CALM paper. The baseline is random erasing with equal probabilities. If we remove pixels, we kill information,[25] so we should still see a drop in accuracy, just not as fast. The used metric is the relative drop in accuracy when erasing according to an attribution method, compared to random erasing. A method is *better* if it results in a faster drop in accuracy. Methods corresponding to curves enveloping others from below are supposed to be more sound explanation methods. Sometimes, we also measure the AUC for this plot, where lower is better. One can also consider the unnormalized plot, where we do not compare against a random baseline.

Our observation is that CALM gives a sound explanation. It gives a huge drop in accuracy for pixels with high attribution scores. For SmoothGrad, the pixels attributed to being most important were not the most important ones, as we see a smaller drop.

One might ponder why most of the methods get worse than random erasing for larger values of k. Filling in gray/black pixels is not the best way to kill information. It can also *introduce* information. We address this in Section 3.7.8.

3.7.8 Missingness Bias

A recent phenomenon named *missingness bias* was reported in a recent paper titled "Missingness Bias in Model Debugging" [86]. Figure 3.44 aims to provide some insights. There

---

[25]As we will see in 3.7.8, while this is generally true, there are cases where we *introduce* information by encoding missingness.

is no common understanding of what the SotA for erasing information is. We argue that in-painting and blurring are good candidates. However, the choice of the inpainter and the exact blurring method are both hyperparameters that have important implications and might raise new problems. It is hard to explain what exactly is happening and what might be confusing textures for different architectures. It is, however, important to be aware of missingness bias and encode missing information in a suitable way.
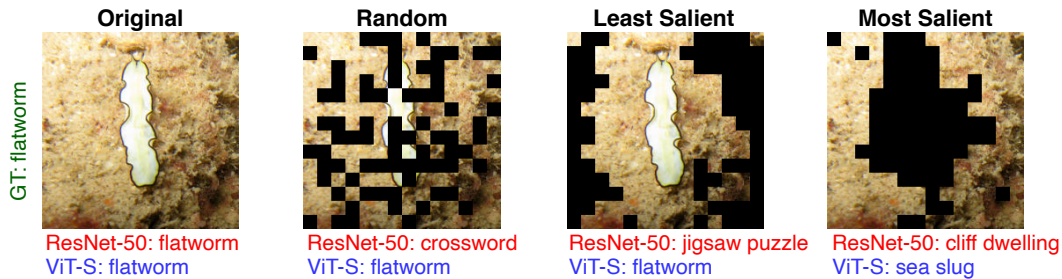

Figure 3.44: Illustration of the missingness bias. "Given an image of a flatworm, we remove various regions of the original image. Irrespective of what subregions of the image are removed (least salient, most salient, or random), a ResNet-50 outputs the wrong class (crossword, jigsaw puzzle, cliff dwelling). A closer look at the randomly masked image shows that the predicted class (crossword puzzle) is not totally unreasonable, given the masking pattern. The model seems to rely on the masking pattern to make the prediction rather than the image's remaining (unmasked) portions. Conversely, the ViT-S either maintains its original prediction or predicts a reasonable label given remaining image subregions." [86] Figure taken from [86]. Replacing pixels with mean values (or any fixed value) does not necessarily remove information. It may add further information (crossword) or kill unnecessary information. We also see that Transformers suffer a lot less from this phenomenon. Thus, depending on different models, the attribution methods might see different success rates. Remove-and-classify is not the perfect soundness evaluation metric. However, it is the most popular and one of the best ways to evaluate soundness.

We now address the curious behavior in Figure 3.32. When we only remove information by erasing pixels, we should see the random baseline as the worst-case removing strategy (in expectation). We see the jump above the baseline in Figure 3.32 because the model can predict based on the "removed" patterns for random removal, which can introduce greater changes in classification than removing pixels in an orderly fashion. Thus, random removal might add information that confuses the model more. If we erase according to CAM, we will see something like the right of Figure 3.44.

3.8   Soundness is Not The End of the Story

There are many other criteria, like soundness, simplicity, generality, contrastivity, socialness, interactivity, but also relevance. The latter depends on the final goal. Is it to debug? Is it to understand? Is it to gain trust? This is an essential criterion, as we are not looking at explanation methods for the sake of themselves, but we rather treat them as an intermediate step towards a final goal.

3.8.1   Various End Goals for Explainability

Model debugging as the end goal. Here, we wish to identify spurious correlations (why a model has made a mistake) and then fix them (we have seen methods for both in 2.10 and 2.11). This can improve generalization to OOD data. There are no successful/commercialized explanation tools yet that are specialized in debugging. It is not yet clear how to help an engineer fix a general problem with the model, and we have yet to see a *successful* use case of XAI for model debugging. There is still so much more to be researched for ML explanations. Attribution does not always guarantee successful debugging.

Understanding as the end goal. Do humans understand the idiosyncrasies ("odd habits") of a model? If a model is doing something odd, understanding why it is doing so could be an interesting objective. Can humans predict the behavior of a model based on the provided explanation? Do humans learn new knowledge based on the explanation? If the model is doing something new that humans cannot do yet, transferring that knowledge to humans would be essential.

Enhancing human confidence, gaining trust as the end goal. Does the explanation technique help persuade doctors to use ML models? Many doctors are still very averse to ML-based advice; they have no trust. Explanations could help them incorporate ML techniques. Does the explanation technique convince people to use self-driving cars (even though safety stays the same – or, as seen, worse because of the trade-offs)?

The "soundness" criterion does not fully align with the previous end goals and desiderata. The current evaluation is too focused on soundness (and qualitative evaluations). Given an explanation, we still have some end goals:

- ML Engineer: "Now I know how to fix model f."
- Scientist: "Now I understand the mechanism behind the recognition of cats."
- Doctor: "Now I can finally trust this model for diagnosing cancer."

Soundness focuses only on the explanation itself, which is an intermediate step. We need evaluation with the end goal in mind. There is no way we do not have to use human-in-the-loop (HITL) evaluation at some point.

3.8.2   Human-in-the-Loop (HITL) Evaluation

> **Definition 3.15: Human-in-the-Loop Evaluation**
>
> Human-in-the-loop evaluation refers to any evaluation technique for explainability that incorporates humans and measures how well the explanations help them achieve their end goals.

Let us now turn to discussing human-in-the-loop (HITL) evaluation. In particular, we will consider the paper "What I Cannot Predict, I Do Not Understand: A Human-Centered Evaluation Framework for Explainability Methods " [40].

An overview of the settings the authors consider is given in Figure 3.45. They address three real-world scenarios, each corresponding to different use cases for XAI.
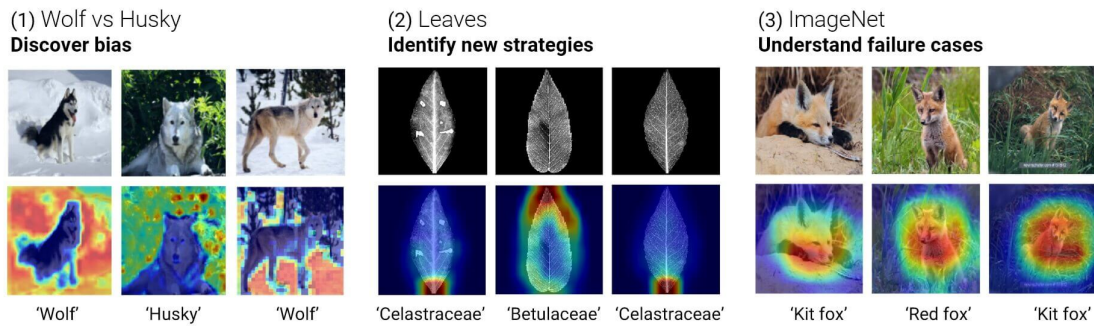
Figure 3.45: Overview of different tasks considered in [40]. The authors are evaluating recent explainability methods directly through the end goals of XAI – practical usefulness. They are trying to see whether explanations are actually helping humans in achieving their end goals. Figure taken from [40].

1. **Husky vs. Wolf**. Here, debugging is the end goal. Can the explanations help the user identify sources of bias in the model? Examples include background bias (snow, grass) instead of the animal.

2. **Real-World Leaf Classification problem**. Here, understanding is the end goal. Can the explanations help the user (non-expert) learn what parts of the leaf to look for to distinguish different leaf types? The humans want to adopt the strategy of the model.

3. **Failure Prediction Problem**. Here, understanding is the end goal again. This dataset is a subset of ImageNet. It consists of images, of which half have been misclassified by the model. Can the explanations help the user understand the failure sources of the (otherwise high-performing) model?
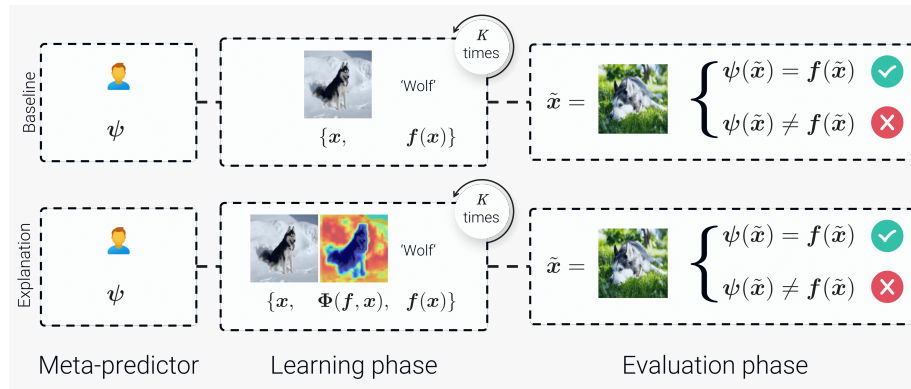
Figure 3.46: Overview of the stages of the method considered in [40]. The authors use a human-centered framework for the evaluation of explainability methods. The evaluation pipeline consists of (1) the predictor f, which is a black-box model, (2) an explanation method Φ, and (3) the meta-predictor, a human subject ψ whose task is to understand the behavior of f based on samples (i.e., the rules that the model uses for its predictions). First, the meta-predictor is trained using K triplets $(x, \Phi(f, x), f(x))$, where x is an input image, $f(x)$ is the model's prediction and $\Phi(f, x)$ is the explanation of the model's prediction. Second, for the Husky vs. Wolf and the Failure Prediction problems, the meta-predictor is evaluated on how well they can predict the model's outputs on new samples $\tilde{x}$. This is done by comparing the meta-prediction $\psi(\tilde{x})$ to the true prediction $f(\tilde{x})$. For the leaf classification problem, the meta-predictor is evaluated on how well they can classify the leaves after observing the explanations. The meta-prediction $\psi(\tilde{x})$ is compared to the GT label y. Figure taken from [40].

Figure 3.46 gives a detailed description of how the explainability methods are evaluated in all three scenarios. If the model makes a mistake on the evaluation image, the human should be able to pinpoint the mistake the model will make. Similarly, humans should be able to learn to classify leaves based on the knowledge encoded by the networks, and they should also be able to identify biases under the assumption that the explanation method works well. The paper uses the term *simulatability*. A model is explainable when its output can be predicted following the explanations.

First, we train humans on dataset $\mathcal{D} = \{(x_i, f(x_i), \Phi(f, x_i))\}_{i=1}^K$. For new samples, we let the humans predict the model predictions. The value $\psi^{(K)}(x)$ is the human prediction of the model prediction after training with K samples). The Utility-K score is calculated as follows:

$$\text{Utility-K} = \frac{P(\psi^{(K)}(x) = f(x))}{P(\psi^{(0)}(x) = f(x))}.$$

In words, the utility score is the relative accuracy improvement of the meta-predictor trained with or without explanations. The baseline factors out the contribution of explanations for educating humans. Humans for the baseline predictions are trained on dataset $\mathcal{D} = \{(x_i, f(x_i))\}_{i=1}^K$. To make the evaluation meaningful for Husky vs. Wolf and Failure Prediction, the authors mixed correct and incorrect model predictions 50-50% during evaluation.
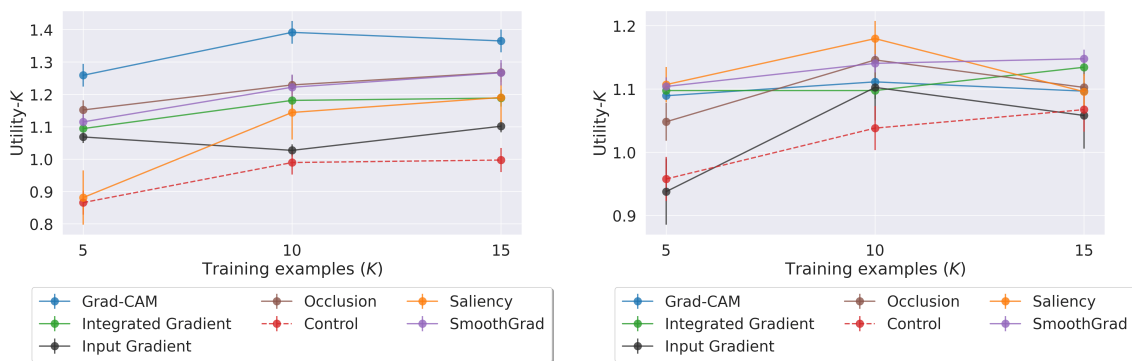
Figure 3.47: Results on the Wolf vs. Husky task (left) and the Leaf Classification task (right). For the Leaf Classification task, the Utility-K value is the normalized accuracy of the human predictor on test leaf images after observing the explanations. For the Husky vs. Wolf task, Grad-CAM, Occlusion, and SmoothGrad are seemingly useful. For the Leaf Classification task, Saliency, Smoothgrad, and Integrated Gradients seem to perform best. Figure taken from [40].

Task (i): Husky vs. Wolf

For Husky vs. Wolf, results are shown in the left panel of Figure 3.47. The control group is shown some score map, called bottom-up saliency, that is not an explanation (it is independent of model f). This is used to rule out the possibility that people try harder to solve the task if any explanation is provided to them. CAM achieves good results (same story as before); SmoothGrad and Occlusion are also good. All attribution methods used are better than the control 'method'. More training samples mean further knowledge of what the model might be doing.

Task (ii): Leaf Classification

The results for Leaf Classification are shown in the right panel of Figure 3.47. This is an example of using ML to educate humans. In this case, Utility-K is the normalized accuracy of the meta-predictor after "training". Humans do not know how to distinguish these leaf types in the beginning. By showing where the model is looking (that can solve the task well), humans also learn how to classify the leaves (as they learn useful cues). SmoothGrad is consistently good for educating humans for the task. CAM helps a bit less ideally compared to Smooth-Grad. Saliency also performs well but does not scale well to more training samples ($K = 15$). Integrated Gradients is on par with CAM and also scales better.

Task (iii): Failure Prediction

On the ImageNet dataset, *none of the methods tested exceeded baseline accuracy*. These results made the authors suspicious that the explanation methods might not be sound: If the user observes explanations from a method that is not sound, it will not gain enough insight into the model's internals. The authors compared Utility scores (AUC scores under the (K, Utility-K) curve for various K values) to corresponding faithfulness scores. The results are shown in Figure 3.48.
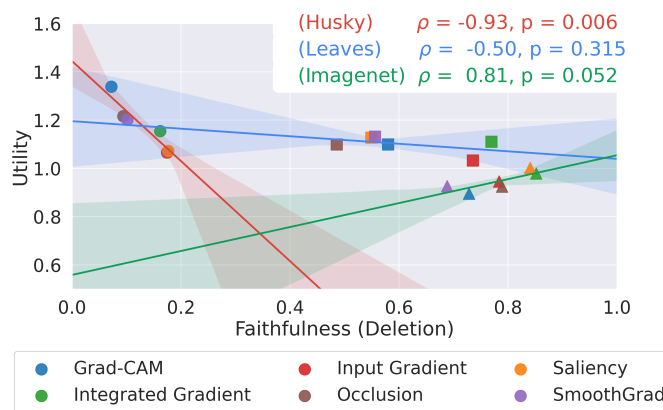
Figure 3.48: Correlation of the faithfulness and the utility score in [40]. HITL Utility does not correlate in general (across datasets) with the evaluation's soundness (faithfulness). (Correlations change across tasks. It seems very random.) Faithfulness metrics are poor predictors of end-goal utility. For the Husky vs. Wolf and leaves datasets, a negative correlation can be observed, meaning *high soundness might even come with the price of less end-goal utility*. Figure taken from [40].

Conclusion for HITL Evaluation

SmoothGrad is doing a great job in helping humans with the end goals considered in the benchmark. If we care about how humans can understand attributions and learn from attributions (explanations), then soundness evaluation is not a good proxy for choosing between methods. Thus, HITL evaluation cannot be replaced with soundness evaluation.

> **Additional Information: HITL**
>
> Out of the three downstream use cases mentioned in the book, HITL evaluation seems to be more tailored toward understanding. How could it measure how much trust is given to the model? How could we measure how much a method helps fix a model (if the answer is not spurious correlations)? While HITL is a step forward compared to soundness evaluation in the sense that it measures how humans understand better, it still does not measure the end-to-end metric of how much more trust is given or how much better a model is debugged after an explanation in general. End-goal-tailored explainability is a young field with many questions to be answered.

3.9 Towards Interactive Explanations

Previously, we have seen that for HITL evaluations, we need to include human participants to evaluate how useful a method is for human subjects and their end goals. *Interactive explanations* are also deemed necessary by decision-makers.

3.9.1 A Survey on Explanations

> **Additional Information: Quant**
>
> A quant, short for quantitative analyst, is a person who analyzes a situation or event (e.g., what assets to buy/sell in a hedge fund), specifically a financial market, through complex mathematical and statistical modeling.

We consider a survey for decision-makers using ML, titled "Rethinking Explainability as a Dialogue: A Practitioner's Perspective" [114]. The survey aims to find answers to the question, "What kind of features do you need from explanations?"

Desiderata for Interactive XAI

Let us now discuss the survey for domain experts using ML in detail. In particular, we consider exact statistics from the survey. **Note**: There is only a small number of respondents, but as they are experts, conducting such surveys is expensive. The quotes are imaginary and only illustrate the discussed desiderata. The list also does not mean that there are technologies already satisfying these desiderata. We are far away from many aspects still.

24/26 respondents wish to eliminate the need to learn and write the commands for generating explanations. "We do not want to care about writing code. We need a more natural-language-based interaction with the system."

24/26 respondents prefer methods that describe the accuracy of the explanation in the dialogues. A notion of *uncertainty* is needed.

23/26 respondents wish to use explanation tools that preserve the context and enable follow-up questions. "If we do not understand something in the previous round, we should be able to ask for follow-up explanations." A key characteristic of a dialog-based system is that the machine should remember previous topics/conversations.

21/26 respondents would like real-time explanations. "Do not take several hours to answer our questions. We want an experience as if we were talking to a human." This is a rather basic requirement for efficiency.

17/26 respondents would let the algorithm decide which explanations to run. Users should not have to ask for a specific explainability algorithm. "We do not wish to decide ourselves, as so many of them exist. We do not want to build a benchmark, compare all attribution methods, and decide on an appropriate one for the use case. The system should determine the best algorithm for our domain."

Key Takeaways from the Desiderata

Decision-makers prefer *interactive* explanations. Explanations are preferred in the form of *natural languages*. Experts want to treat machine learning models as "another colleague" they can talk to. For example, a hedge fund might find good use of ML: They might wish to have a virtual human (a quant) sitting next to them who can answer questions like "Why do you think this trend is happening?" or "Why did you buy/sell the stocks?" They want to ask the models' *opinion* or what they had in mind when making a decision. In particular, they want models that can be held accountable by asking why they made a particular decision through expressive and accessible natural language interactions.

3.9.2 Generating Counterfactual Explanations with Natural Language

The "Generating Counterfactual Explanations with Natural Language" [75] paper is a work of Hendricks et al. An overview of the method is given in Figure 3.49. This is a step towards interactive explanations for humans.

3.9.3 e-ViL

e-ViL is an explainability benchmark introduced in the paper "e-ViL: A Dataset and Benchmark for Natural Language Explanations in Vision-Language Tasks" [94]. A test example and the outputs of various VL models are given in Figure 3.50. An overview of the architectures benchmarked in this work is shown in Figure 3.51.

3.9.4 Summary of Interactive Explanations

We only touched on interactive explanation techniques. These are very new, and there are only a few works. However, it has much potential. We recommend working in this domain. To work forward, we need to take humans into account – whether XAI is helpful for the end user (HITL evaluation) and view XAI systems as socio-technical systems. We also want users to be able to interact with the explanation algorithm and to make the interface more natural for humans (e.g., having a natural-language-based "chat" about the explanation).
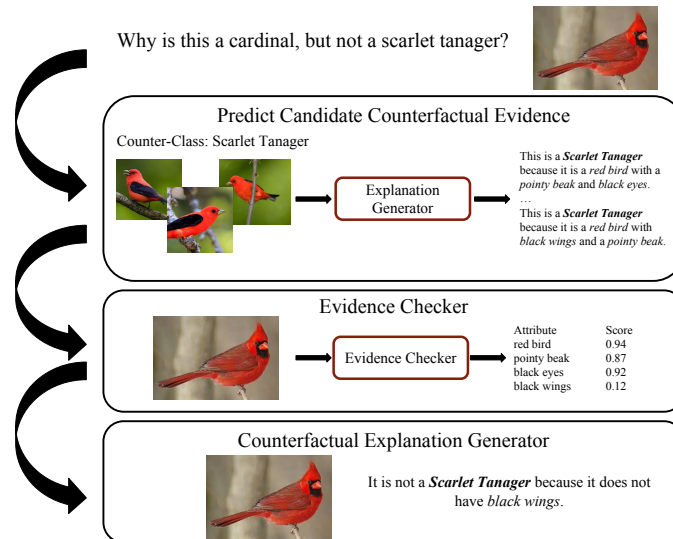
Figure 3.49: Overview of the counterfactual explanation pipeline of [75]. The method allows users to generate explanations based on high-level concepts. We have a bird classifier available. There is also an explanation generator. This is different from image captioning, as image captioning only talks about what is in the image, but the explanation generator first makes a prediction (Scarlet Tanager) for the set of counter-class images and describes in natural language why it thinks it is that class. The evidence checker checks how many characteristics extracted from the explanation generator's explanations are present in the current input (i.e., it makes a list of scores). It is checking for evidence of these characteristics in the current image. The counterfactual explanation generator can use the evidence to answer the counterfactual question. Figure taken from [75].

3.10 Attribution to Model Parameters

As we have seen, a model is a function that is an output of a training algorithm (which, in turn, is another function of the training data and other ingredients). The model takes training data as input implicitly through the training procedure. This is a hidden part of the model that is not used for explanations when we only focus on the attribution to the test sample features. It can very well happen that the model is making a bizarre decision not because of a specific feature in the test sample but because of strange (defective) training samples. It is difficult to rule this possibility out, and it is, therefore, meaningful to look at training samples.

We write the model prediction as a function of two variables:

$$Y = \mathrm{Model}(X; \theta) = \mathrm{Model}(X; \theta(\{z_1, \ldots, z_n\}))$$

where Y is our prediction, X is the test input, θ are the model parameters, and $\{z_1, \ldots, z_n\}$ is the training dataset. We use z because these can both correspond to inputs and input-output pairs. The prediction of our model is implicitly also a function of the training data.

As we discussed before, explaining our prediction against features of x is not always sufficient (but is very popular). We might also be interested in the contribution of

1. individual parameters θ_j of the model, and
2. individual training samples z_i in the training set

Hypothesis: The people are flying kites at the beach.
Answer: Contradiction
RVT: People can't be riding kites while they are flying kites.
PJ-X: People cannot be flying and flying at the same time.
FME: People cannot be walking and flying kites at the same time
e-UG: People cannot be flying kites while they are standing on a street.
GT Explanation: construction site is different from the beach

Figure 3.50: A test example from the e-SNLI-VE dataset [94]. *Contradiction* means the hypothesis contradicts the image content. *RVT, PJ-X, FME, and e-UG* are explanation methods. They provide natural language explanations (NLEs). Explanations are not trained on any GT. (If they were, that would be another predictive model, and there is no guarantee it would explain the model's way of prediction.) They instead extract information from a vision-language (VL) model into a human language format. The *GT Explanation* is a human-generated explanation for the answer collected by the authors. **Task**: Given an image and a hypothesis, decide if the hypothesis is aligned with the image. The machine also has to explain why they might be contradictory. VL-NLE models predict *and* explain. Figure taken from [94].
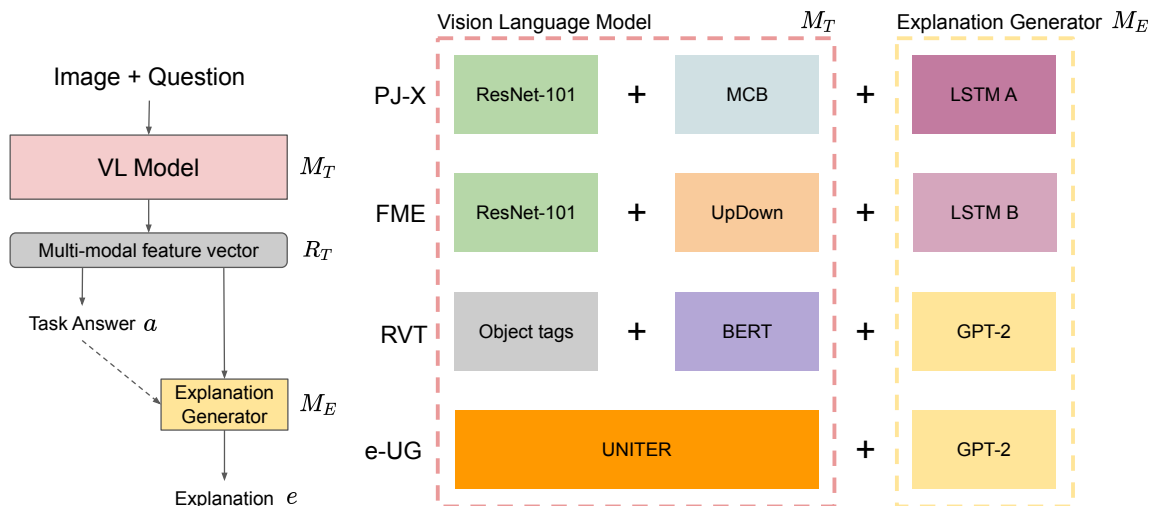
Figure 3.51: Overview of the structure of general VL models (left) and detailed subparts of individual models benchmarked in [94]. Figure taken from [94].

to the final prediction of the model. First, we look at the contribution of model parameters. Then, we discuss attribution methods to training samples.

3.10.1 Explanation of Model Parameters θ

For DNNs, model parameters are simply millions of raw numbers. They are complicated to understand. Explaining a prediction w.r.t. these raw numbers is seemingly a tricky problem. This is in contrast with the input-level features x and labels y. Inputs to a DNN are usually sensory data (image, sound, text), so humans can naturally understand them.

Thus, inputs and outputs to a DNN are often human-interpretable. However, the parameters are not, at least not directly. To understand the parameters θ, we "project" them onto the input space; i.e., we give visualizations of them (or explain them in text for NLP methods).
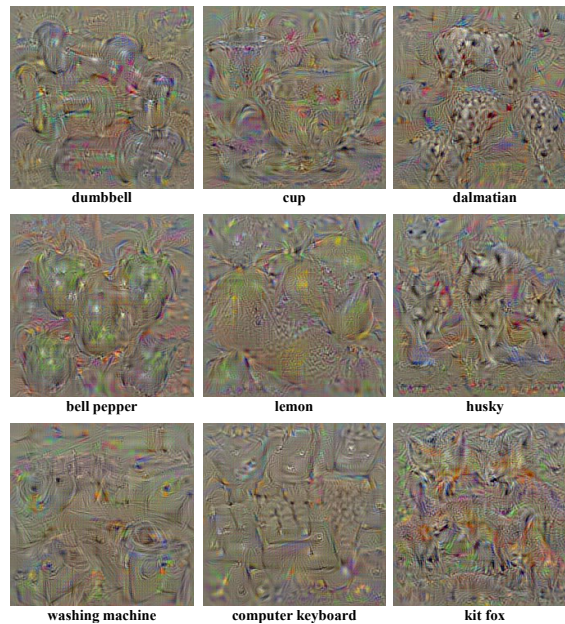
Figure 3.52: Various weight visualizations w.r.t. different target classes from [182]. We ask "What is the most likely image for the class dumbbell?" from the model, or "What excites a certain neuron most?" One can employ several regularization techniques (e.g., TV) to make the visualizations more interpretable. For these samples, the model predicts a very high score for the respective classes. These are preliminary results from a seminal paper about turning model parameters into an image in the input space. Figure taken from [182].

Examples for turning parameters into samples from the seminal paper "Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps" [182] are given in Figure 3.52. We generate these samples by solving an optimization problem in the pixel space. We maximize the score for class c in the input space in a regularized fashion:

$$\arg\max_{I} S_c(I) - \lambda\|I\|_2^2,$$

where $S_c(I)$ is the prediction score (logit value, pre-activation of the output layer) for class c and image I from the network. L_2 regularization prevents a small number of extreme pixel values from dominating the entire image. It results in smoother and more natural (more interpretable) images. We can also regularize the discrete image gradient (e.g., with the TV regularizer), which is also a popular choice. This mitigates the noise issue even more.[26] The

---

[26]The L_2-regularized image can still be very noisy, just a bit less than the original because of the reduction in magnitude.

objective of adversarial attack algorithms is very similar to this optimization problem. However, attacks try to minimize the score for a specific class. Here we are trying to maximize it, e.g., using gradient descent for the loss (less often used) or using gradient ascent for the logit value (popular).

3.10.2   More examples of turning parameters into samples

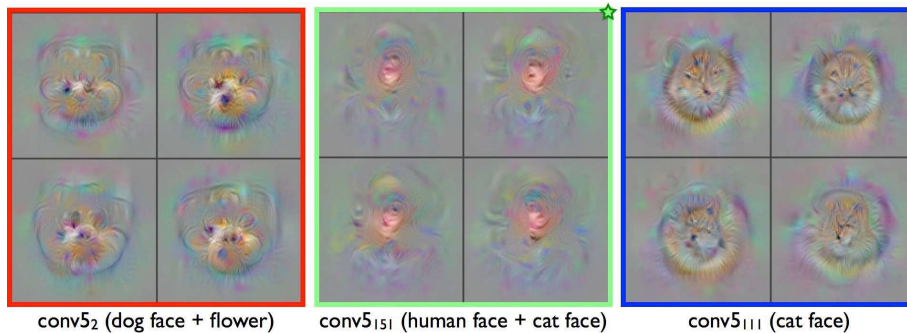Let us discuss two more examples of turning parameters into samples.

conv5$_2$ (dog face + flower)   conv5$_{151}$ (human face + cat face)   conv5$_{111}$ (cat face)

Figure 3.53: Visualizations of a deeper layer of an AlexNet-like architecture from [218]. The synthesized images resemble a mixture of animals, flowers, and more abstract objects. The indices correspond to different feature channels of the conv5 pre-activation tensor. Each grid corresponds to four runs of the same optimization problem. Figure taken from [218].

First, we consider "Understanding Neural Networks Through Deep Visualization" [218]. We can perform the previous optimization procedure on different intermediate layers as well. Instead of maximizing the score of a certain class, we maximize an intermediate feature activation for one of the units of a layer or maximize the entire layer's activation. We then recognize patterns in the generated images. These are interpreted as the patterns that the corresponding neurons have learned and respond to. The optimization problem here is

$$x^* = \arg\max_{x} \left(a_i(x) - R_\theta(x) \right),$$

where $a_i(x)$ can be an activation for a particular unit in a particular layer, or we can also maximize the mean, min, and max activation in a layer. That leads to similar results. (Not done in this work.) $R_\theta(x)$ is the regularization term. In this work, the authors use

$$x \leftarrow r_\theta \left(x + \eta \frac{\partial a_i(x)}{\partial x} \right),$$

which is more expressive. For example, for L_2 decay one can choose $r_\theta(x) := (1-\theta) \cdot x$. An example collage is shown in Figure 3.53. Please refer to the full paper for various visualizations across many layers, which we discuss below.

When we maximize the output of an early neuron, its receptive field is usually smaller than the entire input image. Thus, when we visualize the optimized input, we will see only the small corresponding region changing in the input. The other input regions are left as we initialized them.

Typically, we will not see any interpretable pattern for many of the neurons. In many cases,

people cherry-pick to generate these images. The results are also heavily dependent on the initial image of the optimization. One should be careful with how they interpret them.

At higher layers, we visualize more semantic content (e.g., cup, garbage bin, goose). To visualize entire layers, one can take one image for each channel in the corresponding layer's feature map output. As we go down the layers, we see more and more generic patterns. These are smaller, more common patterns that are found in many objects.

When we visualize the optimized inputs for the *first* convolutional layer's neurons, we roughly see the filters (see Gabor filter) of the corresponding channel for each neuron in that channel. These contain single colors or combinations of a few repetitive textures. If we use a single convolution, we just have a sparse linear network (S_c becomes linear w.r.t. I). The regularized activation-maximizing inputs are nearly the same as the filters themselves. If we take a channel of the filter (of shape $(3, H, W)$) corresponding to the output channel of choice, we can directly visualize it. When doing so, we will see very similar visualizations to the visualizations of the regularized activation-maximizing inputs.

Consider a 3×3 convolutional layer with a single channel. Then the filter K is of shape $(1, 3, 3, 3)$. The operation for a single neuron a in the output is simply

$$a = \sum \left(I^a \odot K \right),$$

where I^a is the receptive field of the neuron a, of shape $(3, 3, 3)$. If we perform unregularized optimization, we obtain

$$I^{a*} = \mathbf{1}(K > 0).$$

By using, e.g., L_2 regularization, we roughly get $I^{a*} \approx K$, with the outline of the generated image being the same as the corresponding filter channel.

As our second example, we look at "Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images" [148]. One can use the previously introduced optimization problem in the pixel space to generate images corresponding to high activations (e.g., maximize prediction score for the class of choice). We get significantly different images depending on the regularization of the image generation (effectively, the search space). This work introduces a different generation technique that has astounding results. A teaser is shown in Figure 3.54. Please refer to the paper for an extensive collection of visualizations.

The provided visualizations allow us to, e.g., look into the texture bias of the network. For example, the activation-maximizing input (in the modified search space) for "baseball" contains a very similar pattern as a baseball. However, we would not say that this is a baseball as humans. Nevertheless, the model predicts "baseball" with very high confidence. The technique allows us to see into the model's decision-making process, which might often be quite surprising.

In the paper's oral talk, the authors also showed that classifying images through their mobile phones gives the same result (they made an app for live demonstration). This shows that the visualizations are stable representations of the classes for the DNN in question.

The authors provide multiple visualization techniques. All visualizations correspond to highly confidently predicted generated images for the classes $0 - 9$. We get astonishing results on
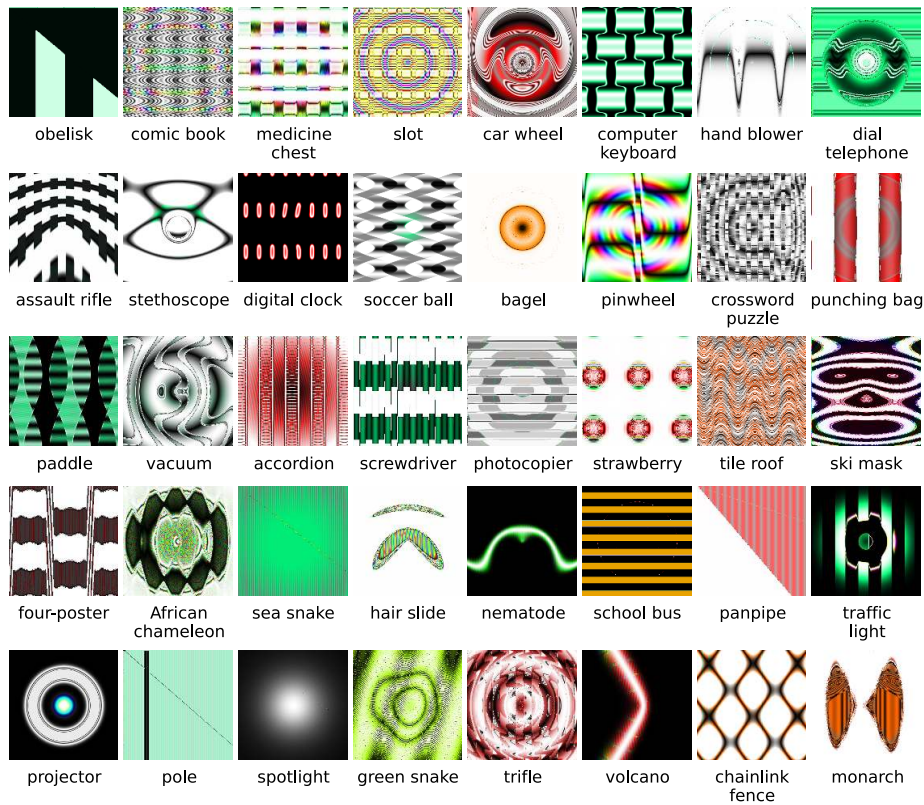
Figure 3.54: Visualizations of high-confidence images for different labels using the novel generation technique of [148]. Figure taken from [148].

even a simple dataset like MNIST.

The linked resource is a recommended blog post for feature visualization.

3.10.3   Criticism of feature visualization

While feature visualization can give impressive results (as can be seen in the Distill blog post), there is criticism about the utility of such methods. Even though we can find visualizations for units of a neural network that correspond to a human concept (cf. baseball example above), many visualizations are not interpretable [142]. We are not guaranteed to find something. Moreover, when we look at feature visualization as explanations to humans about what *causes* a CNN to activate, Zimmermann et al. found that feature visualizations do not provide better insight into model behavior than e.g. looking at data samples directly [225]. The visualizations are interesting but simpler approaches can provide the human with the same intuitions. A recent work [59] shows that feature visualizations are not reliable and can easily be fooled by an adversary while keeping the predictive performance of the model. They also prove that feature visualizations cannot guarantee to deliver an understanding of the model. Feature attribution gives an interesting tool for exploratory analysis, but it is not necessarily suitable as is for explaining model behavior to humans.

3.11 Attribution to Training Samples

This part is much more critical than attributing to individual weights, for reasons clarified just below.

3.11.1 Why attribute to training samples?

We saw that we had to eventually map our parameters onto the input space to visualize what was happening inside the model. It is a natural question to ask "Why don't we just look at the raw ingredients for the parameters then?" These are exactly the training samples. Model parameters are *not interpretable*, so it makes sense to attribute *directly* to the training samples. Model parameters θ are also built on the training samples $\{z_1, \ldots, z_n\}$. Therefore, attributing to training samples is sufficient. Training data explanations are more likely to give more *actionable* directions to improve our model. If we can trace the model's error or strange behavior back to the training samples, we can fix/add/remove training samples to resolve erratic behavior. If we find out that the model made a mistake because some of the labels were wrong, we can (1) relabel these samples that were attributed to the strange behavior, (2) remove them if the GT labeling $P(Y \mid X = x)$ is too stochastic (faulty sample), or (3) we can even add new samples to the training set if we think there is no strong sample supporting the right behavior of the model we wish to see.

3.11.2 Basic Counterfactual Question for Attribution to Training Data – Influence Functions

We look at influence functions, first used for deep learning in the paper "Understanding Black-box Predictions via Influence Functions" [104]. These find influential training samples for the model prediction on a particular test sample. The influence here is an answer to the counterfactual question "What happens to the current model prediction for a test input x if one training sample z_j was left out of the training set $\{z_1, \ldots, z_n\}$?" This is a minor change in the training set, as typically, the training set size is in the range of millions to billions. Leaving out one sample will generally not greatly affect the overall behavior of the model. However, for a particular test sample z, we can still be interested in the training samples that made the largest impact on the test sample through the optimization procedure. Such training samples are likely to be visually similar to the test sample. We want an algorithm to measure the impact of each training sample on this particular test sample.

Notation

The notation is introduced in Table 3.3. To find out $L(z, \hat{\theta}_{\setminus j}) - L(z, \hat{\theta})$, we could retrain the model on the dataset without z_j. However, this is infeasible for real-life scenarios. To study the impact of every training sample on every test sample, we would need to train ("number of training samples" + 1) DNNs and evaluate the differences in the losses for all test samples we want to consider.

Note: We assume that $\hat{\theta}$ and $\hat{\theta}_{\setminus j}$ are *global* minimizers of the respective empirical risks. This is a strong assumption, but we will see that relaxations of the resulting method still work well in practice.

Table 3.3: Notation used for discussing attribution to training samples. A large positive $L(z, \hat{\theta}_{\setminus j}) - L(z, \hat{\theta})$ implies that z_j is very useful for test sample z. Conversely, when it is a large negative number, z_j is highly harmful to z.

Notation	Description
$z = (x, y)$	Test input and output of interest.
$z_i = (x_i, y_i)$	Training samples.
$L(z, \theta)$	Test loss for test sample z with parameters θ.
$\frac{1}{n} \sum_i L(z_i, \theta)$	Training loss (empirical risk) for training sample z_i.
$\hat{\theta} := \arg\min_\theta \frac{1}{n} \sum_i L(z_i, \theta)$	Solution parameters for the original training set.
$\hat{\theta}_{\setminus j} := \arg\min_\theta \frac{1}{n} \sum_{i \neq j} L(z_i, \theta)$	Solution for the training set after removing z_j.
$L(z, \hat{\theta}_{\setminus j}) - L(z, \hat{\theta})$	Change in the test loss for z after removing training sample z_j.

Lesson from attribution methods: take the gradient!

We only had to do a single backpropagation to determine the contribution of all pixels modified by an infinitesimal amount (separately) to the infinitesimal change in the output. Here, we also take the gradient of the test loss $L(z, \hat{\theta})$ w.r.t. the training sample z_j, where the two values are connected through the entire optimization procedure. In particular,

$$\hat{\theta} = \arg\min_\theta \frac{1}{n} \sum_i L(z_i, \theta).$$

$\hat{\theta}$ is, therefore, a function of z_j through the optimization we employ, and the dependency between the test loss and z_j is exactly through $\hat{\theta}$. We are interested in the change in the test loss when we make a small change in the training sample z_j.

We need a few tricks to compute the gradient

$$\frac{\partial L(z, \hat{\theta})}{\partial z_j}.$$

Taking the gradient through an optimization procedure has been done in subparts of ML quite a few times. There are algorithms like "gradient descent by gradient descent" [11].[27] This is overall a great technique to know.

First, we generalize the notion of "removal" into a continuous procedure. Removing z_j is a discrete procedure and is thus non-differentiable. Instead, we take the loss of z_j into a separate term:

$$\hat{\theta}_{\epsilon, j} := \arg\min_\theta \frac{1}{n} \sum_i L(z_i, \theta) + \epsilon L(z_j, \theta).$$

Note that the first term still contains a z_j term.

- $\epsilon = 0 \in \mathbb{R}$: We recover the original minimizer of the training loss, $\hat{\theta}_{0,j} = \hat{\theta} \in \mathbb{R}^d$.
- $\epsilon = -1/n$: We obtain our previous notion of "removal".

---

[27]This algorithm aims to find the optimal LR without cross-validation through another GD algorithm. Optimal here means good for generalization to the held-out validation set. For this, we also need backpropagation through the optimization procedure.

We further assume that the loss L is twice differentiable and strictly convex w.r.t. θ (so that the Hessian matrix of the model parameters is PD). For DNNs, there is usually no unique $\hat{\theta}_{\epsilon,j}$ and $\hat{\theta}$ because of weight space symmetries and other contributing factors that make the loss landscape highly non-convex, with many equally good minima. So we further enforce strict convexity to have a unique optimum. This is a rather typical trick in research: We assume that everything is simple during theoretical derivations. In practice, we ignore the assumptions and hope our method still works.

One can obtain the following derivative (with annotated shapes for clarity):

$$\underbrace{\frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0}}_{\in \mathbb{R}^d} \approx - \underbrace{H_{\hat{\theta}}^{-1}}_{\in \mathbb{R}^{d \times d}} \underbrace{\nabla_\theta L(z_j, \hat{\theta})}_{\in \mathbb{R}^d}.$$

where

$$H_{\hat{\theta}} = \frac{1}{n} \sum_i \nabla_\theta^2 L(z_i, \hat{\theta}).$$

In words, $\frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\big|_{\epsilon=0}$ is the derivative of the weights w.r.t. ϵ, evaluated at $\epsilon = 0$ when $\hat{\theta}_{\epsilon,j} = \hat{\theta}$. This gives the relative change in the globally optimal weights using the original objective if we change the additional influence of z_j by an infinitesimal amount from 0. The last term is the gradient of z_j loss w.r.t. θ, evaluated for weights $\hat{\theta}$. $\nabla_\theta^2 L(z_i, \hat{\theta})$ is the Hessian matrix of L:

$$\left(\nabla_\theta^2 L(z_i, \hat{\theta}) \right)_{ij} = \frac{\partial^2 L(z_i, \hat{\theta})}{\partial \theta_i \partial \theta_j}.$$

Why is the derivative formula well-defined, i.e., why is this average Hessian matrix invertible? It is a well-known fact that the average of symmetric, positive definite (PD) matrices is symmetric PD. The Hessians are symmetric because of Schwarz's theorem (the loss has continuous second partial derivatives w.r.t. θ everywhere). The Hessians are also PD (i.e., they only have positive eigenvalues, and there is strictly positive curvature in all directions) because the function is strictly convex by assumption. Therefore, $H_{\hat{\theta}}$ is symmetric PD.

> **Additional Information: Interpreting the Hessian**
>
> If we have 10^6 parameters, then $\nabla_\theta L(z_i, \hat{\theta}) \in \mathbb{R}^{10^6}$ gives us how the function value changes in each principal axis direction *relative* to an infinitesimal change. To obtain the relative change in the loss value in a particular input direction v, one can consider
>
> $$\nabla_\theta L(z_i, \hat{\theta})^\top v \in \mathbb{R}, \quad \|v\| = 1.$$
>
> Similarly, $\nabla_\theta^2 L(z_i, \hat{\theta}) \in \mathbb{R}^{10^6 \times 10^6}$ gives us how the gradient of the loss at z_i changes in the neighborhood of θ along all canonical axes. This is why it is a matrix. In each axis direction, we measure the relative change in the gradient vector (in each of its entries) w.r.t. each principal axis. To get the rate of change of the gradient (curvature) in a particular

input direction v, one can consider

$$v^\top \nabla_\theta^2 L(z_i, \hat{\theta}) v \in \mathbb{R}, \quad \|v\| = 1.$$

When Hessians are symmetric (which is almost always the case in ML settings), they are orthogonally diagonalizable. In this case, the diagonal entries of the diagonalized Hessian give the rate of change of the gradient (curvature) in the eigenvector directions. Let

$$\nabla_\theta^2 L(z_i, \hat{\theta}) = Q\Lambda Q^\top$$

where Λ is diagonal and Q is orthogonal. Then, if v_i is the ith eigenvector direction, we have

$$v_i^\top \nabla_\theta^2 L(z_i, \hat{\theta}) v_i = v_i^\top Q\Lambda Q^\top v_i = v_i^\top Q\Lambda e_i = \lambda_i v_i^\top Q e_i = \lambda_i v_i^\top v_i = \lambda_i.$$

Our story does not end here, as we wish to see the influence of z_j on the test loss for test sample z. Given the previous result, we compute the influence of sample z_j on the loss for test sample z as $\mathrm{IF}(z_j, z) \in \mathbb{R}$,

$$
\begin{aligned}
\mathrm{IF}(z_j, z) &:= \frac{\partial L(z, \hat{\theta}_{\epsilon,j})}{\partial \epsilon}\bigg|_{\epsilon=0} \\
&= \nabla_\theta L(z, \hat{\theta})^\top \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} \\
&= -\nabla_\theta L(z, \hat{\theta})^\top H_{\hat{\theta}}^{-1} \nabla_\theta L(z_j, \hat{\theta}).
\end{aligned}
$$

This is the formulation for IF that the referenced paper uses. The IF value gives the relative change in the test loss value if we increase ϵ by an infinitesimal amount from $\epsilon = 0$. This "upweighing" represents the removal of z_i from the loss computation. It is large and positive when upweighting z_j a bit increases the loss by a lot (harmful) $\iff$ when downweighting z_j a bit decreases the loss by a lot. It is large and negative when upweighting z_j a bit decreases the loss significantly (helpful). This formulation refers to *negative influence*. *As we would intuitively expect a high influence value for a sample that decreases the loss a lot, both this book and the Arnoldi paper [172] consider the definition*

$$\mathrm{IF}(z_j, z) = \nabla_\theta L(z, \hat{\theta})^\top H_{\hat{\theta}}^{-1} \nabla_\theta L(z_j, \hat{\theta}).$$

Let us consider some remarks. Using a first-order Taylor approximation, it is also clear that

$$L(z, \hat{\theta}_{\epsilon,j}) = L(z, \hat{\theta}) + \epsilon \frac{\partial L(z, \hat{\theta}_{\epsilon,j})}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon).$$

One can study the behavior of the test loss when perturbing sample z_j by an infinitesimal amount in a clear way using the above formula. The IF formula is also very symmetrical: we are taking a modified dot product between the gradient of loss of the test sample and the training sample.

From now on, we will use the latter definition for the influence function (without the negative

sign). We have discussed that $H_{\hat{\theta}}$ is symmetric and PD. Therefore, it can be orthogonally diagonalized, i.e., we can find a rotation/mirroring such that in this new basis, the average Hessian on the training points is a diagonal matrix:

$$H_{\hat{\theta}} = Q\Lambda Q^{\top}$$

with an orthogonal matrix Q (with ortho*normal* columns), and its inverse is given by

$$H_{\hat{\theta}}^{-1} = Q\Lambda^{-1}Q^{\top}.$$

Therefore,

$$
\begin{aligned}
\text{IF}(z_j, z) &= \nabla_\theta L(z, \hat{\theta})^{\top} H_{\hat{\theta}}^{-1} \nabla_\theta L(z_j, \hat{\theta}) \\
&= \nabla_\theta L(z, \hat{\theta})^{\top} Q\Lambda^{-1}Q^{\top} \nabla_\theta L(z_j, \hat{\theta}) \\
&= \left(Q^{\top}\nabla_\theta L(z, \hat{\theta})\right)^{\top} \Lambda^{-1} \left(Q^{\top}\nabla_\theta L(z_j, \hat{\theta})\right) \\
&= \left\langle Q^{\top}\nabla_\theta L(z, \hat{\theta}), Q^{\top}\nabla_\theta L(z_j, \hat{\theta})\right\rangle_{\Lambda^{-1}}.
\end{aligned}
$$

To calculate $IF(z_j, z)$, we rotate/mirror the gradient vectors to transform them into the eigenbasis of the average Hessian, then compute a generalized dot product between the gradients expressed in the eigenbasis, weighted by the corresponding diagonal entries of Λ^{-1} (the inverse curvatures in each direction of the eigenbasis). The dot product is, therefore, calculated in a *distorted space*, where directions with the flattest curvature in the loss landscape are given more weights. To get a high influence value, having large positive values in these directions in the gradient vectors expressed in the eigenbasis is more important.

The caveat is that *we might have millions or billions of parameters*. Let $p :=$ number of parameters $= \mathcal{O}(\text{millions-billions})$ and $n :=$ number of training samples $= \mathcal{O}(\text{millions-billions})$. Then, the naive $H_{\hat{\theta}}^{-1}$ computation is $\mathcal{O}(np^2 + p^3)$ where the np^2 part corresponds to computing $H_{\hat{\theta}}$ and p^3 corresponds to computing its inverse. Computing $H_{\hat{\theta}}^{-1}$ dominates the IF computation when n is not significantly larger than p. In practice, naive computation is prohibitive and infeasible.

> **Additional Information: Proof of the Derivative Formula**
>
> We start with the definitions
>
> $$\hat{\theta} = \arg\min_\theta \frac{1}{n}\sum_i L(z_i, \theta)$$
>
> and
>
> $$\hat{\theta}_{\epsilon,j} = \arg\min_\theta \frac{1}{n}\sum_i L(z_i, \theta) + \epsilon L(z_j, \theta).$$
>
> Following the strict convexity assumption, both of these values are unique (we consider $\epsilon > -1/n$ s.t. all terms in the sum are strictly convex). Fermat's theorem tells us that every extremum of a differentiable function is a stationary point. Thus, a necessary condition for the optimality of a differentiable function is that the gradient at the optimum must

be 0. (This is not sufficient, however: stationary points can also be maxima and saddle points.) Therefore, the previous optimality assumptions imply

$$\nabla_\theta \left(\frac{1}{n} \sum_i L(z_i, \hat{\theta}) \right) = \frac{1}{n} \sum_i \nabla_\theta L(z_i, \hat{\theta}) = 0$$

(because the gradient is a linear operator) and

$$\nabla_\theta \left(\frac{1}{n} \sum_i L(z_i, \hat{\theta}_{\epsilon,j}) + \epsilon L(z_j, \hat{\theta}_{\epsilon,j}) \right) = \frac{1}{n} \sum_i \nabla_\theta L(z_i, \hat{\theta}_{\epsilon,j}) + \epsilon \nabla_\theta L(z_j, \hat{\theta}_{\epsilon,j}) = 0.$$

These are ingredients (1) and (2).

We also make use of the Implicit Function Theorem. $\hat{\theta}_{\epsilon,j}$ is differentiable w.r.t. ϵ at $\epsilon = 0$. (The optimal point of the modified loss is also differentiable w.r.t. another variable of that function.) Therefore, one can consider the first-order Taylor expansion (by linearizing $\hat{\theta}_{\epsilon,j}$ in ϵ around $\epsilon = 0$):

$$\hat{\theta}_{\epsilon,j} = \underbrace{\hat{\theta}}_{\hat{\theta}_{\epsilon,j}\big|_{\epsilon=0}} + \epsilon \underbrace{\frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0}}_{\in \mathbb{R}^d} + o(\epsilon).$$

This is ingredient (3).

- We use linearization often, just like when attributing to test input features.
- $f(\epsilon) = f(0) + (\epsilon - 0) \cdot \frac{\partial f(0)}{\partial \epsilon} + o(\epsilon)$ is the Taylor expansion of $f(\epsilon) := \hat{\theta}_{\epsilon,j}$ around $\epsilon = 0$.
- $o(\epsilon)$ specifies $\lim_{\epsilon \to 0} \frac{R_1(x)}{\epsilon} = 0$. The remainder term converges to 0 faster than ϵ itself.

We compute $\nabla_\theta L(z_i, \hat{\theta}_{\epsilon,j})$ in terms of $\nabla_\theta L(z_i, \hat{\theta})$ as follows (by plugging in ingredient (3)). We calculate the Taylor expansion of $\nabla_\theta L(z_i, \hat{\theta}_{\epsilon,j})$ in θ, around $\hat{\theta}$.

$$\nabla_\theta L(z_i, \hat{\theta}_{\epsilon,j}) = \nabla_\theta L \left(z_i, \hat{\theta} + \epsilon \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon) \right)$$

$$= \nabla_\theta L(z_i, \hat{\theta}) + \nabla_\theta^2 L(z_i, \hat{\theta}) \left(\epsilon \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon) \right) + o(\epsilon)$$

$$= \nabla_\theta L(z_i, \hat{\theta}) + \epsilon \underbrace{\nabla_\theta^2 L(z_i, \hat{\theta})}_{\in \mathbb{R}^{d \times d}} \underbrace{\frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0}}_{\in \mathbb{R}^d} + o(\epsilon)$$

- This formulation is, of course, given to later get rid of the $o(\epsilon)$ terms and provide an approximation. The approximation is justified because $\hat{\theta}_{\epsilon,j}$ is very similar to $\hat{\theta}$ anyways for small ϵ. The difference is

$$\epsilon \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon).$$

We plug

$$\nabla_\theta L(z_i, \hat{\theta}_{\epsilon,j}) = \nabla_\theta L(z_i, \hat{\theta}) + \epsilon \nabla_\theta^2 L(z_i, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon)$$

into the second ingredient

$$\frac{1}{n} \sum_i \nabla_\theta L(z_i, \hat{\theta}_{\epsilon,j}) + \epsilon \nabla_\theta L(z_j, \hat{\theta}_{\epsilon,j}) = 0.$$

This results in

$$\frac{1}{n} \sum_i \left(\nabla_\theta L(z_i, \hat{\theta}) + \epsilon \nabla_\theta^2 L(z_i, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon) \right)$$

$$+ \epsilon \left(\nabla_\theta L(z_j, \hat{\theta}) + \epsilon \nabla_\theta^2 L(z_j, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + o(\epsilon) \right) = 0$$

$$\iff \frac{1}{n} \sum_i \nabla_\theta L(z_i, \hat{\theta}) + \epsilon \frac{1}{n} \sum_i \nabla_\theta^2 L(z_i, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0}$$

$$+ \epsilon \nabla_\theta L(z_j, \hat{\theta}) + \epsilon^2 \nabla_\theta^2 L(z_j, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + (\epsilon+1)o(\epsilon) = 0$$

$$\iff \frac{1}{n} \sum_i \nabla_\theta L(z_i, \hat{\theta}) + \epsilon \frac{1}{n} \sum_i \nabla_\theta^2 L(z_i, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + \epsilon \nabla_\theta L(z_j, \hat{\theta}) + o(\epsilon) = 0$$

$$\overset{(1)}{\iff} \epsilon \frac{1}{n} \sum_i \nabla_\theta^2 L(z_i, \hat{\theta}) \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + \epsilon \nabla_\theta L(z_j, \hat{\theta}) + o(\epsilon) = 0$$

$$\iff \epsilon H_{\hat{\theta}} \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + \epsilon \nabla_\theta L(z_j, \hat{\theta}) + \underbrace{o(\epsilon)}_{\lim_{\epsilon \to 0} \frac{R_1(\epsilon)}{\epsilon} = 0} = 0$$

$$\iff H_{\hat{\theta}} \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + \nabla_\theta L(z_j, \hat{\theta}) + \underbrace{\frac{o(\epsilon)}{\epsilon}}_{\lim_{\epsilon \to 0} \frac{R_1(\epsilon)}{\epsilon^2} = 0} = 0$$

$$\iff H_{\hat{\theta}} \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + \nabla_\theta L(z_j, \hat{\theta}) + o(\epsilon^2) = 0$$

$$\overset{\epsilon \text{ small}}{\implies} H_{\hat{\theta}} \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} + \nabla_\theta L(z_j, \hat{\theta}) \approx 0$$

$$\iff \frac{\partial \hat{\theta}_{\epsilon,j}}{\partial \epsilon}\bigg|_{\epsilon=0} \approx -H_{\hat{\theta}}^{-1} \nabla_\theta L(z_j, \hat{\theta})$$

Generally, the focus of research in influence function computation is how to speed things up while keeping approximations accurate. Interestingly, one can speed things up a lot.

3.11.3 LISSA

LISSA [3] is a method the authors of "Understanding Black-box Predictions via Influence Functions" [104] use to keep the inverse average Hessian calculation tractable. The LISSA algorithm uses an iterative approximation to approximate the inverse Hessian vector product

(iHVP):

$$H_{\hat{\theta}}^{-1} \nabla_\theta L(z, \hat{\theta}).$$

(Note the symmetricity of the average Hessian matrix.) For each test point of interest, they can precompute the above vector, and then they can efficiently compute the dot product between it and

$$\nabla_\theta L(z_i, \hat{\theta})$$

for each training sample z_i. This also helps with the quadratic scaling of the size of the average Hessian, as instead of computing the inverse average Hessian directly, they approximate the Matrix-Vector (MV) product through the iterative procedure.

The iterative approximation uses the fact that

$$A^{-1} = \sum_{k=0}^{\infty} (I - A)^k$$

for an invertible matrix A with all eigenvalues bounded below 1. At the small cost of inaccuracy, we gain a lot of speedup by this. The authors have another speedup by subsampling the training data in the summation (like how we do SGD-based optimization):

$$H_{\hat{\theta}} \approx \frac{1}{|I|} \sum_{i \in I} \nabla_\theta^2 L(z_i, \hat{\theta}).$$

Averaging Hessians through all training samples is infeasible. If we have a good representation of our training samples, we do not need to do a complete pass through the training samples. Random samples very likely give a good representation.

The final procedure estimates the inverse Hessian-Vector Product (HVP) as

$$H_i^{-1}v = v + (I - H_{\hat{\theta}})H_{i-1}^{-1}v,$$

where $H_{\hat{\theta}}$ is approximated on random batches (of size one or a small enough size), $v = \nabla_\theta L(z, \hat{\theta})$, and $i \in [t]$ is a particular iteration of the method ($H_0^{-1}v = v$). Using this technique, the authors reduce the time complexity of computing IF(z_j, z) for all training points and a single test point to $\mathcal{O}(np + rtp)$ where r is the number of independent repeats of the iterative HVP calculation (where they average the results from the r runs) and t is the number of iterations.

Note: LISSA already existed before the seminal IF paper – the authors adapted it to their method.

3.11.4 Arnoldi

Arnoldi, introduced in the paper "Scaling Up Influence Functions" [172], is a method for speeding up influence function calculations and reducing its memory requirements. Calculating and keeping a billion-dimensional vector (number of parameters) $H_{\hat{\theta}}^{-1} \nabla_\theta L(z, \hat{\theta})$ in memory is still very restrictive, and very coarse approximations (e.g., considering only a subset of pa-

rameters) are needed. If we consider the diagonalized formula for IF, $H_{\hat{\theta}}$ is written as

$$H_{\hat{\theta}} = Q\Lambda Q^{\top}$$

and the formula is

$$\mathrm{IF}(z_j, z) = \left\langle Q^{\top}\nabla_{\theta}L(z, \hat{\theta}), Q^{\top}\nabla_{\theta}L(z_j, \hat{\theta}) \right\rangle_{\Lambda^{-1}}.$$

Here, $Q \in \mathbb{R}^{p\times p}$ which is infeasibly large for efficient use. Setting $G = Q^{\top}$ to contain the k eigenvectors of $H_{\hat{\theta}}$ that correspond to its largest eigenvalues as rows (i.e., it is the projection matrix onto the span of the "top k eigenvectors"), we obtain

$$\mathrm{IF}(z_j, z) = \left\langle G\nabla_{\theta}L(z, \hat{\theta}), G\nabla_{\theta}L(z_j, \hat{\theta}) \right\rangle_{\Lambda_k^{-1}}$$

where

$$H_{\hat{\theta}} \approx G^{\top}\Lambda_k G.$$

By using this formulation, we map the gradients to a much lower-dimensional (k) space, and computations (dot product) become notably faster. G and Λ_k are calculated once and then cached.

The top k eigenvalues of $H_{\hat{\theta}}$ are the smallest k eigenvalues of its inverse, thus we take the eigenvalues that have the least influence in calculating IF. Very curiously, the authors report that selecting the top k eigenvalues of the inverse (corresponding to the dominant terms of the dot product) performs worse. DNN loss landscapes are highly non-convex, and the Hessian can have negative eigenvalues. The authors select the top k eigenvalues *in absolute value*.

The actual Arnoldi method is much more detailed and sophisticated in obtaining the referenced matrices, but the main idea is the same as was introduced here.

> **Additional Information: Using a Subset of the Parameters**
>
> Instead of the entire model, one can also use only the final or initial layers for θ. This dramatically reduces p by orders of magnitude. It has two drawbacks: the choice of layers becomes a hyperparameter, and the viable values of the number of parameters kept will depend on the model architecture. Using just one layer can result in different influence estimates than those based on the whole model. This is deemed suboptimal and is not used in Arnoldi but was used in earlier work, e.g., [104].

> **Additional Information: Reducing the Search Space**
>
> The following speedup is also compatible with Arnoldi, although the authors do not use it. (They use Arnoldi for retrieval of wrong labels. It is not aligned with the goal.) It is used in FastIF.
>
> IF(z_j, z) is already quite expensive. We should not calculate it for all j. The end goal is usually to retrieve influential training samples z_j for the test sample z. Instead of computing

IF(z_j, z) for all training samples z_j, we first reduce the search space (for candidate training samples that are likely to influence our test samples) via cheap, approximate search. This greatly reduces computational load. For example, we can perform L_2-distance k-NN using the last layer features to retrieve candidates. This is a typical trick for deep metric learning: We take the last layer activations from a network as a good representation of our sample in a lower-dimensional space. We compute the Euclidean distance in this space and retrieve the top k semantically similar samples from the training set to the test sample. Instead of taking the top k samples, we could also threshold by the L_2 radius. Both ways reduce the search space by a lot; thus, they reduce computational time.

3.11.5   LISSA vs. Arnoldi

Table 3.4: "Retrieval of mislabeled MNIST examples using self-influence for larger CNN. For TracIn the C value is in brackets (last or all). All methods use full models (except the LISSA run on 10% of parameters Θ)." [172] TracIn[10] gives the best results while staying feasible to compute. RandProj is also a surprisingly strong method. Table is adapted from [172].

Method	$\tilde{p}$	T, secs	AUC	AP
LISSA, $r = 10$	-	4900	98.9	95.0
LISSA, $r = 100$ (10% Θ)	-	32300	98.8	94.8
TracIn[1]	-	5	98.7	94.0
TracIn[10]	-	42	**99.7**	**98.7**
RandProj	10	0.2	97.2	87.7
RandProj	100	1.9	98.6	93.9
RandSelect	10	0.1	54.9	31.2
RandSelect	100	1.8	91.8	72.6
Arnoldi	10	0.2	95.0	84.0
Arnoldi	100	1.9	98.2	92.9

Results [172] of Arnoldi and various other methods are given in Table 3.4. The Arnoldi authors compute AUC and AP for the retrieval of wrong labels. They try to retrieve the wrongly put labels in the training set using self-influence. The task is not exactly aligned with removing training samples and retraining (precise estimation of IF) – this is why RandProj can also perform quite well. In fact, it performs better than Arnoldi. (It does not need to give precise IF estimates!) In RandProj, G is a random Gaussian matrix: It does not correspond to the eigenvectors of the top k eigenvalues. Eigenvalues are all considered to be one. In RandSelect, the eigenvalues are also all considered to be one, and we select the (same) elements of the two gradient vectors randomly. It needs a much larger k than RandProj. Arnoldi is $10^3 - 10^5$ faster than LISSA while being only a couple of percent worse on AUC and AP. TracIn [158] performs best on AUC and AP, but Arnoldi and RandProj are an order of magnitude faster.

3.11.6   TracIn

Are we asking the right question in the previous set of methods to attribute to training samples? We only measure the change in the loss w.r.t. the optimal model, given an infinitesimal
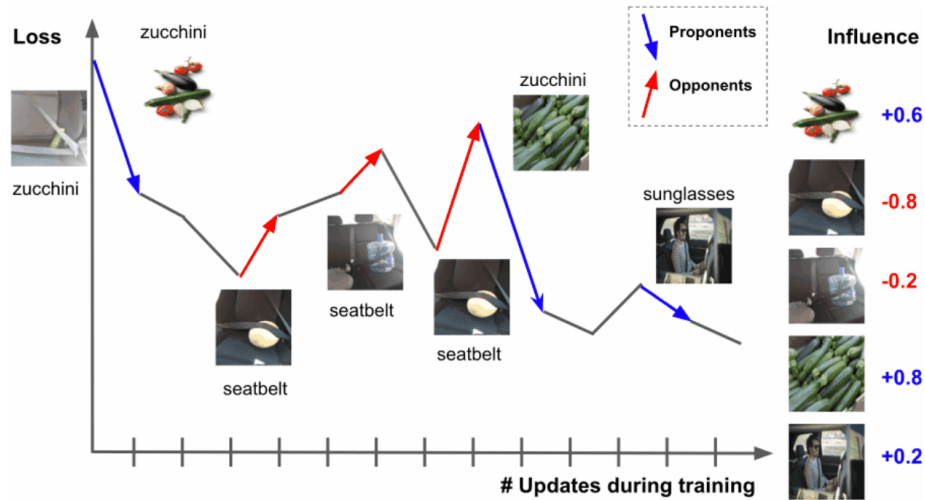
Figure 3.55: Loss value of a 'zucchini' sample over the course of training. The initial loss value for the test sample at the beginning of training is shown on the left. Test losses are evaluated *after* being presented with the shown images. Different training samples have a different impact on the test loss of interest. When we have a similar image but the label is different, the loss goes up for the test sample. Samples that increase the loss are called opponents to the test sample of interest. Samples that decrease the loss are the proponents of a test sample. We can find proponents and opponents for each training sample during training by considering the changes in the loss value. The 'Zucchini' training image results in a lower test loss, as the model learns to detect zucchinis better. The 'Sunglasses' training image is similar to the seatbelt images (car interior shown) but has a non-car-related label. The model has to focus on a small part of the image to predict correctly. This implicitly helps the prediction of 'zucchini'. **Note**: During training, the general trend of the loss should be downwards, but because of the noisy behavior of SGD (e.g., update after every training sample) and the possibility of overfitting, the test loss of sample z does not have to decrease at every gradient step.

change in the weight of one of the training samples. This sounds super naive and irrelevant in practice. Who would want to introduce an infinitesimal change in the weight of one of the samples?

TracIn, presented in the paper "Estimating Training Data Influence by Tracing Gradient Descent" [158] is another approach from 2020: We decompose the final test loss $L(z, \theta_T)$ of the trained model θ_T minus the baseline loss $L(z, \theta_0)$ of the randomly initialized model θ_0 into contributions from individual training samples. This is *global linearization* of the final test loss w.r.t. the update steps.[28]

Figure 3.55 gives an intuitive introduction to TracIn – it considers *useful* and *harmful* examples. *TracIn is the Integrated Gradients for training sample attribution.*[29]

The loss for the test sample at final iteration T can be written as the telescopic sum (i.e., everything cancels):

$$L(z, \theta_T) = L(z, \theta_0) + (L(z, \theta_1) - L(z, \theta_0)) + \cdots + (L(z, \theta_T) - L(z, \theta_{T-1})).$$

It is the sum of the original loss value and the loss differences between consecutive parameter

---

[28]It is global because one can use this linearization for *any* test sample.

[29]This is because the linearization the method admits is very similar to the one of Integrated Gradients.

updates.

Let us first consider the case when θ is updated for every single training sample z_j (batch size $= 1$). There is a clear, unique assignment of which training sample affects which parameter update step. (We never do this in practice, but we assume this for simplicity.) Then there is a natural notion of contribution of z_j to the test loss $L(z, \theta_T)$:

$$\text{TracInIdeal}(z_j, z) = \sum_{t:z_j \text{ used for } \theta_t \text{ update}} L(z, \theta_t) - L(z, \theta_{t-1}).$$

The summation is over the changes of loss for test sample z, where the parameter update was done by training on a sample of interest z_j. There will be millions/billions of iterations. We want to determine which of these iterations corresponds to the training sample of interest; then, we sum up these differences. This results in a completeness property (refer back to the Integrated Gradients method for test feature attribution):

$$L(z, \theta_T) - L(z, \theta_0) = \sum_{t=1}^{T} L(z, \theta_t) - L(z, \theta_{t-1})$$

$$= \sum_{j=1}^{n} \sum_{t:z_j \text{ used for } \theta_j \text{ update}} L(z, \theta_t) - L(z, \theta_{t-1})$$

$$= \sum_{j=1}^{n} \text{TracInIdeal}(z_j, z).$$

This follows from each time step corresponding to a unique training sample. Thus, we have a decomposition of (final loss - baseline loss) into individual contributions.

Of course, the critical issue with this formulation is that, in practice, we update models on a *batch* of training samples. When there is a parameter update, it is hard to attribute the change in loss (due to the update) to individual training samples in the batch. Many training samples are involved in the difference $L(z, \theta_t) - L(z, \theta_{t-1})$.

Each parameter update with SGD looks as follows:

$$\theta_{t+1} = \theta_t - \frac{\eta_t}{|B_t|} \sum_{i:z_i \in B_t} \nabla_\theta L(z_i, \theta_t)$$

where η_t is the learning rate at step t and $|B_t|$ is the size of the batch at step t. This is usually kept fixed, but we often do not drop the last truncated batch that has a smaller size. We average the gradients over the batch. We have a nice decomposition of the parameter update steps as a sum of individual training-sample-wise gradients for the loss (in the batch).

We rewrite the loss $L(z, \theta_{t+1})$ with parameters from time step $t+1$ as

$$L(z, \theta_{t+1}) = L\left(z, \theta_t - \frac{\eta_t}{|B_t|} \sum_{i:z_i \in B_t} \nabla_\theta L(z_i, \theta_t)\right)$$

$$= L(z, \theta_t) + \left(-\frac{\eta_t}{|B_t|} \sum_{i:z_i \in B_t} \nabla_\theta L(z_i, \theta_t)\right)^\top \nabla_\theta L(z, \theta_t) + o(\eta_t)$$

where we performed a Taylor expansion of $L(z, \theta_{t+1})$ around $\eta_t = 0$ ($f(\eta_t) = f(0) + \eta_t f'(0) +$

$o(\eta_t)$) or around θ_t ($f(\theta_{t+1}) = f(\theta_t) + (\theta_{t+1} - \theta_t)\nabla_\theta f(\theta_t) + o(\theta_{t+1} - \theta_t)$). We can choose to do both because they have a linear relationship. This is an *accurate approximation* because we are using a small learning rate ($1\mathrm{e}{-}3$), so θ_{t+1} is close to θ_t.

Therefore,

$$L(z, \theta_t) - L(z, \theta_{t+1}) \approx \frac{\eta_t}{|B_t|} \sum_{i:z_i \in B_t} \nabla_\theta L(z_i, \theta_t)^\top \nabla_\theta L(z, \theta_t).$$

In words, the difference in loss values before and after the update is approximately equal to some constant times a summation over dot products of training sample gradients with the test sample gradient. There is a natural decomposition of the contribution of individual samples in the batch towards the difference in the loss. When this difference is a large positive number, it means the batch samples were useful for the test sample z. This is the particular reason why we "flip the sign" and choose to model $L(z, \theta_t) - L(z, \theta_{t+1})$.

A natural notion of the contribution of sample z_j towards the difference in losses for this particular update is given by

$$\frac{\eta_t}{|B_t|} \nabla_\theta L(z_j, \theta_t)^\top \nabla_\theta L(z, \theta_t)$$

when z_j is included in batch B_t for updating θ and 0 otherwise. Using this approach, we make attributing to individual training samples feasible in practice. This is a constant times the dot product between the test sample of interest gradient and the training sample of interest gradient. This is similar to what we have seen in the previous methods.

Summing over the entire trajectory of model updates, we define the contribution of z_j towards the loss for z as

$$\mathrm{TracIn}(z_j, z) = \sum_{t:z_j \in B_t} \frac{\eta_t}{|B_t|} \langle \nabla_\theta L(z_j, \theta_t), \nabla_\theta L(z, \theta_t) \rangle.$$

This is the final definition of TracIn, the trajectory-based influence of sample z_j towards test sample z. It is simply a summation of all parameter update steps t that contained z_j in the batch. These are the only relevant terms, the others are 0. The smaller the loss becomes on test sample z between steps t and $t + 1$, the more we attribute those training samples that were in the batch of step t.

3.11.7 TracIn vs. IF

These two methods have very similar formulations but also some key differences.

$$\mathrm{IF}(z_j, z) = \left\langle Q^\top \nabla_\theta L(z, \hat\theta), Q^\top \nabla_\theta L(z_j, \hat\theta) \right\rangle_{\Lambda^{-1}}$$
$$\mathrm{TracIn}(z_j, z) = \sum_{t:z_j \in B_t} \frac{\eta_t}{|B_t|} \langle \nabla_\theta L(z_j, \theta_t), \nabla_\theta L(z, \theta_t) \rangle$$

Both use a form of a dot product between parameter gradients for the training and test samples. It is quite impressive that the final formulations end up being so simple, but it is a natural byproduct of linearization. TracIn sums over training iterations (checkpoints) and does not use a Hessian-based distortion of the dot product (to squeeze/expand some of the eigenbasis directions). It is, therefore, cheaper because we do not need to compute the Hessian.

However, it is very memory intensive. In contrast, IF considers only the final[30] parameter and distorts the space using the average Hessian.

Using IF, we are missing out on all contributions on the way during training. Intuitively, TracIn makes more sense, but it is hard (if not impossible) to say which one is better conceptually. We can only use empirical evaluation to tell which serves our purpose better. IF has many more assumptions that are also violated in practice. The method considers globally optimal parameter configurations, strict convexity, and twice-differentiability. In practice, the eigenvalues could also become negative (saddle point) or $0 \implies$ invertibility does not hold when we have a 0 eigenvalue (i.e., the loss is constant in some directions). In theory, this can happen during optimization. A small epsilon has to be added.

3.12   Evaluation of Attribution to Test Samples

There are two perspectives of evaluation of such methods: (1) comparing approximate values against their GT counterparts and (2) evaluating such attribution methods based on some end goals/downstream tasks.

3.12.1   Comparison of Approximation Against GT Value

IF approximates the remove-and-retrain algorithm (remove a certain training sample, retrain, and see how much that influences the loss value for the test sample of interest). One can measure *soundness* by comparing influence values against the actual remove-and-retrain baseline. This is an evaluation of soundness. To see the correspondence, consider the first-order Taylor approximation of $L(z, \hat{\theta}_{\epsilon,j})$ again around $\epsilon = 0$. To avoid confusion, we stick to the definition of IF where a larger positive value signals positive influence.[31] We have

$$L(z, \hat{\theta}_{\epsilon,j}) - L(z, \hat{\theta}) = \epsilon \underbrace{\left. \frac{\partial L(z, \hat{\theta}_{\epsilon,j})}{\partial \epsilon} \right|_{\epsilon=0}}_{-\operatorname{IF}(z_j, z)} + o(\epsilon).$$

The notion of removal is equivalent to setting $\epsilon = -1/n$, which is generally a very small number, so the linear approximation stays reasonably faithful to the actual loss function. We finally obtain

$$L(z, \hat{\theta}_{\setminus j}) - L(z, \hat{\theta}) \approx \frac{1}{n} \operatorname{IF}(z_j, z).$$

To benchmark IF on how faithful it is to the remove-and-retrain algorithm, we can compare the left quantity to the right one. We might be interested in removing not just one sample but a group of them. This is not modeled by the most naive version of remove-and-retrain that IF approximates.

Soundness results of IF are shown in Figure 3.56. We can also try leaving a group of samples out from the training set and seeing how the model reacts regarding the change in the loss for a test sample. This is shown in Figure 3.57 from the FastIF paper [70], which is yet another paper on how we can speed up IF computations. The task is MNLI, a 3-class natural language

---

[30]Strictly speaking, IF considers the globally optimal parameter configuration in the formulation.
[31]Interestingly, the FastIF paper considers the original IF definition without flipping the sign.
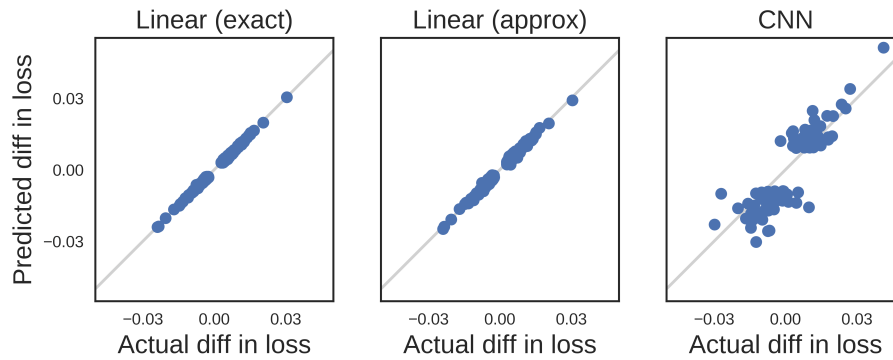
Figure 3.56: Comparison of the predicted difference in loss after the removal of a sample (i.e., the IF value) against the actual difference in loss. Figure taken from [104]. The benchmark measures faithfulness to leave-one-out retraining on MNIST. For every training and test sample, we measure the change in test loss by actually removing a training sample. This is quite fast for a linear model. We also have the predicted difference in loss through the IF computation. We compare the two values. *Left.* The gradient-based approximation of the influence (times $1/n$ to model removal) gives nearly the same result as the actual difference in the loss for the linear model (logistic regression). It is a good sanity check that the exact Hessian computation performs well approximation-wise. *Middle.* "Linear (approx)" still considers logistic regression but uses the LISSA approximations to speed up the Hessian computation. Even if we use LISSA to approximate the average Hessian, we do not lose much accuracy. *Right.* To evaluate on CNNs, we must take a leap of faith. The logistic regression optimization is strictly convex, but the CNN one is, of course, not. They apply the method to a small CNN on MNIST. We can see some correlation, but many things are seemingly not working anymore. Two groups follow the overall trend, but we do not see much correlation between the actual and the predicted value *within* each group. We have mixed results.

inference task with classes entailment, neutral, and contradiction. The group of samples we remove is determined by the influence values (we sort all training samples according to their influence values). The influence value has parity: it can be positive or negative. Using this book's IF definition, positive means that including the sample helps, and negative means that by including this training sample ($\epsilon > 0$), we are increasing the loss. The general trend is that removing helpful samples increases the loss. (It is harmful to remove the samples with a high IF value.) Similarly: removing harmful samples decreases the loss. (It is useful to remove the samples with a low IF value.) By just removing random samples, we do not see much change in the test loss. "Full" means we use the entire dataset. The KNN versions correspond to selecting representative samples from the training set. We can see that this can even be beneficial.

3.12.2 Focus on the End Goal: Mislabeled Training Data Detection

IF and TracIn are eventually serving certain end goals. Remove-and-retrain may not be very useful as the end goal. For example, when the actual end goal is to debug/improve the model/dataset, faithfulness to the remove-and-retrain algorithm is not of particular interest. It is just an intermediate step (a proxy) for using the method for improving models.[32] We need to evaluate based on more reasonable end goals, e.g., mislabeled training data detection.

---

[32]Suppose that a self-driving car killed a pedestrian. We need to find out which data sample was responsible for the incorrect (sequence of) predictions. Remove-and-retrain is not the end goal in this case, we do not care about how well we approximate it or whether we even approximate it at all.
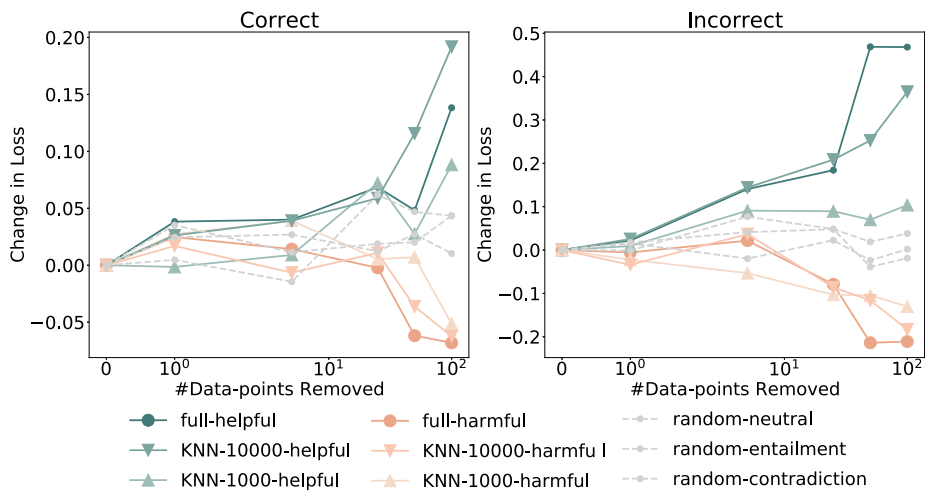
Figure 3.57: Leave-M-out results on MNLI [209]. "Change in loss on the data point after re-training, where we remove $m_{\text{remove}} \in \{1, 5, 25, 50, 100\}$ data-points [either positives or nega-tives]. We can see that the fast influence algorithms [(the KNN versions)] produce reasonable quality estimations at just a fraction of computation cost." [70] Correct and incorrect mean that the original predictions were correct/incorrect. Figure taken from [70].

We will see how people can use influence functions to detect mislabeled training samples. Checking, e.g., how well our method approximates remove-and-retrain might be good to check whether our proposed idea works. Then, we evaluate the method using the actual end goal. This can change the conclusion of which method is better for us.

> **Definition 3.16: Self-Influence**
>
> Self-influence is a metric used in training sample attribution methods that measures how much contribution a particular training sample z_j has to its own loss.
>
> **Example**: Using influence functions, the self-influence score for sample z_j is $\text{IF}(z_j, z_j)$. Using TracIn, we can use $\text{TracIn}(z_j, z_j)$ as a self-influence score.

We make use of self-influence scores for mislabeled training data detection. If a sample is one of its kind, then it only has itself to decrease its loss, therefore, we expect a high self-influence score. Looking at that exact sample is the only way to decrease the loss of that sample. On the other hand, if the sample is just like other data points in the training set, then it is among many that decrease its loss, therefore, we expect a low self-influence score: Including it or not has little influence. Mislabeled data are typical examples of "one of its kind" data. As such, we expect high self-influence scores for them.[33] By measuring self-influence, we should be able to tell which samples are mislabeled.

Figure 3.58 shows benchmark results on mislabeled training data detection from the TracIn paper. Mislabeled training data detection is a typical binary detection task: we want to classify mislabeled/not mislabeled. We know the ground truth in the benchmark; we try to retrieve the mislabeled ones in the training set. We can use detection metrics like AUROC and AP (AUPR), which are typical evaluation scores for retrieval tasks. Our only feature is the attri-

---

[33]Assuming that the training set has many correctly labeled data and a few mislabeled data points (i.e., there is no systematic mislabeling).
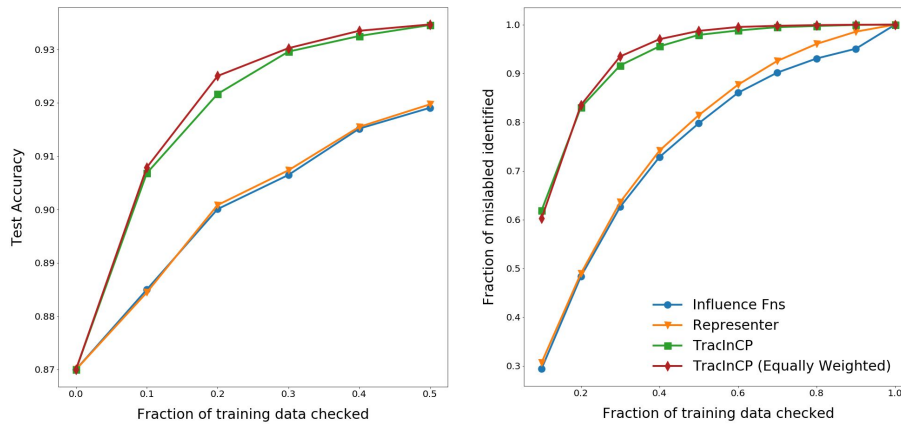
Figure 3.58: Results of using self-influence to detect mislabeled training samples on CIFAR. *Left*. Fixing the mislabeled data found within a certain fraction of the training data results in a larger improvement in test accuracy for TracIn compared to the other methods. *Right*. TracIn retrieves mislabeled samples much better than IFs. Figure taken from [158].

bution score. The question is, "Is there a threshold that is extremely good for separating mislabeled samples from not mislabeled?" In this benchmark, however, they are *not* doing that. Instead, they sort training samples according to self-influence values and then decrease the threshold from top to bottom and see how many mislabeled samples are retrieved.
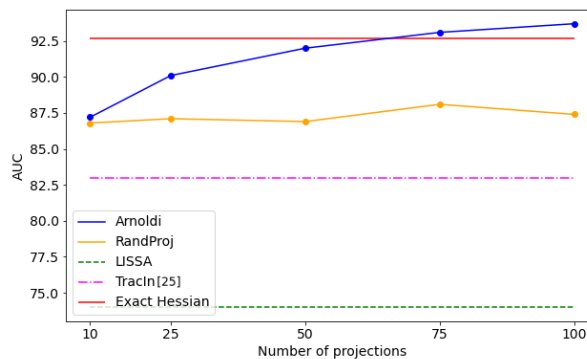
Figure 3.59: Results of using self-influence to detect mislabeled training samples on MNIST using a small CNN. AUC for retrieval of mislabeled MNIST examples as a function of the number of eigenvalues (projections), $\tilde{p}$. Figure taken from [172].

We also discuss using self-influence to detect mislabeled training samples on MNIST. The results are shown in Figure 3.59. The task is not perfectly aligned with IF computation: the exact method can be surpassed.

Finally, we discuss the retrieval of mislabeled MNIST examples using self-influence for a larger CNN. As discussed before, AUC and AP are usual detection metrics for mislabeled samples. The results are shown in Table 3.4.

3.13 Applications of Attribution to Test Samples
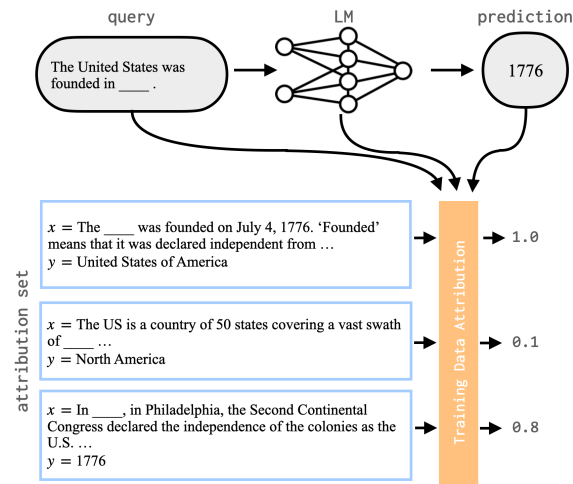
Figure 3.60: Illustration of using training data attribution scores for fact tracing. Figure taken from [5].

Fact Tracing

We discuss fact tracing, an important application of test sample attribution, as shown in the paper "Towards Tracing Factual Knowledge in Language Models Back to the Training Data" [5]. Suppose we have a language model that is trained to predict missing words using actual facts, and we have built a dataset with the GT fact attributions in the training set. Then we can measure fact retrieval performance: We evaluate any Training Data Attribution method on its ability to identify the so-called true proponents, i.e., the true training sample information sources. We want to retrieve the true proponents out of a large set of training examples, which is, again, a classical retrieval task. This is illustrated in Figure 3.60.

It is a natural question to ask the model, "Did you just make this up? Which training datum did you look at to make this decision?" Nowadays, fact tracing matters a lot, and training data influence can be readily used for it. LLMs are critical candidates for this method. We cannot be sure how it would scale, but it is something to keep an eye out for.

Membership Inference

Given a model and arbitrary data we give to the model, we wish to see whether that data was included in the training of that model.

- "Was this image used for training the DALL-E model?"
- "Was this image used for generating the current image that I got?"

Being able to answer such questions could be a nice tool for dealing with copyright issues for large-scale generative models. It would also be possible to use influence functions and training attribution in general. Suppose we had access to the training set. Then, we could use the scores to sort decreasingly and manually check whether the sample was used (soft filtering). Alternatively, if we are really searching for exact matches, we could search for matches according to the ordering given by influence function scores. Hashing already works for checking for exact matches very efficiently, but it would not work for matches that are not exact (e.g. when JPEG encoding/decoding is applied). There are only very few papers in this area

so far, but it is gaining traction. Large companies are probably also already working on this problem.

Chapter 4

Uncertainty

4 Uncertainty

4.1 Introduction to Uncertainty Estimation

Uncertainty is everywhere. Having complete information and a perfect understanding of a system can only happen in simple and closely controlled environments. The world around us is not such an environment. Humans learn to build complex internal models of uncertainty to cope with incomplete information and react robustly to events that either have not happened yet or are only partially observed.

Understanding, quantifying, and evaluating uncertainty is of crucial importance in our everyday lives, but also in fields specialized to cope with and leverage uncertainty. Examples include financial analysis, economic decisions, general statistics, probabilistic modeling, and also machine learning. Classical ML theory usually did not aim to *make systems know when they do not know* – the main goal was to find methods and solutions that work well, considering them as standalone components. These days, accuracy in most applications is not the biggest concern – most ML solutions provide reasonably good accuracy in several tasks. Instead, there is an ever-increasing demand to quantify sources of uncertainty in ML models and make them understand their own limitations. As we will soon see, uncertainty quantification is a crucial requirement whenever we want to incorporate an ML solution into a certain pipeline.

In the Uncertainty chapter of the book, we are going to further motivate the need for uncertainty estimation, quantify sources of uncertainty, consider methods that can give us different kinds of uncertainty estimations, and learn about methods to evaluate uncertainty predictions for DNNs.

4.1.1 Motivation

We first consider a meeting with another business, based on a real story of one of the authors when they were working at a company. Teams without ML knowledge tend to downplay the difficulty of doing technical things. There are always typical subjects in such meetings:

- "Why does your AI system not return how sure it is about the output?"
- "Is it not kind of trivial to make the system predict confidences?"
- "We cannot plug your system into our pipeline if there is no such estimate."
- "We really need it, cannot you just do it?"

Unfortunately, solving such tasks is not at all trivial. However, they are prevalent (as there are many such requests) and valid desires; we will see methods to achieve these goals.

4.1.2 Uncertainty estimation is a critical building block for many systems.

When an ML model is part of a bigger modular pipeline, uncertainty estimation is very beneficial and often required. For an ML-based data-driven module, it is not easy to trust everything the model outputs. Such models are never perfect and extra care is needed to use the model's prediction in downstream modules. This is also true when the downstream module in question is a human – people do not (and should not) trust every prediction of the model.

Let us suppose for a moment that the model already knows about itself how certain it is. We consider some example downstream use cases of reporting uncertainty (in later modules of the pipeline).

Human in the loop. We only want humans to intervene when the ML confidence is low, as human knowledge is expensive. When the model's confidence is low, the model can say, "I am not sure about the result." When humans need to intervene, they can take control and handle certain requests themselves (i.e., they can fix the model's prediction).

Risk avoidance. When there are great risks involved in the model's task, the ML system should only act when it is confident. If the model is unsure, the processing pipeline should stop (or fall back to some other safe state), as the situation is deemed too risky. An example of this is a learning-based manufacturing robot for cars. When the robot is uncertain about its next action, there is a high risk it is going to make a mistake which could also result in it damaging or destroying the car. We want the model to be able to say, "We should probably not take care of this input and just stop."
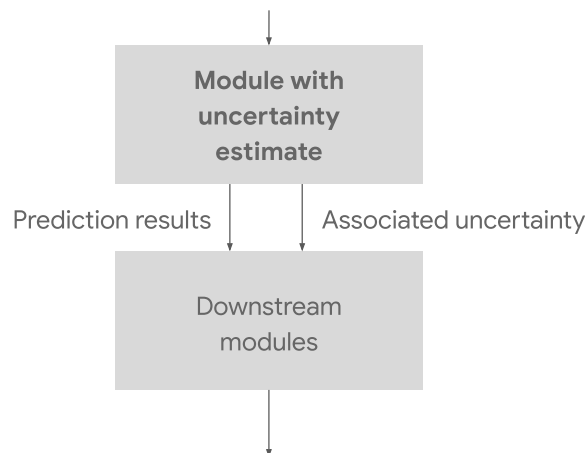
Figure 4.1: Simplified flowchart of the ideal integration of ML models into modular pipelines. In addition to the prediction results, we also wish to obtain associated uncertainty estimates to efficiently use the predictions in downstream tasks.

Thus, it is very beneficial for our model to output *two* predictions when it is part of a pipeline: the prediction results and also the associated uncertainty estimate(s), as illustrated in Figure 4.1. This gives us many more choices of what to do later in the pipeline in downstream modules.

When do we need confidence estimation?

In general, confidence estimation is needed when the outputs of a model cannot be treated equally – outputs for certain samples are more confident and some of them are less trustable.

It is not needed when the system always returns perfect answers. Why would we need it? If such a time would come when AI systems were always giving the right answers, this study would become useless. We will learn about whether that can happen... (Spoiler: It cannot, as in almost any sensible scenario, there is some level of stochasticity we cannot get rid of.)

4.1.3   Example Use Cases of Uncertainty Estimation

The following examples of uncertainty estimation are inspired by [115].

Image search for products

In this example, we do not consider the old Google image search. We consider products like Google Lens. Such products do not only search for similar images – they also take the user and context into account:

> "Google Lens is a set of vision-based computing capabilities that can understand what you're looking at and use that information to copy or translate text, identify plants and animals, explore locales or menus, discover products, find visually similar images, and take other useful actions. [...] Lens always tries to return the most relevant and useful results. Lens' algorithms aren't affected by advertisements or other commercial arrangements. When Lens returns results from other Google products, including Google Search or Shopping, the results rely on the ranking algorithms of those products." [65]

Companies are usually also very motivated to link image search results to actual products to make money. Customers can also get quick answers from such image search results.

Given an image, the task is to find the product that is shown. What should happen if the photo taken by the user is of poor quality? Regular algorithms would search for the most likely product anyway, which is usually a very poor suggestion. If the model is equipped with uncertainty estimates, when the confidence is low, it can

1. ask the user to take another photo, and/or
2. show different results from all products that could match with high probability.

What if the photo *does not contain* any product of interest? Again, regular algorithms would simply return poor results. Uncertainty estimation can allow the system to determine whether the provided photo is relevant. When there is no object of interest, the system can output suggestions such as "User should be posing the camera differently." or "Try to focus on an object of interest." This feedback loop can ensure that the model can perform correctly and does not mislead the user with unconfident predictions.

High-stake decision making

The prime example of a high-stake decision-making application is healthcare, where the model has to determine whether there is anything wrong with our body. We can use model

uncertainty to decide when to trust the model or defer to a human. This is a crucial ability of a model in general cost-sensitive decision-making, where mistakes can potentially have huge costs. Costs include potential lawsuits, the death of a patient, or fatal road accidents. The task is to provide a binary prediction of healthy/diseased from the input image. Ideally, the model should make a prediction and output confidence estimates as well. One should only trust the model's predictions when they are confident.[1] When the model is not confident enough, we defer to a human. For example, we can ask a human doctor to come in and take a look.

Table 4.1: Example cost table for decision making in healthcare. Predicting 'healthy' for a diseased person has the highest cost, as such cases can even lead to the death of a patient. Table recreated from [115].

| | | True Label | |
		Healthy	Diseased
	Predict Healthy	0	10
Action	Predict Diseased	1	0
	Abstain "I don't know"	0.5	0.5

In discrete cost-sensitive decision-making problems, we usually have *cost tables*, depicted in Table 4.1. We have a very high stake in false negative disease diagnoses. We incur huge costs. Thus, we want to predict 'healthy' only when the system is very certain, and we even prefer the answer 'I don't know' over predicting false positives. Predicting 'I don't know' defers to a human doctor. An example of such a scenario is diabetic retinopathy detection from fundus images [115], illustrated in Figure 4.2.
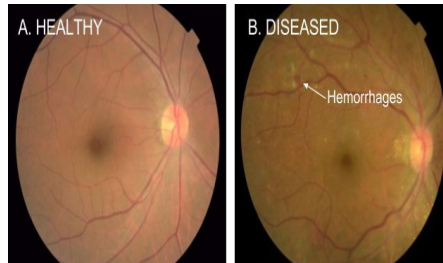
Diabetic retinopathy detection from fundus images
Gulshan et al, 2016

Figure 4.2: Diabetic retinopathy detection from fundus images. Predicting 'healthy' can be catastrophic if the patient is actually diseased. Figure taken from [115].

The field of self-driving cars also requires uncertainty estimates. It also qualifies as high-stake decision-making, as people's lives are at stake.[2] We do not want our current self-driving systems to drive *in all cases*. In self-driving scenarios, we often experience *dataset shift*. We want to make sure that our car does not crash in such cases. Examples include changes in

- time of day/lighting (driving at night vs. in the morning),
- geographical location (inner city vs. suburban location),
- weather conditions (thunderstorm vs. clear weather),

---

[1] For example, we might accept the model's prediction when the provided confidence estimate is above a certain tuned threshold.

[2] Self-driving and healthcare usually come in pairs when discussing high-stake ML use cases.

- or traffic conditions (traffic jams, construction sites, clear highways).

In such cases, we wish to take over control and drive responsibly. By using uncertainty estimation, the car can tell us when it is uncertain.

Open-set recognition

Open-set recognition is a different scenario that is more specific to the classification task. In the development (dev) stage (Section 2.3), we can pre-define a set of classes, e.g., the 100 most popular skin condition classes. When deploying in the real world, there can be very rare diseases as test inputs for which we do not have classes. If the model predicts 'normal skin' in such cases, it is very harmful. However, the other scenario is not better either: "Well, it does not look normal, but since I need to pick one from the known cases, I will just guess Acne." A classification system should also be able to say, "This is something I have not learned before, a new class. This is none of the above." This can either be an explicit class, or it can be signaled by low predictive confidence.

Open-set recognition considers different ways to deal with new classes in deployment. There are generally two variants of open-set recognition: models trained with or without OOD data. When they are trained with OOD data, they also usually contain a separate dimension in the output probability vector for indicating the probability of OOD (explicit introduction of the 'I don't know.' class). When they are trained only with known classes, there is no data to train this extra dimension and, therefore, it is not added. Even in this case, the model can be trained to predict calibrated uncertainty estimates that can then be used to determine the 'I don't know.' class in an implicit fashion. Of course, without explicit supervision, the latter case will likely produce worse results.
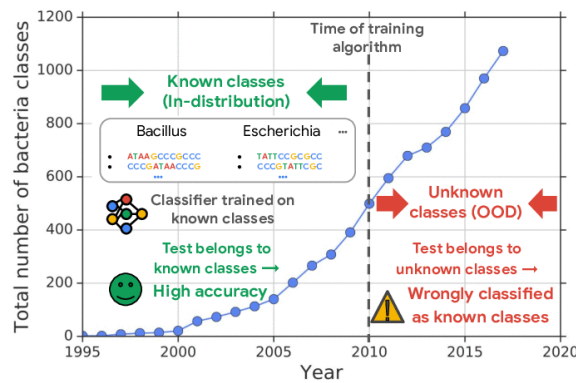
Figure 4.3: Example for the need for uncertainty estimation in the classification of genomic sequences. "A classifier trained on known classes [without proper uncertainty calibration] achieves high accuracy for test inputs belonging to known classes, but can wrongly classify inputs from unknown classes (i.e., out-of-distribution) into known classes with high confidence." [164] Figure taken from [164].

The same story goes for "growing field" cases. An example is the classification of genomic sequences. We discover more and more bacteria classes in biology research – new entries are coming to our database of bacteria. We usually have high ID accuracy on known classes, but this is not sufficient. We wish to be prepared for new bacteria classes in the future (unknown classes, OOD scheme), but we can only train on classes that are currently in the database. We

need to detect inputs that do not belong to any of the known classes. We wish to assign an 'I don't know.' label for future cases. This scenario is depicted in Figure 4.3. Samples predicted as 'I don't know.' can be used later on for further training the model: we can put labels on them once we discover them. For example, we can initialize a new row in the classifier layer's weight matrix, add a new bias scalar, and then we can predict one more output class after learning to predict such samples. The keyword here is *class-incremental learning*, which deals with efficiently increasing the number of classes over time without sacrificing the original classification score.

Active learning
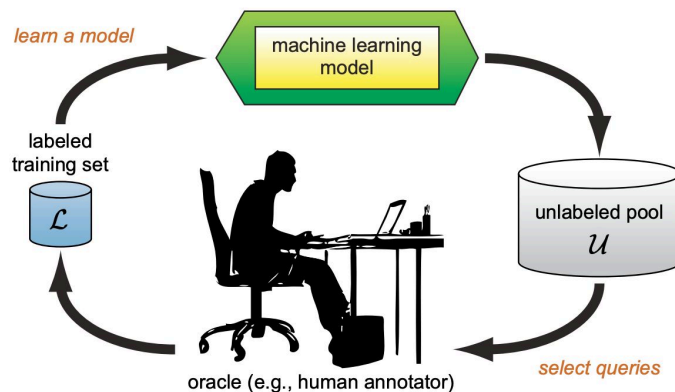
Image source: Active Learning Literature Survey, Settles 2010

Figure 4.4: General overview of active learning. We can get away with labeling significantly fewer samples for our model if we label the "right" ones. Figure taken from [115].

Active learning, illustrated in Figure 4.4, is concerned with finding samples to label smartly. Instead of going through a huge set of unlabeled samples to label everything, we pick the samples the model is very likely to be confused about, and then ask for human feedback on those samples in an iterative fashion. This way, we maximize the utility of humans (that are expensive). We can use model uncertainty to improve data efficiency and the model's performance in "blind spots". To tell which of the unlabeled samples is most likely to have the highest return when annotated by a human, we should rely on a notion of uncertainty and confidence values.

Hyperparameter optimization and experimental design

Hyperparameter optimization and experimental design are widely used across large organizations and the sciences. Such methods often employ *Bayesian optimization*. Examples include photovoltaics, chemistry experiments, AlphaGo, electric batteries, and material design.

The setup is as follows. We are searching through a huge (combinatorial) space of possibilities for configurations/settings. For example, in a very naive hyperparameter search for an ML model, we might have

$$5 \text{ learning rates} \times 4 \text{ numbers of layers} \times 5 \text{ net widths} \times 3 \text{ weight decays}$$
$$\times 10 \text{ augmentations} \times 3 \text{ numbers of epochs} \times 3 \text{ optimizers} = 27000$$

possible hyperparameter settings to iterate over. Usually, we have thousands or millions of possible combinations, even in quite simple cases. It is clearly infeasible to consider all possible configurations. Bayesian optimization reduces the uncertainty of performance in this complex landscape while also choosing performant configurations. By observing a few data points where the configurations were chosen smartly (i.e., considering the trade-off between uncertainty reduction and exploitation), it constantly updates its beliefs based on the training results of the well-studied configurations. This reduces uncertainty over time, and eventually, we find a configuration that will likely maximize our return. To explore the space most efficiently, we need a notion of uncertainty. An example use of Bayesian optimization for experimental design is shown in Figure 4.5.
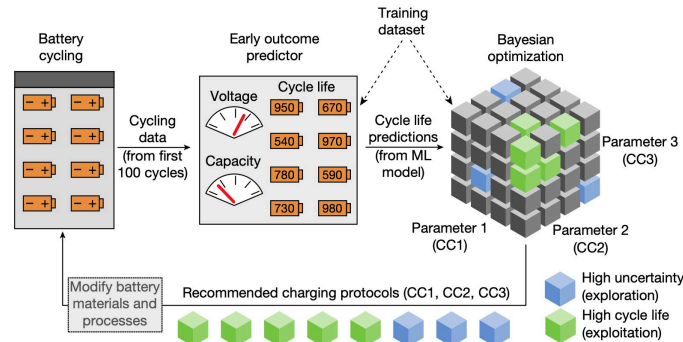
Image source: Attia et al. 2020 Closed-loop optimization of fast-charging protocols for batteries with machine learning

Figure 4.5: Role of uncertainty in optimizing battery charging protocols with ML. "First, batteries are tested. The cycling data from the first 100 cycles (specifically, electrochemical measurements such as voltage and capacity) are used as input for an early outcome prediction of cycle life. These cycle life predictions from a machine learning (ML) model are subsequently sent to a BO algorithm, which recommends the next protocols to test by balancing the competing demands of exploration (testing protocols with high uncertainty in estimated cycle life) and exploitation (testing protocols with high estimated cycle life). This process iterates until the testing budget is exhausted. In this approach, early prediction reduces the number of cycles required per tested battery, while optimal experimental design reduces the number of experiments required. A small training dataset of batteries cycled to failure is used to train the early outcome predictor and to set BO hyperparameters." [14] The linear model the predicts cycle life of a battery (and also gives a CI for the predictions). The GP relates protocol x to cycle life y through its internal parameters θ. Here, the GP outputs uncertainties naturally. Figure taken from [14].

Object detection pipeline

In object detection, we produce a bounding box and a class label for each object. Two-stage detectors (propose then refine) use multiple modules by construction. We will likely require confidence scores whenever we have multiple modules in any ML setting. Fast(er) R-CNN [61], the most popular object detection pipeline, is illustrated in Figure 4.6. In Faster R-CNN, we have the following stages.

1. Propose boxes with confidence scores. (Between 10^3 and 10^6 boxes are proposed.) This is the *objectness score*.
2. Prune boxes by thresholding w.r.t. the confidence/objectness scores. We return only the most likely boxes containing any objects. Then we further perform non-maximum suppression.
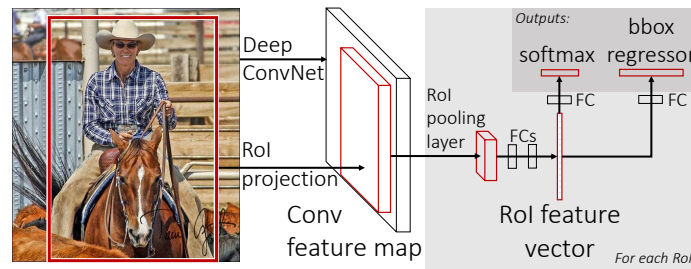
Figure 4.6: Fast(er) R-CNN is a renowned model in object detection. One of its distinguishing features is its modularity. When proposing bounding boxes for objects, referred to as Regions of Interest (RoIs), the method also provides a confidence or "objectness" score for each box. This score is crucial; it allows the system to prune less likely boxes before it refines and classifies the remaining ones, ensuring both accuracy and efficiency. Figure taken from [61].

3. Classify the pruned boxes and refine the boxes.

4.2 Types and Causes of Uncertainty

In this section, we aim to discuss different sources of uncertainty and how they relate to each other. In particular, we will discuss the terms *predictive uncertainty*, *epistemic uncertainty*, and *aleatoric uncertainty*. In the last paragraph of each of the subsections discussing these sources, we give an introduction to *how* we can evaluate these.

4.2.1 Predictive Uncertainty

> **Definition 4.1: Predictive Uncertainty**
>
> Predictive uncertainty refers to the degree of uncertainty or lack of confidence that a machine learning model has in its predictions for a given input.
>
> In particular, predictive uncertainty is typically referred to as the probability of the prediction's correctness. If for a fixed input sample x we define the indicator variable $L : \Omega \to \{0, 1\}$,
>
> $$L = \begin{cases} 1 & \text{if prediction } f(x) \text{ is correct} \\ 0 & \text{otherwise,} \end{cases}$$
>
> then predictive uncertainty is usually defined as
>
> $$c(x) = P(L = 1) = \text{probability that } f(x) \text{ is correct.}$$
>
> **Note**: Here, $f(x)$ denotes a single prediction from model f, *not* a distribution over predictions.

To summarize the above definition, predictive uncertainty tries to measure if we are likely to make an error in our prediction. Most of the ML uncertainty literature specifies two possible typical causes of predicting 'I am not sure.' – i.e., two main sources of predictive uncertainty.

First, we give an informal description of these two sources, and then we discuss them in more detail.

1. **Epistemic uncertainty**: "I am not sure because I have not seen it before."
2. **Aleatoric uncertainty**: "I have experienced it before, I know what I am doing, but I think there is more than one good answer to your question, so I cannot choose just one."

Evaluation always requires quantification – a quantified definition of the concept. Without evaluation, we cannot progress. How should we quantify whether a specific uncertainty estimate is reasonable? For discussing the basic evaluation of predictive uncertainty, we stick to the scalar confidence values introduced in the definition of predictive uncertainty, where we equate $c(x)$ to the probability of the prediction's correctness. Predictive uncertainty depends on both the model and the data. In particular, it increases both when the input is ambiguous and when the model is uncertain in its parameters (arising from the undefined behavior in no-data regions of the input space). Some evaluation metrics measure the true likelihood of the model failing and compare that to the given predictive uncertainty estimation. This is a direct way to benchmark predictive uncertainty estimates. In later sections, we will consider exact methods.

4.2.2 Aleatoric Uncertainty

> **Definition 4.2: Aleatoric Uncertainty**
>
> Aleatoric uncertainty is uncertainty that arises due to the inherent variability or randomness in the data or the environment. This type of uncertainty cannot be reduced by collecting more data or improving the model, as it is an intrinsic property of the system being modeled. Examples of sources of aleatoric uncertainty include measurement noise, natural variability in the data, or incomplete information.

Intuitively, aleatoric uncertainty translates to "I do not know because there are multiple plausible answers." For a predictive task of predicting Y from X, aleatoric uncertainty takes place whenever the *true distribution* $P(Y \mid X = x)$ is non-deterministic (according to human knowledge), thus has a non-zero entropy. We have aleatoric uncertainty when $Y \mid X = x$ has some entropy. It simply means that a sample x accommodates multiple possible ys.[3]

Examples from the CIFAR-10H [155] dataset are shown in Figure 4.7. For some samples (lower ship and bird), humans are quite uncertain, even without time constraints. We have high aleatoric uncertainty; the true $Y \mid X = x$ (according to human knowledge) has high entropy.[4] The approximation of it by several human inspectors (47-63 per image for the CIFAR-10H dataset [155]) has a high entropy (non-deterministic). They have disagreements.

---

[3]In many cases, we could equivalently say that we have aleatoric uncertainty when the variance of $Y \mid X = x$ is non-zero. However, if we want to be precise, we have to consider that variance is undefined for *nominal/categorical* variables.

[4]For general variables, multimodality is perhaps the most extreme case of aleatoric uncertainty. However, for discrete distributions (corresponding to our categorical variable $Y \mid X = x$ here), multimodality is synonymous with having multiple possibilities, which is synonymous with having a non-zero entropy.
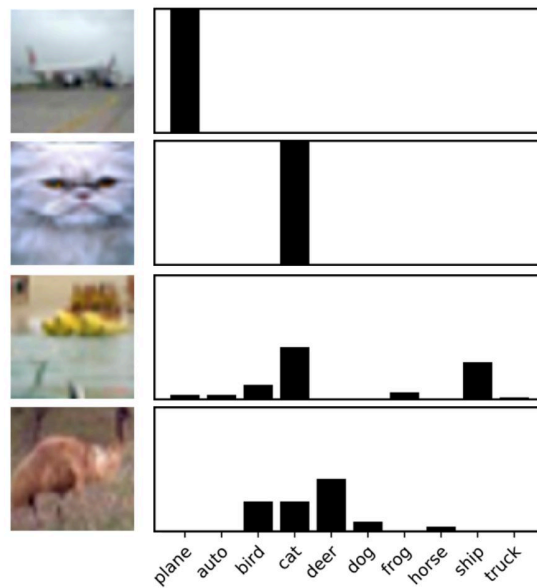
Image source: Battleday et al. 2019 "Improving machine classification using human uncertainty measurements"

Figure 4.7: Example of the absence and presence of aleatoric uncertainty. Examples of images and their human choice proportions are given. For many images (upper plane and cat), the label choices are unambiguous. We have very low aleatoric uncertainty, i.e., the true $Y \mid X = x$ has a very low entropy. The approximation of it by ten human inspectors has no entropy (deterministic); they all agree on the label. The bottom samples accommodate various labels. The single GT label does not always exist. Figure taken from [115].
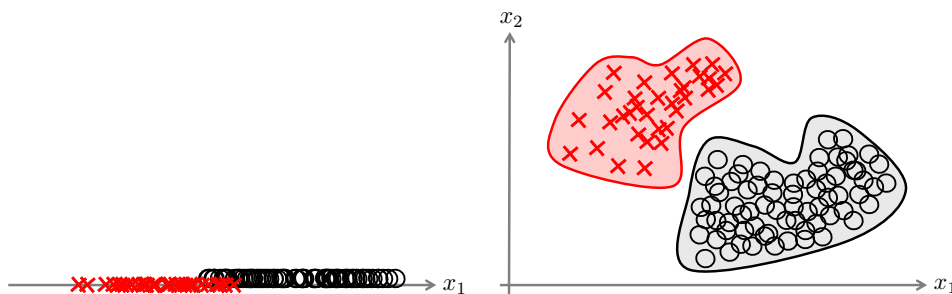
Figure 4.8: Missing features can introduce overlaps in two classes. When we have both features x_1 and x_2, points of the same label are well separated. For example, we know all the pixels in the N-digit MNIST example. We have no aleatoric uncertainty. As soon as we remove some features (e.g., a part of the image), we have overlaps between the classes. Figure taken from [84].

Many faces of aleatoric uncertainty

First, we consider *ambiguity in the observation*. This can arise, e.g., when features are missing (lack of information). An illustration of how missing features can introduce overlaps in two classes is shown in Figure 4.8. This is a typical source of aleatoric uncertainty. We can also take N-digit MNIST samples [150] and consider intentionally corrupted versions of them, shown in Figure 4.9. The input goes through corruption/occlusion that removes some features. Then,
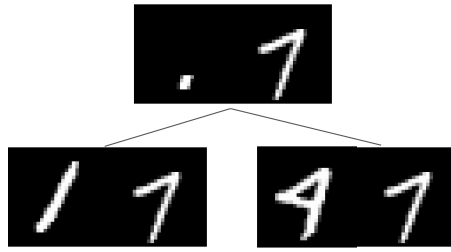
Figure 4.9: Sample from the N-digit MNIST dataset. There are multiple possibilities for the original image. Figure adapted from from [150].

multiple labels might make sense (e.g., $41, 11$). For larger corruptions, we might have

$$P(Y = 41 \mid X = x) = 0.5 = P(Y = 11 \mid X = x).$$

Many people also have poor handwriting, and it is generally difficult to tell a 1 apart from a 7. No artificial perturbations are required in these cases, as the observation already had an inherent ambiguity.

We can also refer back to the CIFAR images from Figure 4.7. When the photo of the ship was taken, it went through extra corruption (resolution reduction) to obtain thumbnails. This removes information and introduces aleatoric uncertainty. If objects are seen in the real world (all features are present), then there is probably no ambiguity.[5]

Out-of-focus images are also examples of ambiguity in the observation. Here, we have measurement noise. This also introduces missing features (information). We cannot tell how many people are in the image if it is severely corrupted.

Let us now consider *ambiguity in the question*. In general, the task may be formulated so that multiple answers are naturally plausible. In the ImageNet-1K dataset, there are several such examples. Consider an image of a desk with many objects on it, illustrated in Figure 4.10. The ImageNet-1K label is 'desk', but other ImageNet-1K categories also make sense: 'screen', 'monitor', or 'coffee mug'. It is quite likely, in general, that multiple classes are present on a single image. In such cases, $P(Y \mid X = x)$ is multimodal. This dataset is not a "solvable" problem, as all labels mentioned are plausible, and neither could not be deemed wrong. Annotators, in this case, will arbitrarily choose one category among them. They are only allowed to provide a single label per image. Referring back to the question of whether neural networks will ever become perfect predictors, it is now clear why the answer is negative. Inherent aleatoric uncertainty is *irreducible*, and correct quantification of uncertainty is, therefore, always needed.

Another example of inherent ambiguity in the question/task is image synthesis. Consider DALL-E image synthesis for the caption "crayon drawing of several cute colorful monsters with ice cream cone bodies on dark blue paper" illustrated in Figure 4.11. Here, $P(\text{image} \mid \text{caption})$ is highly multimodal – we expect multiple good answers. DALL-E generates multiple plausible outputs for the caption, and all of them make sense. Thus, we have aleatoric uncertainty – we do not have a single good answer. (Many images fit the caption, as decided by humans.) In a real dataset, we will not see the same caption

---

[5]This is not true for the handwriting case, where even if we see the handwritten digits in real life, we might be unable to tell a 1 apart from a 7.

twice. We do not exactly have this multitude of possible images given the same caption in the dataset, but it can be an indicator if we see a very close caption corresponding to a completely different GT image. This "approximate multimodality" of our outputs is also counted as aleatoric uncertainty.

Figure 4.10: ImageNet-1K sample with label 'desk'. Aleatoric uncertainty arises naturally because many objects corresponding to different ImageNet-1K labels are present in the image. There is no *single* good answer to this task, therefore, networks should also not be overconfident in one particular prediction. Figure taken from [176].

Figure 4.11: Four samples from DALL-E for the prompt "crayon drawing of several cute colorful monsters with ice cream cone bodies on dark blue paper". Each of the synthesized images is a plausible image given the prompt, leading to the presence of aleatoric uncertainty in $Y \mid X = x$ where Y is the image and x is the exact prompt.

In summary, when we have ambiguities and multiple plausible answers for a task, whatever the source is, we call it aleatoric uncertainty.

Reducing aleatoric uncertainty

Unfortunately, we cannot reduce aleatoric uncertainty by observing more data.[6] When $Y \mid X = x$ has a non-zero entropy, an infinite amount of data will present data samples with *mixed supervision*. For the same x, different supervision signals y will be given. Of course,

---

[6]We will soon see the key difference between epistemic and aleatoric uncertainty: epistemic uncertainty *can* be reduced to 0 with an infinite amount of data, sampled from the right distribution $P(X)$ (considering underexplored regions, too).

for finite datasets like ImageNet-1K, we do not see the same image with different labels but very similar images with different labels. By seeing ambiguities multiple times, we do not reduce them. The model learns to see similarities between images and gets confused if it sees similar images but with very different labels.

To address aleatoric uncertainty, one must...

1. ...formulate a model architecture that accommodates multiple plausible outputs. That is normal for classifiers but not for usual regressors. They usually predict a single number/vector, not a set of plausible answers.

2. ...adopt a learning strategy that lets the model learn multiple plausible outputs rather than sticking to one. This is true for the CE loss for classification. However, it is not true for the L_2 loss for regression! It learns the mean of the labels.

Even though aleatoric uncertainty does not depend on the model, the only possible way to approximate it for a general test input is to use a data-driven model. Then, the focus becomes to formulate models that give reliable aleatoric uncertainty predictions. If we know the generative process (i.e., the true distribution $P(Y \mid X = x)$) or have multiple samples from it, then we can compare aleatoric uncertainty predictions against the true "spread" of $P(Y \mid X = x)$ or the empirical spread, e.g., by comparing against its variance or entropy. Proxy tasks can also be used for benchmarking aleatoric uncertainty predictions. For example, even though aleatoric uncertainty differs from predictive uncertainty, one might want to evaluate the aleatoric uncertainty predictions on predictive uncertainty benchmarks. One reason is practicality. If we do not have access to $P(Y \mid X = x)$, benchmarking against predictive uncertainty is better than not benchmarking at all. Another reason is correlation. Predictive uncertainty necessarily monotonically increases by increasing aleatoric uncertainty. If we assume that epistemic uncertainty (discussed in 4.2.3) does not vary too much on the test samples, we can use the true predictive uncertainty values as ground truth for *ranking* the test samples, and we can measure how well the ranking based on aleatoric uncertainty estimates agrees with it. This is a strong assumption, and such an evaluation is usually used as a heuristic.

4.2.3 Epistemic Uncertainty

Epistemic uncertainty is uncertainty from lack of experience: "I do not know because I have not experienced it." Let us first consider an example of epistemic uncertainty in a binary classification setting to motivate the formalism that follows.

Example of Epistemic Uncertainty: Training Data for Binary Classification

We consider a toy example that showcases the presence of epistemic uncertainty, shown in Figure 4.12. There are several possible classifiers compatible with the data we have observed. While they agree on the data we have observed, we are epistemically uncertain about how to classify points where the models disagree. We wish to sample data from underexplored regions[7] to increase our certainty in the choice of the model.

---

[7]We still want to stay on the data manifold – sampling from underexplored regions that are very implausible is not useful.
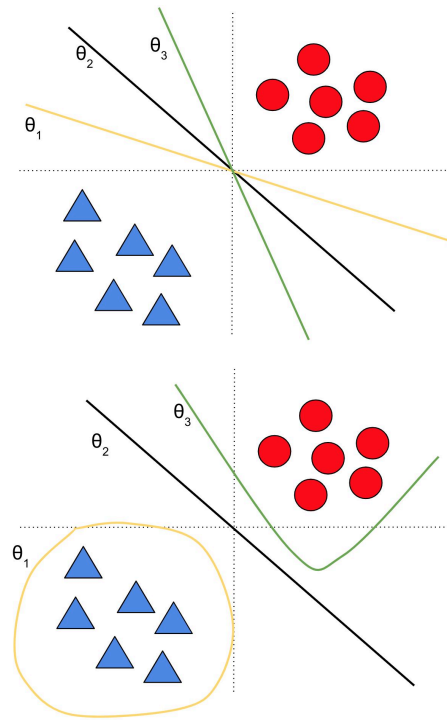
Figure 4.12: Example of the presence of epistemic uncertainty arising from underexplored data regions. The dataset accommodates many models. Models can be from the same hypothesis class (e.g., linear classifiers in the top subfigure or belong to different hypothesis classes (bottom subfigure). To increase our certainty in the "correct" model from the model (= hypothesis) space, we wish to obtain more data from the underexplored regions. Figure taken from [115].

Formal Treatment of Epistemic Uncertainty

Let us consider a more formal definition of epistemic uncertainty than the intuitive description given at the beginning of Section 4.2.3.

> **Definition 4.3: Epistemic Uncertainty**
>
> Epistemic uncertainty is a reducible source of uncertainty that arises due to a lack of knowledge or information. This type of uncertainty can be reduced by collecting more data or improving the model class, as it is a result of the limitations of the current knowledge or understanding of the process being modeled. Examples of sources of epistemic uncertainty include model parameter uncertainty or model *structure* uncertainty.

During learning, we "reduce the possible list of models" to ones that agree with the data (Figure 4.23). One popular way of encoding a "list of plausible models" is via the uncertainty over network parameters in *Bayesian machine learning*:

$$\underset{\substack{\text{No experience}\\\text{Prior over parameters}}}{} \xrightarrow{\substack{\text{Observations}\\\text{Likelihood of data}}} \underset{\substack{\text{Prediction based on experience}\\\text{Posterior over parameters}}}{}$$

We start from our prior knowledge. The prior that encodes our initial beliefs about plausible

models is usually broad and has many possibilities for θ. We have high epistemic uncertainty in regions with no observations, so in the beginning, we have high epistemic uncertainty in general. Then, we accumulate observations. By doing so, we reduce epistemic uncertainty. The likelihood of the data $\mathcal{D}$ is a stack of likelihoods of each data point X_i (IID assumption). By merging our prior knowledge with the observations, we obtain our *posterior beliefs*. Finally, we can make our prediction based on our posterior beliefs using the posterior predictive distribution.

$$P(\theta \mid \mathcal{D}) \propto P(\theta)P(\mathcal{D} \mid \theta) \overset{\text{IID}}{=} P(\theta) \prod_{i=1}^{n} P(X_i \mid \theta)$$

Typically, the entropy for θ decreases with multiple observations.

Model Misspecification and Effective Function Space

The uncertainty arising from the restriction of the model class we are learning over (e.g., all linear models or all GPs), i.e., the uncertainty about *choosing the right model family*, is a part of epistemic uncertainty.[8]

> **Definition 4.4: Model Misspecification**
>
> Model misspecification in ML happens when the inductive biases and prior assumptions injected into the model disagree with the (usually stochastic) process that generated the data.

We leave model misspecification out in the remainder of the book, always assuming that the model class includes the true $P(Y \mid X = x)$ so that the epistemic uncertainty can be reduced to 0.[9] We quickly formalize this below.

> **Definition 4.5: Function Space**
>
> The function space corresponding to a neural network architecture is the set of all functions we can represent using different parameterizations of the architecture:
>
> $$\mathcal{H} = \{f_\theta \colon \mathcal{X} \to \mathcal{Y} | \theta \in \Theta\}$$
>
> where
> - θ is a particular parameterization,
> - $\mathcal{X}$ is the input space and $\mathcal{Y}$ is the output space,
> - and Θ is the space of all possible parameterizations. For example, for a linear regressor with input $x \in \mathbb{R}^n$, $\Theta = \mathbb{R}^n$.
>
> The above definition does *not* consider the training algorithm, the regularizers, or the

---

[8]It could also be treated as a separate source of uncertainty when considering a different definition of epistemic uncertainty.

[9]Most existing uncertainty quantification methods also do not model misspecification as an additional source of uncertainty.

optimizer.

Definition 4.6: Effective Function Space

The effective function space of a neural network is a subspace of the function space that the network can represent. It is a set of functions that the network can actually learn or achieve, given the training procedure, optimization algorithm, and other hyperparameters.

The effective function space of a neural network is influenced by the dataset. For example, a dataset with high noise may require more regularization or early stopping to prevent overfitting, which may limit the effective function space of the network. Conversely, a well-structured and informative dataset may allow the network to explore a wider effective function space as we vary the dataset.

The effective function space of a neural network is also influenced by the optimization algorithm and the training procedure. Different optimization algorithms, such as stochastic gradient descent or Adam, may converge to different local minima (or saddle points), which may affect the set of optimal parameters that the network can achieve. Similarly, the training procedure, such as the choice of learning rate, batch size, or data augmentation, may affect the set of optimal parameters that the network can reach.

When leaving model misspecification out, we can give an alternative definition for epistemic uncertainty: Epistemic uncertainty arises when multiple models out of our *effective* function space can fit the training data well. So, epistemic uncertainty is uncertainty in the *set of plausible models* (but not a property of each individual model). But let's return to the gist of it:

Epistemic uncertainty is reducible.

Considering the appropriate distribution, epistemic (= model) uncertainty *vanishes* (reduces to 0) in the limit of infinite data (= observations).[10] One can thus completely rule out specific models in the limit, and in fact, we can uniquely determine which model is the right one, i.e., which one "generated the data".[11]

Definition 4.7: Data Manifold

Informally speaking, the data manifold is a region of the input space where elements look more natural and realistic.

As a more formal definition, data manifold refers to the underlying geometric structure of the (usually high-dimensional) data that is being modeled. It describes the intrinsic, underlying structure of the data in a lower-dimensional space that captures the essential

---

[10]This depends on what we consider a "model". If we consider the models as the parameters, then this statement is subject to model identifiability. For DNNs, because of weight space symmetries and other factors, many models can correspond to the same function. If we equate models to the functions, then this statement always holds.

[11]We emphasize that this only holds under the simplifying assumptions we (and many other authors) make in this book; namely that the generative model is contained in the effective function space.

> features and relationships between the data points.
>
> The data manifold is typically assumed to be smooth and continuous, and it is usually modeled as a lower-dimensional submanifold embedded in the high-dimensional feature space. The dimensionality of the data manifold is determined by the number of intrinsic degrees of freedom in the data, which is almost always lower than the dimensionality of the original feature space, especially in the case of sensory data.

If the data distribution we sample from does not cover specific areas of interest in the input space (the *data manifold*), then we will still have uncertainty there in limit. It is, therefore, important to sample from *underexplored regions* $P(X)$ of the data manifold that are still realistic but underrepresented in the original training data to achieve this (OOD samples). However, we do not care about images that are purely Gaussian noise or that are away from the data manifold. As soon as we collect and label many OOD samples, we can reduce epistemic uncertainty as much as we like.[12]

Example: Active learning reduces epistemic uncertainty efficiently by acquiring supervision on underexplored samples. We can use epistemic uncertainty to sample from regions where the model needs the most samples. For such a scheme, the model must provide us with well-calibrated epistemic uncertainty estimates.

Example Sources of Epistemic Uncertainty in Practice

Let us first consider two possible sources of epistemic uncertainty that often arise in practice.

Distribution shifts. For example, a self-driving car was mostly trained on daylight videos, but it is deployed in a night scenario. On OOD samples, we (usually) have high epistemic uncertainty.

Novel concepts. For example, new objects, words, or classes (open set recognition). These naturally have high epistemic uncertainty (but not always – this highly depends on the employed inductive biases).

For epistemic uncertainty, many definitions exist (e.g., refer to [175, 200, 112]), and it is not exactly clear what the best way is to properly benchmark such estimates. One possibility is to employ proxy tasks that should be reasonably well correlated with epistemic uncertainty. Another possibility is to consider a binary OOD/not OOD prediction task. This is only a proxy task for epistemic uncertainty because the true "OOD-ness" of a sample is independent of any model. However, we still expect epistemic uncertainty to be higher on OOD samples, so the use of such benchmarks is justified to some extent. This is further discussed in 4.3.1.

4.2.4   Epistemic vs. Aleatoric Uncertainty

Aleatoric uncertainty is *data uncertainty*. It means there is a *multiplicity of possible answers*. When class-conditional distributions overlap, $P(Y \mid X = x)$ has a considerable entropy.[13]

---

[12]New types of realistic OOD data (e.g., counterfactual data) did not matter so much before, so they were not collected. This is precisely the reason they *stayed* OOD. With the rising popularity of the field of ML robustness, these samples also matter a lot (refer back to OOD generalization), so we want to perform well on these samples, too.

[13]For example, if we have two class-conditional Gaussians, we necessarily have variance/uncertainty in the largely overlapping region, but it reduces considerably outside of this region.

Aleatoric uncertainty is inherent to the data distribution.

Epistemic uncertainty is *model uncertainty*. It means there is a *multiplicity of possible models*. It arises from underexplored data regions. Epistemic uncertainty is inherent to the dataset that allows multiple possible hypotheses.

Treating epistemic and aleatoric sources of uncertainty separately is not only done for philosophical reasons. If we only obtained new samples based on regions with high *predictive* uncertainty, it could very well happen that the epistemic uncertainty was actually *low* in that region but a high *irreducible* value of aleatoric uncertainty caused the high predictive uncertainty. For the sake of intuition, we might consider predictive uncertainty as simply the sum of epistemic and aleatoric uncertainty. Later we will see that the decompositions are not this straightforward and require many assumptions. However, we can still say in general that both aleatoric and epistemic uncertainty influence predictive uncertainty.

In essence, both types of uncertainty arise from the data. While epistemic uncertainty depends on the set of plausible models, aleatoric uncertainty does not, as it only depends on the entropy/variance of the true $Y \mid X = x$ variable. This also agrees with the statement that epistemic uncertainty is reducible, while aleatoric is inherent to the data generating process and, therefore, is irreducible. However, they both *influence* predictive uncertainty.

In general, it is hard to disentangle general predictive uncertainty into aleatoric and epistemic sources and is an open research topic. [175, 200, 112]

Additional Information: Epistemic vs. Aleatoric Uncertainty in Computer Vision

We consider a method that models both epistemic and aleatoric uncertainty in computer vision, introduced in the paper "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" [96]. This is illustrated in Figure 4.13. The task is semantic segmentation, which is pixel-wise classification. This method is capable of measuring epistemic and aleatoric uncertainty at the same time.

Aleatoric uncertainty arises at boundaries between classes (e.g., pavement/road). People annotate pixel-wise, and mistakes usually take place around boundaries. Mixed supervision around boundaries leads to high aleatoric uncertainty in these regions.

Epistemic uncertainty arises at parts of the image the model has not seen before. It seems that the model has not seen similar pavements before.

4.3 Connection of Uncertainty Estimates to Earlier Chapters

The subfields of Trustworthy Machine Learning are very interconnected. Here, we briefly discuss some of the connections to OOD generalization and explainability.

| (a) Input Image | (b) Ground Truth | (c) Semantic Segmentation | (d) Aleatoric Uncertainty | (e) Epistemic Uncertainty |
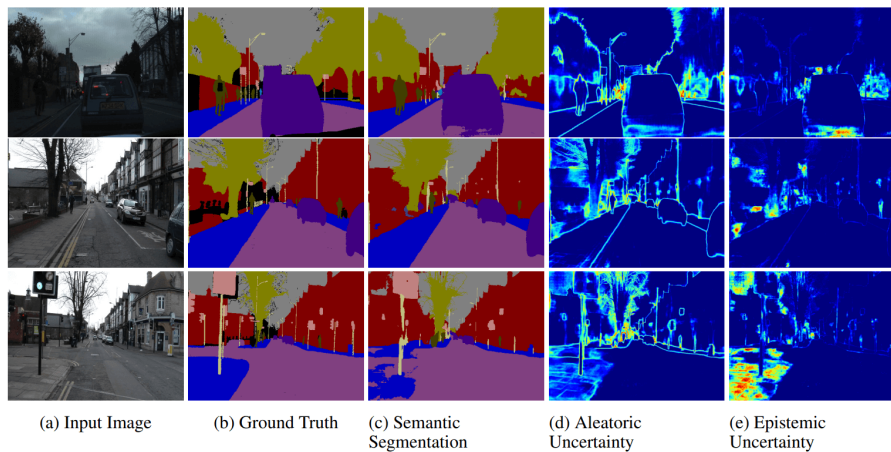
Figure 4.13: Example application of epistemic and aleatoric uncertainty estimation in computer vision. These two sources of uncertainty are fundamentally different, which is further highlighted by the uncertainty maps in the Figure. Figure taken from [96].

4.3.1 Connection of Epistemic Uncertainty to OOD Generalization

Epistemic uncertainty and OOD generalization have many connections, though they should not be treated interchangeably, as discussed previously. Still, epistemic uncertainty should be high for OOD samples. If we have access to M models in the form of an ensemble, then the epistemic uncertainty for a sample x is closely linked to the diversity of predictions $f_1(x), \ldots, f_M(x)$ by the set of trained models. Let us assume that we have a diagonal dataset (Section 2.7.1) and multiple plausible models that are fit to this dataset. As the models are all trained on the training samples (ID data), they all perform well on the training samples (given sufficient expressivity). However, as the models still differ, they will generally not agree on off-diagonal samples. This is emphasized even more if the models are *regularized* to be diverse. Therefore, the off-diagonal samples will have high output variance (high epistemic uncertainty), and the training samples will have very low output variance (low epistemic uncertainty). We can, therefore, measure epistemic uncertainty by training multiple models and seeing how much they agree on a particular sample. This is the essence of Bayesian ML: training multiple models simultaneously more smartly and efficiently, and checking their divergence on certain test samples.

4.3.2 Connection of General Uncertainty Estimation to Explainability

When a model returns its predictive confidence or other uncertainties and is well-calibrated, it is a great way for the user to learn about the model and the output. Such uncertainty estimates are great explanation tools. Some interesting questions that relate explainability to uncertainty estimation are listed below.

- How uncertain was the model?
- Due to which factor was the model uncertain? (If there are multiple factors, see above.)
- What additional training data will make the model more confident? (What regions suffer from high epistemic uncertainty?)

4.3.3 Trustworthiness and Confidence Estimates

A critical component of the trustworthiness of an ML system is the "truthfulness" of the confidence estimates $c(x)$. The most popular demand for *predictive uncertainty* estimates is that $c(x)$ must quantify the actual probability of the model to get the prediction right (known as true predictive uncertainty). A model needs to address two tasks now: (1): Predicting the GT label y, and (2): Predicting the correctness of the prediction L. Then, we want to obtain $c(x) = P(L = 1)$.

4.4 Formats of Uncertainty

Let us first consider different approaches people use to represent/estimate uncertainty. What is the appropriate data format for uncertainty? In the following sections, we will refer to confidence and uncertainty "interchangeably", with "confidence $= 1 -$ uncertainty".

The simplest form: a scalar.

The model f on input x produces an output $f(x)$ *and* a scalar confidence score $c(x) \in [0, 1]$, where

$$c(x) = \text{probability that } f(x) \text{ is correct.}$$

This is a type of *predictive uncertainty* which subsumes aleatoric and epistemic uncertainty (and also the model not being expressive enough).

Note: Whenever we have a scalar in $[0, 1]$, we can treat it as a probability.
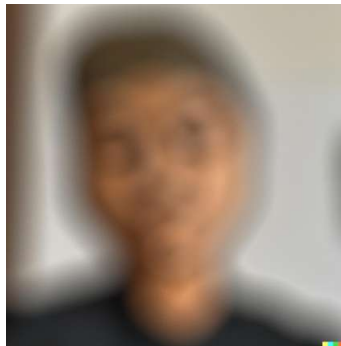
Figure 4.14: A blurry image of a person, generated by DALL-E. The blurriness corresponds to high aleatoric uncertainty.

A vector.

The model can also report $c(x) \in \mathbb{R}^d$, an array of scalars, as a representation of uncertainty. The question we ask is "Which attributes/features/concepts does the model lack confidence in?" We attach a confidence value to each attribute (evidence) of the sample.

Consider a *person identification task*. Let the prediction $f(x) :=$ person name. A possible input x is shown in Figure 4.14. We might obtain the following confidence values over various

evidence:

$$c_{\text{hair color}} = 0.99 \qquad \text{(we kind of see it)}$$
$$c_{\text{eye color}} = 0.39 \qquad \text{(we cannot see it well)}$$
$$c_{\text{ear shape}} = 0.1 \qquad \text{(we are not sure at all)}.$$

The value c can model predictive uncertainty like here (how sure the model is in the correctness of its prediction, broken down into confidences along various evidence), but analogously also aleatoric uncertainty (how much variance does the true $Y \mid X = x$) have along various evidence), or epistemic uncertainty (how much uncertainty there is arising from the lack of observations in the sample along various evidence). For these, different evaluations exist.

A matrix and a vector.

This section is inspired by the work "Probabilistic Embeddings for Cross-Modal Retrieval" [39]. Uncertainty cannot only arise in the outputs of discriminative models that aim to model $Y \mid X = x$. If we want to embed our data into a lower-dimensional space using probabilistic methods, modeling uncertainty has several advantages. We discuss probabilistic embeddings in Section 4.14.1; here, we only consider the representation of uncertainty.

One can have $c(x) = [\mu_\theta(x), \Sigma_\theta(x)]$ interpreted as parameters of a distribution/density. The prediction of the network is a distribution, not a single point. We obtain the posterior over the embedding (probabilistic embeddings), which represents aleatoric uncertainty:

$$P(z \mid x) = \mathcal{N}\left(\mu_\theta(x), \Sigma_\theta(x)\right)$$

with

$$\Sigma_\theta(x) \in \mathbb{R}^{D \times D}.$$

The network outputs a Gaussian for each x, just like a Gaussian Process (GP) would. $\Sigma_\theta(x)$ is a representation of the aleatoric uncertainty in the embedding (covariance of $\mathcal{N}$). This is a more complicated way of uncertainty representation.

A "disentangled" representation.

We consider the work "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" [96] to highlight the possibility to separately obtain aleatoric uncertainty estimations $c_{\text{al}}(x)$ and epistemic uncertainty estimations $c_{\text{ep}}(x)$.[14] Then, we can give our *approximate* predictive uncertainty as $c(x) = c_{\text{al}}(x) + c_{\text{ep}}(x)$.

Consider a regression task, and in particular, the problem of monocular depth estimation, where the network has to output per-pixel depth estimates from a single image. Suppose that we have a distribution $Q(W)$ over the weights W of the model by using dropout (discussed in detail in Section 4.11.4), and each model outputs a mean prediction and a variance term

---

[14]These only approximate the true aleatoric and epistemic uncertainties. Their faithfulness is subject to evaluation.

that measures aleatoric uncertainty. In the referenced paper, the authors calculate these as

$$c_{\mathrm{al}}(x) = \frac{1}{T} \sum_{t=1}^{T} \hat{\sigma}_t^2 \approx \mathbb{E}_q \left[\hat{\sigma}_t^2 \right]$$

$$c_{\mathrm{ep}}(x) = \frac{1}{T} \sum_{t=1}^{T} \hat{y}_t^2 - \left(\frac{1}{T} \sum_{t=1}^{T} \hat{y}_t \right)^2 \approx \mathrm{Var}_q \left[\hat{y} \right]$$

where

$$\left\{ \hat{W}_t \right\}_{t=1}^{T} \sim Q(W), \qquad \left[\hat{y}_t, \hat{\sigma}_t^2 \right] = f^{\hat{W}_t}(x).$$

$c_{\mathrm{al}}(x)$ is the average learned spread (variance) of $Y \mid X = x$) by the ensemble members and c_{ep} is the variance among the ensemble predictions. Here, $\hat{y}_t$ is a single output scalar, corresponding to the mean prediction of model t for a particular input pixel. These uncertainties are calculated for all pixel-wise depth predictions $\hat{y}_t$ of the different networks $\left\{ f^{\hat{W}_t} \middle| t \in \{1, \ldots, T\} \right\}$. Thus, when performing monocular depth estimation, we have as many aleatoric and epistemic uncertainty scalars as there are input pixels.

The Bayesian training for a single input image x is then performed by minimizing the following loss function. This is learned loss attenuation (attenuating the L_2 loss with the learned weight of error σ^2).

$$\mathcal{L}_{\mathrm{BNN}}(\theta) = \frac{1}{D} \sum_{i=1}^{D} \left[\frac{1}{2\hat{\sigma}_i^2} (y_i - \hat{y}_i)^2 + \frac{1}{2} \log \hat{\sigma}_i^2 \right],$$

where

$$\hat{W} \sim Q(W), \qquad \left[\hat{y}, \hat{\sigma}^2 \right] = f^{\hat{W}}(x).$$

The likelihood is Gaussian and heteroscedastic (w.r.t. pixels and samples). $\hat{y}$ is the predicted mean, and $\hat{\sigma}^2$ is the predicted variance (aleatoric uncertainty), both vectors with as many dimensions as there are pixels. q is the approximate posterior over the weights modeled by dropout, which corresponds to epistemic uncertainty. Not only does the formulation allow for modeling epistemic and aleatoric uncertainty, but it also improves accuracy.

4.5 Proper Scoring Rules

We discuss a useful and general framework for training and measuring uncertainty estimates: the framework of *proper scoring rules*. Considering the simplest case of scalar uncertainty estimates, we generally want to learn a value $c(x)$ for a particular sample x that corresponds to the true probability (be it predictive, epistemic, or aleatoric uncertainty). Luckily, there is a class of scores/losses that *ensures this automatically*.

In subsections 4.5.1, 4.5.2, and 4.5.3, we do not make connections to ML concepts, such as the correctness of prediction $L = 1$. We will simply aim to match a predicted probability q of a binary event $Y = 1$ to the true probability p of it. Later, in subsection 4.5.4, we will see that this is indeed very useful for matching probabilities corresponding to different sources of uncertainty in neural networks.

4.5.1   Motivation: Binary Forecasting Task

Consider a simple weather forecasting task. We let subjects bet on the probability of rain tomorrow, which is a binary random variable $Y : \Omega \to \{0, 1\}$ according to the distribution $P(Y)$. The prediction is the scalar $q \in [0, 1]$. We want to encourage the prediction of the correct probability among people. To this end, we give $S(q, Y)$ USD to subjects. S is a function of the reported probability q and the true outcome Y. $Y = 1$ if it actually ends up raining and $Y = 0$ otherwise. Let us assume that the subjects are rational, i.e., they maximize the expected money they get. We want to give the maximum amount of money to people who predict the actual probability of rain. How should we design S?

The expected reward for the subject is

$$\mathbb{E}_{P(Y)} S(q, Y) = S(q, 0) P(Y = 0) + S(q, 1) P(Y = 1),$$

as Y is a binary random variable. Depending on the actual outcome, we get a different amount of money. We wish to find a function S such that

$$\max_{q \in [0,1]} \mathbb{E}_{P(Y)} S(q, Y)$$

is attained iff $q = P(Y = 1)$, i.e., the predicted probability truly represents the probability of rain. That is,

$$\mathbb{E}_{P(Y)} S(q, Y) \leq \mathbb{E}_{P(Y)} S(P(Y = 1), Y) \ \forall q \in [0, 1]$$

and the equality implies $q = P(Y = 1)$. Such a function is called a *strictly proper scoring rule*, formally defined below.

Definition 4.8: Proper/Strictly Proper Scoring Rule

Let us consider a function $S \colon \mathcal{Q} \times \mathcal{Y} \to \mathbb{R}$ where $\mathcal{Q}$ is a family of probability distributions over the space $\mathcal{Y}$, called the label space. For a particular distribution $Q(Y) \in \mathcal{Q}$ and a sample y from a GT distribution $P(Y)$, the function S outputs a real number.

Proper Scoring Rule

S is called a proper scoring rule iff

$$\max_{Q \in \mathcal{Q}} \mathbb{E}_{P(Y)} S(Q, Y) = \mathbb{E}_{P(Y)} S(P, Y),$$

i.e., P is one of the maximizers of S in $\mathcal{Q}$.

Strictly Proper Scoring Rule

S is a *strictly* proper scoring rule iff
   - S is a proper scoring rule and
   - $\arg\max_{Q \in \mathcal{Q}} \mathbb{E}_{P(Y)} S(Q, Y) = P$ is the *unique* maximizer of S in $\mathcal{Q}$.

Note: The family of distributions $\mathcal{Q}$ can be a parameterized distribution with parameters $\theta \in \mathbb{R}^n$ that uniquely define the distribution. In this case, the scoring rule can also be defined over the space parameters Θ instead of the space of distributions $\mathcal{Q}$.

According to the note in the definition of proper scoring rules, instead of considering the family of Bernoulli distributions $\mathcal{Q}$ in the above example, we considered its parameter q for working with proper scoring rules.

Luckily, many often-used loss functions fulfill this criterion.[15] We will look at some examples below.

4.5.2 The Log Probability is a Strictly Proper Scoring Rule

Define

$$S(q,y) \overset{(1)}{:=} \begin{cases} \log q & \text{if } Y = 1 \\ \log(1-q) & \text{if } y = 0 \end{cases} \overset{(2)}{=} y\log q + (1-y)\log(1-q).$$

Using this definition, a very confidently wrong prediction gives $-\infty$ "reward". The "reward" is non-positive in this case. We can think of it as "I will take less money from you if you get the prediction right."

Note: The two expressions (1) and (2) above have different domains. The first one has $D_S = ((0,1] \times \{1\}) \cup ([0,1) \times \{0\})$, whereas the second one has $D_S = (0,1) \times \{0,1\}$. If we take the expectation of both expressions w.r.t. Y, both domains become $D_{\mathbb{E}_{P(Y)}S} = (0,1) \times \{0,1\}$.

The expected reward for the subject is $\mathbb{E}_{P(Y)}S(q,Y) = P(Y=0)\log(1-q) + P(Y=1)\log q$.

Claim 4.1. *S defined above is a strictly proper scoring rule.*

Proof. For the score S to be well-defined, we have to restrict its domain to $D_S := (0,1) \times \{0,1\}$. (Otherwise, we could obtain "$0 \cdot -\infty$" parts in the expectation below. The case distinction formulation of the score makes $S(1,1)$ and $S(0,0)$ also well-defined, but the expectation below would *not* be well-defined if we included $q \in \{0,1\}$.)

Let $a := P(Y=1)$. Then $\mathbb{E}_{P(Y)}S(\cdot,Y)\colon (0,1) \to \mathbb{R}$,

$$\mathbb{E}_{P(Y)}S(q,Y) = P(Y=0)S(q,0) + P(Y=1)S(q,1)$$
$$= (1-a) \cdot \log(1-q) + a \cdot \log(q).$$

To show that S defined above is a strictly proper scoring rule, we can leverage the first-order

---

[15] Strictly speaking, the negative loss functions fulfill this criterion, as scores are meant to be *maximized*.

condition for optimality w.r.t. $q \in (0,1)$ when $a \in (0,1)$.

$$\frac{\partial}{\partial q}\mathbb{E}_{P(Y)}S(q,Y) = -\frac{1-a}{1-q} + \frac{a}{q} \overset{!}{=} 0$$

$$\Longleftrightarrow$$

$$\frac{a}{q} = \frac{1-a}{1-q}$$

$$\Longleftrightarrow$$

$$a - aq = q - aq$$

$$\Longleftrightarrow$$

$$a = q$$

$$\Longleftrightarrow$$

$$P(Y=1) = \hat{P}(Y=1).$$

$q = a$ is the only stationary point when $a \in (0,1)$. To verify that it corresponds to the global maximizer of $\mathbb{E}_{P(Y)}S(q,Y)$, we can use the second derivative test:

$$\frac{\partial^2}{\partial q^2}\mathbb{E}_{P(Y)}S(q,Y) = -\underbrace{\frac{\overbrace{1-a}^{>0}}{\underbrace{(1-q)^2}_{>0}}} - \underbrace{\frac{\overbrace{a}^{>0}}{\underbrace{q^2}_{>0}}} < 0,$$

which verifies that $\mathbb{E}_{P(Y)}S(q,Y)$ is strictly concave in q for $a \in (0,1)$ and $q = P(Y=1)$ is thus the unique maximizer.

Strictly speaking, when $a \in \{0,1\}$, there are no stationary points of the above formulation as $q \in (0,1)$, according to the domain of the score. However, in these cases, we can trivially simplify $\mathbb{E}_{P(Y)}S(q,Y)$, which allows us to extend the domain to allow $q = a$ even in these extreme cases:

$$a = 0: \qquad \mathbb{E}_{P(Y)}S(q,Y) = \log(1-q), \qquad \text{unique maximizer is } q = 0,$$
$$a = 1: \qquad \mathbb{E}_{P(Y)}S(q,Y) = \log(q), \qquad \text{unique maximizer is } q = 1.$$

This concludes the proof that S is a strictly proper scoring rule. $\qquad\square$

4.5.3 The Brier Score is a Strictly Proper Scoring Rule

Define $S(q,y) := -(q-y)^2$ where q is our belief in a binary event $Y = 1$, and y is an actual outcome of the event (0 or 1) according to random variable Y. The reward is higher when our belief matches the outcome. But in proper scoring maximization, we want to maximize the *expectation* in random variable Y (and also in X, considering an entire data distribution $P(X)$ and not just a single sample x). The expected reward for the subject is $\mathbb{E}_{P(Y)}S(q,Y) = -P(Y=0)q^2 - P(Y=1)(1-q)^2$.

Claim 4.2. *S defined above is a strictly proper scoring rule.*

Proof. Analogous to the proof of the log probability being a strictly proper scoring rule. $\quad\square$

4.5.4 Role of Proper Scoring Rules

A proper scoring rule encourages a subject to report the true probability p of some binary event $Y = 1$ as q. As such, it also encourages them to report their true beliefs, as this corresponds to their best approximation of the true probability. Intuitively, it does not make sense to lie. Now we turn away from considering general binary events $Y = 1$ and consider a use case of proper scoring maximization for ML. In particular, we can use proper scoring maximization to encourage a model to choose its confidence value $c(x)$ such that it is equal to the probability of getting the prediction for sample x right ($L = 1 \iff Y = \hat{Y}$).

In the case of ML models, predicting the random variable L implicitly conditioned on x is a binary classification task of whether we are going to make a correct prediction. The original problem of predicting $Y \mid X = x$ can be multi-class classification as well.

4.5.5 Binary Cross-Entropy for True Predictive Uncertainty

> **Definition 4.9: Binary Cross-Entropy (BCE) Loss**
>
> Consider a classifier $f \colon \mathcal{X} \to [0,1]$ that, for a particular input $x \in \mathcal{X}$, predicts the probability of x belonging to class 1, i.e., $P(Y = 1 \mid X = x)$. For a GT label y sampled from $P(Y \mid X = x)$, the Binary Cross-Entropy (BCE) loss is defined as
>
> $$\mathcal{L}(f, x, y) = \begin{cases} -\log f(x) & \text{if } y = 1 \\ -\log(1 - f(x)) & \text{otherwise.} \end{cases}$$
>
> This is the most prominent loss for binary classification when training DNNs.

Consider a binary prediction problem of classifying into classes 0 and 1. Let $f(x) \in [0,1]$ be the predicted probability of model f for class 1 on sample x. It follows that $1 - f(x) \in [0,1]$ is the prediction of the model for class 0. We predict class 1 when $f(x) \geq 0.5$. Otherwise, we predict class 0.

We define our confidence measure as $c(x) := \max(f(x), 1 - f(x))$, called the *max-probability* or max-prob confidence estimate between classes 0 and 1. It is easy to see that $c(x) \in [0.5, 1]$. Other confidence estimates also exist, such as entropy-based ones. These also consider probabilities of other classes. (Implicitly, max-prob does, too.)

We wish to make sure that $c(x)$ estimates the probability of the prediction being correct ($L = 1$). As seen in 4.5.2, we can encourage the model to report $c(x) = P(L = 1)$ (the true predictive uncertainty) by letting the model maximize the log probability proper scoring in expectation of L.

Claim 4.3. *The negative of the BCE loss is a proper scoring rule for* $c(x) := \max(f(x), 1 - f(x))$ *to report the true predictive certainty* $P(L = 1)$.

Proof. According to the definition of the log probability proper scoring rule,

$$S(c, L) := \begin{cases} \log c(x) & \text{if } L = 1 \\ \log(1 - c(x)) & \text{if } L = 0. \end{cases}$$

One can observe that

- $f(x) < 0.5, Y = 0 \iff L = 1 \wedge S(c, L) = \log c(x) = \log(1 - f(x))$;
- $f(x) < 0.5, Y = 1 \iff L = 0 \wedge S(c, L) = \log(1 - c(x)) = \log f(x)$;
- $f(x) \geq 0.5, Y = 0 \iff L = 0 \wedge S(c, L) = \log(1 - c(x)) = \log(1 - f(x))$;
- $f(x) \geq 0.5, Y = 1 \iff L = 1 \wedge S(c, L) = \log c(x) = \log f(x)$.

Therefore,

$$S(c, L) = \begin{cases} \log c(x) & \text{if } L = 1 \\ \log(1 - c(x)) & \text{if } L = 0 \end{cases} = \begin{cases} \log f(x) & \text{if } Y = 1 \\ \log(1 - f(x)) & \text{if } Y = 0. \end{cases}$$

Maximizing the expectation of the above encourages the true predictive uncertainty when our confidence measure is $c(x) = \max(f(x), 1 - f(x))$. This is exactly the log-likelihood criterion for binary classification. Maximizing this reward on a training set is equivalent to minimizing the BCE loss (negative log-likelihood). $\qquad\square$

Conclusion: BCE encourages not only the correctness of classification $f(x)$ but also the truthfulness of the max-prob confidence $c(x) = \max(f(x), 1 - f(x))$. BCE is excellent in this regard.

Remarks for binary cross-entropy

When the prediction is correct, $\log c(x)$ reward is given. As $c(x) \geq 0.5$, we can, at worst, obtain $\log 0.5$ reward when our prediction is correct. When the prediction is incorrect, but c is very large, we can obtain an arbitrarily negative reward. We can see the role of aleatoric uncertainty, as Y is random. We can also see the role of epistemic uncertainty, as $P(Y = f(x))$ depends on whether the model has seen such a sample already or not.

Note: Looking at the log probability proper scoring rule, one might mistakenly think that naively setting $c(x) = 1$ is enough to maximize the expected reward on sample x when the model is correct according to one labeling. However, L is a random variable because $L = \mathbf{1}(Y = f(x))$ and Y is a random variable. There is an inherent stochasticity in L whenever $P(Y \mid X = x)$ has a non-zero entropy: We want $c(x)$ to maximize the *expected* reward, not just the reward for one particular observation of L.

Proper Scoring Maximization on Finite Datasets

When performing ERM, we have no expectation over the loss. We have deterministic (x, y) pairs in our training set and minimize BCE on the batches. (Multiples *can* be present in the dataset with different labels. Very similar inputs can also correspond to different labels. But every (x, y) pair we have is deterministic.) In this case, we have no guarantee of recovering the true predictive uncertainty $P(L = 1)$ for all samples. We only have the guarantee of recovering the empirical probabilities $\hat{P}(L = 1)$ based on our dataset. We also have no guarantees of

how faithful our predictive uncertainty scores are on unseen (e.g., OOD) samples, as we can arbitrarily overfit our predictive uncertainty predictions. This is important to keep in mind.

Therefore, the truthfulness of the max-prob confidence estimates is only encouraged w.r.t. the empirical probability of correctness on the training set. When we consider the idealistic case of having infinitely many samples from $P(X)$ (i.e., we optimize the expectation), then we have the guarantee that $c(X)$ will recover $P(L = 1)$ for all samples $X \sim P(X)$.

By optimizing the BCE, our model also becomes better on the training samples (until a certain point, given by how expressive the model is). Therefore, the well-calibratedness – as measured by log probability proper scoring – and the accuracy usually improve hand-in-hand.[16] We saw above that BCE encourages the prediction of the true probability of correctness. We can consider two corner cases here, depending on the expressivity of our model.

1. Consider a shallow model, such as a logistic regression classifier. Further, assume that the dataset's generative model is non-linear; there is model misspecification. Unfortunately, even in the limit of infinite data, training with the BCE loss (and in general with any negative proper scoring rule) *does not ensure* that we get well-calibrated predictive uncertainty estimates. Proper scoring rules only guarantee that they are maximized at the GT distribution in expectation. They do not give any guarantees for calibration when this maximizer cannot be attained in our function class. However, when our estimator is consistent, we are guaranteed to have calibrated predictive uncertainty estimates in the limit of infinite data when using strictly proper scoring rules.

2. Now, let us assume that we have a very expressive model: one that is capable of fitting to the generative model extremely well. When trained with the BCE loss, in the limit of infinite data, the model will give very accurate predictive uncertainty estimates. If we consider a case with low aleatoric uncertainty, these estimates will be very confident in the model being correct – and the model will indeed be correct most of the time.

It is hard to create an expressive model using only this criterion that is well-calibrated but inaccurate, as both are optimized simultaneously.

4.5.6 Multi-Class Cross-Entropy (CE) for True Predictive Uncertainty

> **Definition 4.10: Multi-Class Cross-Entropy (CE) Loss**
>
> Consider a classifier $f \colon \mathcal{X} \to \Delta^K$ that, for a particular input $x \in \mathcal{X}$, predicts an element of the $(K-1)$-dimensional probability simplex, i.e., predicts a vector of probabilities corresponding to each class. For a GT label y sampled from $P(Y \mid X = x)$, the (multi-class) Cross-Entropy (CE) loss is defined as
>
> $$\mathcal{L}(f, x, y) = -\log f_y(x).$$
>
> This is the most prominent loss for multi-class classification when training DNNs.

---

[16]Accuracy is usually highly correlated with the negative loss. However, not all calibration metrics have such a high correlation with accuracy.

In multi-class classification, we usually use CE as our loss function. We will see that it also encourages the correct predictive confidence.

Let $f(x) \in \mathbb{R}^K$ be a vector of probabilities for each class $k \in \{1, \dots, K\}$. That is, $\forall i \in \{1, \dots, K\}$: $f_i(x) \geq 0$ and $\sum_{i=1}^{K} f_i(x) = 1$. We can define our confidence measure as the max-probability among class probabilities: $c(x) := \max_k f_k(x)$. Then, just like before, we could apply the log probability proper scoring rule. This rewards the model for how correct it is on its own most likely prediction. But notice the following, using the shorthand $k_{\max} := \arg\max_k f_k(x)$:

$$
\begin{aligned}
&S(c, L) \\
&= \begin{cases} \log c(x) & \text{if } L = 1 \\ \log(1 - c(x)) & \text{if } L = 0 \end{cases} \\
&= \begin{cases} \log \max_k f_k(x) & \text{if } Y = k_{\max} \\ \log \sum_{k \neq k_{\max}} f_k(x) & \text{if } Y \neq k_{\max} \end{cases} \\
&= \begin{cases} \log f_Y(x) & \text{if } Y = k_{\max} \\ \log \left(f_Y(x) + \sum_{k:k \notin \{Y, k_{\max}\}} f_k(x) \right) & \text{if } Y \neq k_{\max} \end{cases} \\
&\geq \log f_Y(x).
\end{aligned}
$$

The proper scoring rule S for $L = 1$ can be bounded from below with $\log f_Y(x)$, i.e., the log probability the model assigns to the true class. The negative log probability $-\log f_Y(x)$ is the CE loss, one of the most widely used losses for training classifiers. Maximizing the lower bound $\log f_Y(x)$ (minimizing the CE loss) encourages $c(x) = \max_k f_k(x)$ to be the truthful predictive uncertainty (either $\hat{P}(L = 1)$ or $P(L = 1)$, depending on whether we consider the expectation or its Monte Carlo (MC) approximation). While in general, when maximizing a lower bound, we do not have any guarantee that we also maximize the original objective, we can prove just that here: In Section 4.13.3, we will prove that this lower bound is *also* a strictly proper scoring rule for the correctness of prediction (thereby saving the CE loss's reputation). In that chapter, we will also uncover important relationships between proper scoring rules for predictive uncertainty and aleatoric uncertainty.

4.5.7   Strictly Proper Scoring Rules can Behave Differently

We have now discovered two strictly proper scoring rules for the correctness of prediction: the log probability of the model's most likely class and the log probability of the true class. Which one should we use? The important bit is that being strictly proper does not necessarily mean that they are also good training objectives. When training deep neural networks, we are solving a highly non-convex optimization problem. Different objectives might induce noisier and more complex loss surfaces: It could be that one of the scoring rules provides a better regularization of the loss surface (which, perhaps, is smoother). In that sense, it is also meaningful to empirically compare the two scores.

Additional Information: Benchmarking Strictly Proper Scoring Losses

Let us compare training with the objective

$$\mathcal{L}_1(f(x), y) = \begin{cases} -\log \max_k f_k(x) & \text{if } y = \arg\max_k f_k(x) \\ -\log \sum_{k \neq \arg\max} f_k(x) & \text{if } y \neq \arg\max_k f_k(x), \end{cases}$$

to usual CE training using

$$\mathcal{L}_2(f(x), y) = -\log f_y(x).$$

This experiment is conducted in the linked notebook. For a toy dataset like MNIST, a shallow CNN (3 convolutional layers) fits the training data very well with both losses and produces equivalent results across the ECE, log probability, and Brier Score metrics. However, training with $\mathcal{L}_1(f(x), y)$ converges slower, even after tuning hyperparameters to have a fair comparison.

Comparing the losses on a slightly more realistic dataset, CIFAR-10, the model is not expressive enough to get close to interpolating the training dataset. The network trained with CE achieves an accuracy of around 67%. The $\mathcal{L}_1(f(x), y)$ loss variant converges even slower than before, and plateaus much earlier. Even after hyperparameter tuning, it only reaches an accuracy of 54% on average. Even though the solution sets are identical, the loss surface corresponding to $\mathcal{L}_1(f(x), y)$ is considerably noisier. Regarding calibration, the Brier Score and log-probability scores are higher for the NLL-trained network (which is partly expected because it also has a considerably higher accuracy) but the ECE value for the $\mathcal{L}_1(f(x), y)$ loss network is very slightly better. Checking how the uncertainty estimates perform in predicting aleatoric uncertainty would also be a curious research objective.

In conclusion, $\mathcal{L}_1(f(x), y)$ *can* train a model, but generally with worse accuracy and predictive uncertainty estimates (as measured by proper scoring rules). This might come as a surprise, given that minimizing a proper scoring loss directly tries to optimize the metric we evaluate on. However, numerical optimization can be quite unintuitive and is generally unpredictable. Not all strictly proper scoring rules are equally good training objectives.

4.5.8 Multi-Class Brier Score

Some researchers also report the multi-class Brier score:[17]

$$S(f(x), y) = -(1 - f_y(x))^2 - \sum_{k \neq y} f_k(x)^2.$$

Claim 4.4. *The above multi-class Brier score provides a lower bound on the Brier score for the max-prob confidence estimate, $S(c, l) = -(c(x) - l)^2$, where l is a realization of the Bernoulli random*

---

[17]Some people refer to scores even when lower is better. To discuss a unified overview in this book, we refer to scores when we wish to 'maximize', and refer to losses when we wish to 'minimize'. There is a trivial correspondence between scores and losses when taking reciprocals or negatives.

variable L.

Proof.

$$S(c, l) = -(c(x) - l)^2$$

$$= \begin{cases} -(c(x) - 1)^2 & \text{if } l = 1 \\ -c(x)^2 & \text{if } l = 0 \end{cases}$$

$$= \begin{cases} -(\max_k f_k(x) - 1)^2 & \text{if } y = \arg\max_k f_k(x) \\ -(\max_k f_k(x))^2 & \text{if } y \neq \arg\max_k f_k(X) \end{cases}$$

$$= \begin{cases} -(1 - f_y(x))^2 & \text{if } y = \arg\max_k f_k(x) \\ -(\max_k f_k(x))^2 & \text{if } y \neq \arg\max_k f_k(x) \end{cases}$$

$$\geq \begin{cases} -(1 - f_y(x))^2 & \text{if } y = \arg\max_k f_k(x) \\ -\sum_{k \neq y} f_k(x)^2 & \text{if } y \neq \arg\max_k f_k(x) \end{cases}$$

$$\geq - \left[(1 - f_y(x))^2 + \sum_{k \neq y} f_k(x)^2 \right]$$

$$= S(f(x), y).$$

$\square$

Perhaps unsurprisingly, this lower bound is, in fact, also a strictly proper scoring rule for the correctness of prediction. We will show this in Section 4.13.3.

> **Additional Information: Can learning theory be used for uncertainty guarantees?**
>
> We have not yet seen learning theory used for uncertainty prediction. In learning theory, we have many results based on the 0-1 loss and binary classification. In predictive uncertainty, we also have a binary classification problem: Is the prediction correct or not? However, it is not a standalone classification problem. First, we make a prediction, and then based on that, we can make the meta-output of whether the prediction was correct. It would be interesting to have such results, but it is very underexplored at the moment.

4.5.9 Empirical Evaluation of Predictive Uncertainties

Using a Test Set to Measure Generalization

As discussed previously, a good objective does not necessarily imply that the final trained model behaves nicely if we train with that objective. For the training set samples, it trivially does. However, we can still arbitrarily overfit to training set samples (during the optimization, anything can go wrong) and be very confidently wrong on test samples. The model then fails to represent its uncertainty generally. This is already problematic for ERM without uncertainty quantification. So we need some metrics to evaluate the uncertainty estimates on test sets.

Using Proper Scoring Rules to Evaluate Predictive Uncertainties

We need empirical evaluation for predictive uncertainty. For empirical evaluation, we always need a sensible evaluation metric. And what metric could be better than one for which we know it achieves its minimum if and only if the prediction is correct? (Strictly) proper scoring rules to the rescue!

Log probability. As the log probability is a strictly proper scoring rule for the correctness of prediction, we can use the average CE (NLL) over the test samples as the evaluation metric (where lower is better) for multi-class classification:

$$\mathcal{L}_{\text{NLL}} = -\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \log f_{y_i}(x_i).$$

Luckily, many papers report NLL tables besides, say, accuracy or RMSE. This allows judging the correctness of confidence predictions.

In NLP, people use perplexity instead of CE (especially for language models, used in benchmarks), which is very similar to CE:

$$\mathcal{L}_{\text{NLL}} = -\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \log f_{y_i}(x_i)$$

$$\mathcal{L}_{\text{Perplexity}} = 2^{-\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \log_2 f_{y_i}(x_i)}$$

The perplexity is the exponentiated NLL value, using base 2 in both the exponential and the logarithm.[18] It shows the same information but is generally deemed more intuitive because of the following reasons.

1. Perplexity can be interpreted as the weighted average branching factor of a language model [90]. In the context of language models, the branching factor refers to the number of words that can follow a given context (with non-zero probability). The word 'weighted' is used because the language model usually assigns different probabilities to different words that can follow – perplexity takes this into consideration. A lower perplexity means the language model is less "perplexed" or less uncertain, i.e., it is more confident in its predictions. This intuition can be easier to understand compared to the raw log-likelihood.
2. Exponentiating with base 2 "undoes" the $\log_2$ operation, bringing the metric back into the probability space.

Note: One can verify that larger LLMs seem to have lower test perplexities, meaning they *seemingly* give better predictive uncertainty estimates (Figure 4.15. However, the NLL and perplexity metrics mix calibration with accuracy (see above). Therefore, we should only conclude that larger LLMs fit the data distribution better, which is not a surprising outcome.

Multi-class Brier score. As the multi-class Brier score is also a proper scoring rule for the

---

[18]The reader can easily convince themselves that the perplexity is independent of the common base of the exponential and logarithm.
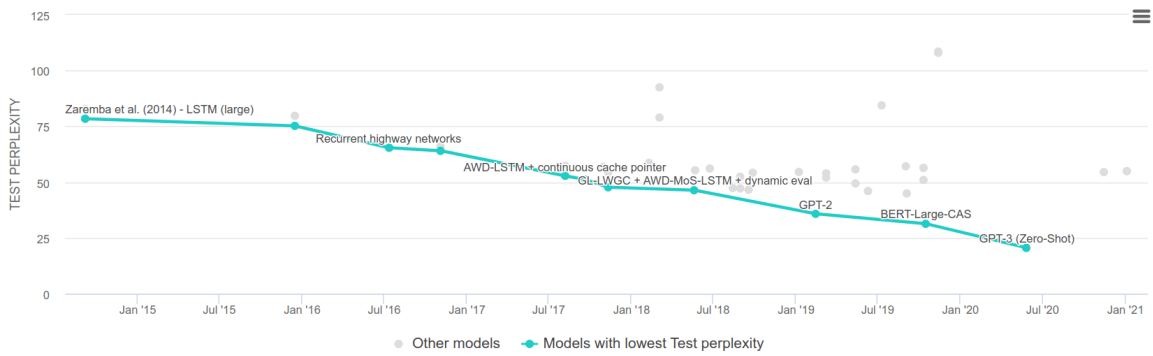
Figure 4.15: Leaderboard of perplexity of Penn Treebank on 04.03.2023 [131]. Test perplexity shows a decreasing trend with increasing model capacity.

correctness of prediction, we can evaluate our predictions using the loss

$$\mathcal{L}_{\text{Brier}} = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \left[(1 - f_{y_i}(x_i))^2 + \sum_{k \neq y_i} f_k(x_i)^2 \right].$$

Remarks for the two previous examples. The lower $\mathcal{L}_{\text{NLL}}$ and $\mathcal{L}_{\text{Brier}}$ are, the better our predictive uncertainty estimates are. However, there are a few important things to keep in mind.

1. We do not know the lowest possible value of these values in expectation over the data generating process. It depends on the aleatoric uncertainty $P(Y \mid X = x)$ on samples $X \sim P(X)$.[19]

2. The NLL can be challenging to interpret. If we take its exponential, then we *roughly* get the average probability assigned to the correct class – not exactly because of the order of sum and exp. For the correctness of prediction, this still does not give rise to an intuitive explanation. Further, it is unbounded from above and bounded from below by the true aleatoric uncertainty, which is generally unknown. The Brier score can be easier to interpret in this regard.

3. In general, proper scoring rules for predictive uncertainty using max-prob *mix good calibration with good accuracy*. Notably, this is not the case for ECE (Section 4.6) that can capture calibration *independently* from accuracy.

4. The pointwise Bayes predictor (the predictor with the minimal pointwise risk), $P(Y \mid X = x)$, is a maximizer of these scoring rules with a max-prob confidence estimate, but it is also a maximizer of proper scoring rules for aleatoric uncertainty. Therefore, epistemic uncertainty is not taken into account – proper scoring only gives statements in expectation over labels, and the Bayes predictor necessarily has an epistemic uncertainty of zero as it only models aleatoric uncertainty.

---

[19] If we do not take an expectation but still have mixed supervision (different labels for the same input x), the lowest possible value is, again, non-zero.

4.6 A New Notion of Calibration

We have seen that proper scoring rules can be used to define a notion of calibration, but their values are often hard to interpret. In this section, we discuss an easily interpretable notion of calibration. However, we will also see that, unlike proper scoring rules, it can be cheated.

4.6.1 Evaluating Calibration

Let us first discuss how we can evaluate calibration [69], quantifying it in an alternative way compared to proper scoring rules. Let the input be $x \in \mathcal{X}$, the output be $y \in \mathcal{Y} = \{1, \ldots, K\}$ (multi-class classification problem) and the model output be

$$h(x) = (\hat{y}, c(x)),$$

which is a pair of the class prediction and the confidence estimate, respectively. $c(x)$ does not have to be a max-prob confidence estimate.

> **Definition 4.11: Perfect Calibration**
>
> A model is *perfectly calibrated* if $P(\hat{Y} = Y \mid C = c) = c \quad \forall c \in [0, 1]$.

Intuitively, for confidence level c, the probability of correct prediction should be c, as the confidence level should faithfully reflect the probability of correctness. This is very similar to what we meant by the correct prediction of predictive uncertainty.

Example for the empirical probability in practice: Predictions for any sample in our dataset with confidence score $c = 0.8$ must only be correct 80% of the time. A rough outline of a procedure that checks for this (refined later) can be given as follows.

1. *Collect all samples in the test dataset with confidence score $c = 0.8$.*
2. Compute accuracy across all samples.
3. Check whether this gives us 80% accuracy.

> **Definition 4.12: Model Calibration**
>
> *Model calibration* is defined as
>
> $$\mathbb{E}_{c \sim C}\left[\left|P(\hat{Y} = Y \mid C = c) - c\right|\right] = \int \left|P(\hat{Y} = Y \mid C = c) - c\right| dC(c).$$

Informally, model calibration quantifies the deviation of our model from perfect calibration. Of course, in practice, we do not have access to the data generating process and, therefore, cannot compute model calibration. If we resort to empirical probabilities, a problem with the rough outline we discussed above is that we never have samples with exactly the same confidence scores, so we cannot calculate the model's accuracy on them this way. An easy fix is to *introduce binning*. The Expected Calibration Error (ECE) metric does exactly that.

Definition 4.13: Expected Calibration Error (ECE)

Expected Calibration Error is a finite approximation of model calibration that uses binning:

$$\mathrm{ECE} = \sum_{m=1}^{M} \frac{|B_m|}{n} \left| \mathrm{acc}(B_m) - \mathrm{conf}(B_m) \right|$$

where

$$\mathrm{acc}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \mathbf{1}\left(\hat{y}_i = y_i \right),$$

$$\mathrm{conf}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} c_i.$$

The ECE measures the deviation of the model's confidence predictions from the corresponding actual accuracies on a test set. It is a weighted average of bin-wise miscalibration. $\mathrm{acc}(B_m)$ is the proportion of correct predictions (the accuracy) in the mth bin, and $\mathrm{conf}(B_m)$ is the average confidence in the mth bin. We take the average of the confidences to ensure we follow the actual confidence values in this range more precisely. Further, we weight by the bin size for the correct approximation of the expectation: $\hat{C}(c) = \frac{|B_m|}{n}$.

Computing the ECE in practice can be done as follows.

1. Train the neural network on the training dataset.
2. Create predictions and confidence estimates using the test data.
3. Group the predictions into M bins (typically $M = 10$) based on the confidences estimates. Define bin B_m to be the set of all predictions $(\hat{y}_i, c_i)$ for which it holds that

$$c_i \in \left(\frac{m-1}{M}, \frac{m}{M} \right].$$

4. Compute the accuracy and confidence of each bin B_m using the above formulas for $\mathrm{acc}(B_m)$ and $\mathrm{conf}(B_m)$.
5. Compute the ECE by taking the mean over the bins weighted by the number of samples in them.

Additional Information: Relationship of the above metrics

What we would ideally want to achieve is that the model returns *truthful predictive uncertainty estimates*, i.e., $c(x) = P(L = 1 \mid x) \forall x$. However, that is impossible to measure. So we measure a necessary (not sufficient!) condition: If the model always returns truthful predictive uncertainty estimates, then it also needs to be *perfectly calibrated* (across all x that have the same $c(x)$.

This condition is quantified by the *model calibration*: The model calibration is zero if and only if the model is perfectly calibrated. To measure this in practice, we need to approximate it by the *ECE*. This is basically a discretized version of the model calibration integral.

> Due to the approximation, we cannot theoretically guarantee that an ECE of 0 implies a model calibration of 0 or vice versa (and, in fact, we show how to game both below). But an ECE close to zero means the model calibration should also be close to zero. This, in return, at least checks one of the boxes a model with truthful predictive uncertainties has to fulfill. It is the best we can do in practice.

While the ECE is a useful metric, for high-risk applications we might be interested in worst-case metrics. The *Maximum Calibration Error* computes such a worst-case discrepancy.

> **Definition 4.14: Maximum Calibration Error**
>
> The *Maximum Calibration Error* is a useful metric for high-risk applications:
>
> $$\mathrm{MCE} = \max_{m \in \{1,...,M\}} |\mathrm{acc}(B_m) - \mathrm{conf}(B_m)|.$$

MCE computes the maximal bin-wise miscalibration (difference between empirical accuracy and average confidence value). This might be a very pessimistic metric if for

$$m' := \argmax_{m \in \{1,...,M\}} |\mathrm{acc}(B_m) - \mathrm{conf}(B_m)|,$$

$\frac{|B_{m'}|}{M}$ is very small, depending on our end goal. For high-risk applications, we could also define the worst-case ECE per class if our concern is per-class performance.

4.6.2   Gaming the ECE Metric

ECE is usually a good *indicator* of whether something is fairly well-calibrated. Its main advantage is that ECE scores are often more interpretable and intuitive than proper scoring rules, as they denote deviations from the perfect calibration in a bounded manner: The ECE is a number between 0 and 1. It tells us how much we are deviating from the $x = y$ line w.r.t. a weighted average. In comparison, NLL scores can be arbitrarily large. When we consider the log probability, the sign flips, which can be confusing. We cannot immediately tell what is good or bad. It is difficult to interpret what the numbers mean, and it heavily depends on the scoring rule of choice.

Although it has many nice properties, *the ECE is not a proper scoring rule*. One can easily achieve ECE $= 0$ (the minimal value) even when the model is not reporting the true predictive uncertainties. This can give us a false sense of calibration and can kill the purpose of the metric. In particular, if we predict a constant c for all samples, where $c = P(\hat{Y} = Y)$ is the global accuracy of the model on the data distribution. Then the conditional probability is only defined for $c = P(\hat{Y} = Y)$, as this is the only value with a positive measure (i.e., we have a Dirac measure at the global accuracy), and for this value, the definition holds by construction. To game the ECE metric, one does not even need access to labeled validation data. All one needs to know is the prior probability of correctness, $P(\hat{Y} = Y)$. The same trick can game the more theoretical notion of model calibration.

Therefore, perfect calibration does *not* imply that $c(x) = P(L = 1)$, i.e., that $c(x)$ is the GT probability of predicting the output correctly for all individual inputs x. Predictive uncertainties can be arbitrarily incorrect per sample ($c(x) \neq P(L = 1)$). This is because the *conditional* probability $P(\hat{Y} = Y \mid C = c)$ *aggregates* all samples with the same value $c(x)$. As long as this group has the correct accuracy on average, it is considered perfect. The intention of ECE and related metrics is still to ensure $c(x) = P(L = 1)$, but they fail to fully encode this requirement. This can be exploited to, e.g., win competitions and benchmarks.

Another important drawback of the ECE metric is that it depends on the binning. Using twenty bins gives us a different score than using ten. There should be an agreed-upon number of bins across papers and methods. This is usually ten but there are several papers using different numbers as well. However, fixing it is probably not a good idea in the long run: There are pros and cons of fixing the number of bins. Eventually, models will be making more and more correct predictions. We should probably make binning more fine-grained near the $90\% - 100\%$ confidence range, as there will probably be a lot more samples there.

4.6.3   Reliability Diagrams

Instead of quantifying calibration in a single number, we can also *visualize* how well-calibrated a model is by leveraging *reliability diagrams* (Figure 4.16).

> **Definition 4.15: Reliability Diagram**
>
> A reliability diagram is a visualization of model calibration that uses binning. It is calculated as follows.
>   1. Bin through different confidence values and take the mean accuracy per bin on the test set: for each bin, calculate $\mathrm{acc}(B_m)$ and $\mathrm{conf}(B_m) - \mathrm{acc}(B_m)$ as defined previously.
>   2. Visualize the discrepancies between the bin-wise accuracies and confidences using a barplot.

Reliability diagrams allow us to judge whether a model is under- or overconfident (or a mixture). While ECE only concerns the distance to the true c, the diagram tells us whether the actual accuracy is higher or lower than the model predicts. If the line is above, then the model is *underconfident*. If it is below, it is *overconfident* (as in Figure 4.16).

Reliability diagrams also allow us to look at the MCE, while ECE can often hide that. But they do not allow inferring the ECE because we do not know the bin sizes (the weights). Seemingly large discrepancies might be weighted with a negligible weight if only a couple of samples are in those bins. If the model on the right had a tiny gap for the last bin, it could have a lower ECE value than the one on the left. Even if the weights are reported as histograms along with the reliability diagrams (the original paper did this, but most follow-ups drop this), the reliability diagram might still give the wrong impression *at first glance*.

The connection between reliability diagrams and the ECE and MCE scores can be seen in Figure 4.17. Note that the plot starts at 0.1 and not at 0. This is not a coincidence: If we use the max-prob class as a prediction, its lowest possible c can only be $1/K$. This becomes even
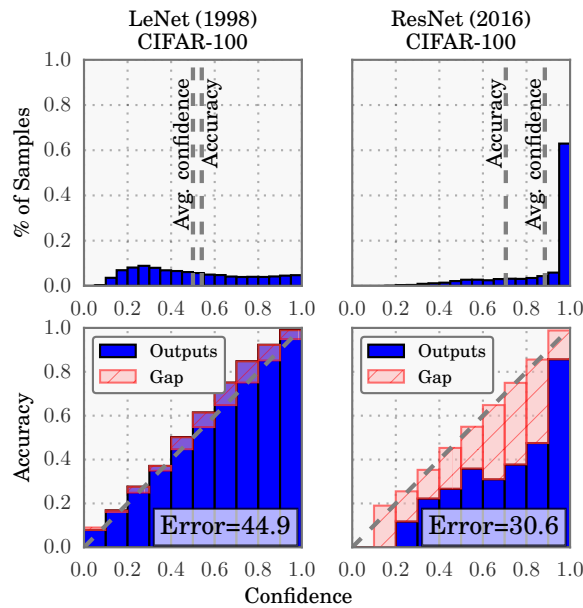
Figure 4.16: Example reliability diagrams (bottom Figures) with confidence histograms (top Figures). Error refers to the Top-1 accuracy of the models, not the calibration error. For each bin, $\mathrm{acc}(B_m)$ in blue and $\mathrm{conf}(B_m) - \mathrm{acc}(B_m)$ in red are plotted as a barplot. The well-calibrated model's (left) gaps are nearly all smaller than the less well-calibrated model's (right) gaps. In the right plot, the bin indices and accuracies for the bins do not match well. The ResNet is much more accurate and also much less well-calibrated. Figure taken from [69].
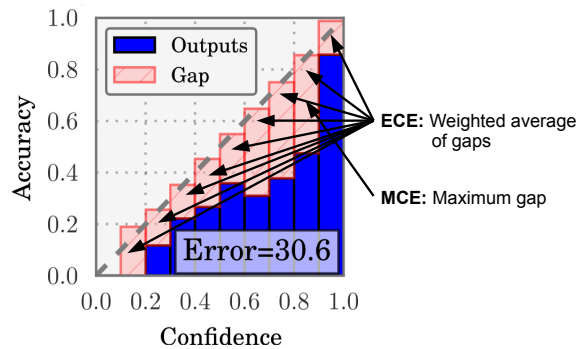
Figure 4.17: Connection between the reliability diagram and the ECE, MCE scores. Accuracy: $P(\hat{Y} = Y \mid C = c)$, confidence: $C = c$. Figure taken from [52].

more visible when we only have 10 or 2 classes.

For binary classification, there is also a second definition of reliability where the y-axis shows the probability of the positive class. Thus, it always starts at 0 and does not include the mind-flip that the confidence may also be the probability of the 0 class. However, it requires a different mind-flip: An underconfident model, in this case, would have an S-shaped diagram. In the definition of [69] above, an underconfident model has a curve that is always above the line. This version of a reliability diagram is common in traditional statistics, where classes are not equal, but the 1 class is more important. So, if one sees a binary reliability diagram, it is better to double-check its axis labels.

4.7 Summary of Evaluation Tools for the Truthfulness of Confidence

Let us provide a collection of evaluation tools for the truthfulness of confidence (predictive uncertainty).

Proper Scoring

As we have seen before, one can use the negative log-likelihood (NLL) loss or the log probability scoring rule on the test dataset to evaluate the truthfulness of predictive uncertainty estimates. Similarly, one can use the Brier score or its multi-class variant on a test dataset. These are all proper scoring rules/losses for the correctness of prediction.[20]

Metrics Based On Model Calibration

One can use the ECE score for an expected deviation from perfect calibration (in a binned fashion). For high-risk applications where we are concerned with the "worst-case bin," one can also employ the MCE score.

It is also possible to visualize calibration by using reliability diagrams. However, it is also important to plot confidence histograms, as reliability diagrams alone can be misleading.

These metrics/visualization tools are all used for predictive uncertainty (correctness of prediction $L = 1$).

4.8 Excourse: How well-calibrated are DNNs?

Let us consider some findings from the literature on DNN calibration.

4.8.1 On Calibration of Modern Neural Networks

We discuss the seminal paper titled "On Calibration of Modern Neural Networks" [69]. In particular, we refer to Figure 4.16. Both LeNet and ResNet are trained with the NLL loss, which is the negative of a lower bound of a proper scoring rule for multi-class predictive uncertainty under max-prob. According to the Figure, LeNet is relatively well-calibrated, and ResNet performs worse than LeNet regarding calibration.

It is important to note that this finding is not a general observation. ResNet-50s usually perform well on calibration benchmarks [55]. Training procedures and best practices since this work have also improved considerably, which might have compounding effects on the results shown in Figure 4.16.

---

[20]The NLL loss and the multi-class Brier score are also strictly proper for aleatoric uncertainty (i.e., the recovery of $P(Y \mid X = x)$), as we will see in Section 4.13.2.

> **Additional Information: The Use of ResNets in Modern DL**
>
> In medium-sized models, ResNets are still among the top performers (see the "What Can We Learn From The Selective Prediction And Uncertainty Estimation Performance Of 523 Imagenet Classifiers" paper [55]. They are often used in practice as "the smallest possible model that still allows experimenting with DL."
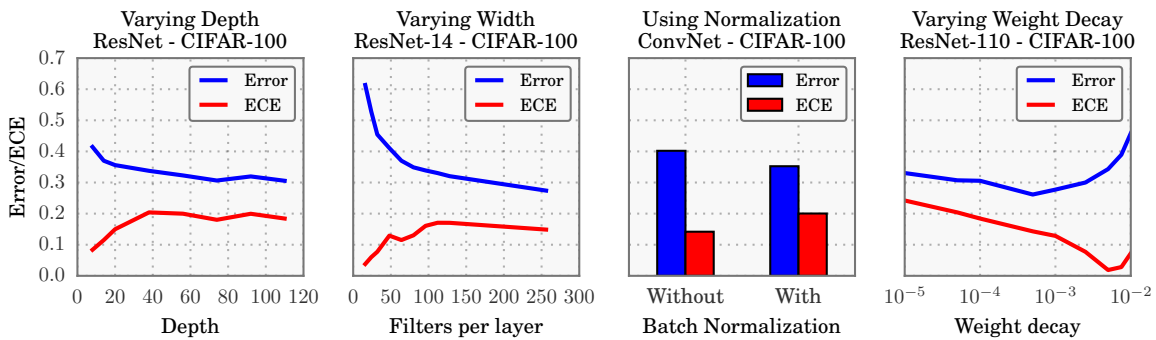
Figure 4.18: Influence of depth, filters per layer, batch normalization, and weight decay on the error and calibration of different ConvNet architectures. Figure taken from [69].

Why is this the case? Let us consider Figure 4.18. Greater model capacity is known to improve model generalizability [62]. We can see a decrease in error as the capacity increases.[21] However, it also leads to greater miscalibration. We can see an increase in ECE. In particular, increasing the depth or the number of filters per layer ("width") both result in decreased calibration.
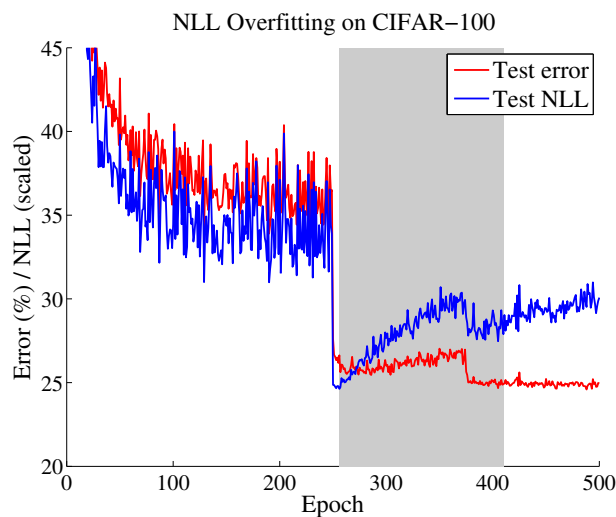
Figure 4.19: Test error and NLL of ResNet-110 over a training run. While the test NLL starts to overfit (i.e., uncertainty estimates become less calibrated), the error keeps decreasing. NLL is scaled in order to fit the Figure. Note the scheduled LR drop at epoch 250. Figure taken from [69].

Let us now turn to Figure 4.19. We can measure predictive uncertainty faithfulness with the

---

[21] Curious readers might find the phenomenon of benign overfitting in the highly overparameterized regime interesting.

test NLL. At epoch 250, we have a scheduled LR drop. Both the test error and test NLL decrease a lot. The grey area is between epochs in which the best validation loss and validation error are produced. The test NLL tends to increase after epoch 250. It shows the overfitting of $c(x)$ to the training samples. It does not go back to epoch 250 levels, not even after the scheduled LR drop at epoch 375. The test error also shows a little overfitting, as it increases by $1 - 2\%$ after epoch 250. However, it drops again after the scheduled LR drop at epoch 375, surpassing epoch 250 levels. The authors draw the following conclusions. "In practice, we observe a disconnect between NLL and accuracy, which may explain the miscalibration in [Figure 4.18]. This disconnect occurs because neural networks can overfit to NLL without overfitting to the 0-1 loss. We observe this trend in the training curves of some miscalibrated models. [Figure 4.19] shows test error and NLL (rescaled to match error) on CIFAR-100 as training progresses. Both error and NLL immediately drop at epoch 250, when the learning rate is dropped; however, NLL overfits during the remainder of the training. Surprisingly, overfitting to NLL is beneficial to classification accuracy. On CIFAR-100, test error drops from 29% to 27% in the region where NLL overfits. This phenomenon renders a concrete explanation of miscalibration: the network learns better classification accuracy at the expense of well-modeled probabilities. We can connect this finding to recent work examining the generalization of large neural networks. Zhang et al. (2017) observe that deep neural networks seemingly violate the common understanding of learning theory that large models with little regularization will not generalize well. The observed disconnect between NLL and 0-1 loss suggests that these high capacity models are not necessarily immune from overfitting, but rather, overfitting manifests in probabilistic error rather than classification error." [69]

4.8.2 Modern Results on Model Calibration
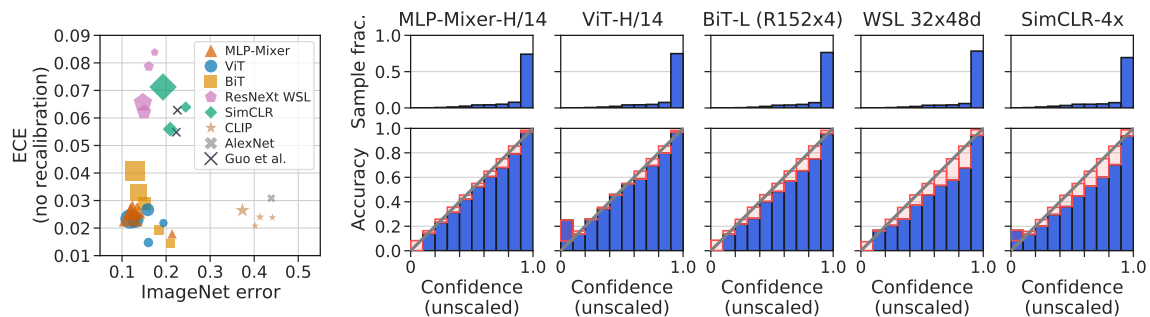
Figure 4.20: The ViT, BiT, and MLP-Mixer architectures are well-calibrated and accurate. *Left.* ECE is plotted against classification error on ImageNet for various classification models. *Right.* Confidence distributions and reliability diagrams of various architectures on ImageNet. "Marker size indicates the relative model size within its family. Points labeled "Guo et al." are the values reported for DenseNet-161 and ResNet-152 in Guo et al. (2017)." [140] Figure taken from [140].

For more recent models, [55] provides an extensive calibration analysis. Several MLP-Mixers [196] (fully connected vision models), ViTs [48] (vision transformers), and BiTs [107] (ResNet-based models) are among the most calibrated *and* accurate models, considering both the NLL loss and the ECE. In particular, knowledge-distilled variants of these usually perform better. This disagreement with the previous study shows that there is no unanimous agreement on the matter of model calibration in the literature.

ViT and Mixer are reported to be well-calibrated [140] in other works as well, however, as shown in Figure 4.20. Notably, no recalibration is performed for the Figure. "Several recent model families (MLP-Mixer, ViT, and BiT) are both highly accurate and well-calibrated compared to prior models, such as AlexNet or the models studied by Guo et al. (2017). This suggests that there may be no continuing trend for highly accurate modern neural networks to be poorly calibrated, as suggested previously. In addition, we find that a recent zero-shot model, CLIP, is well-calibrated given its accuracy." [140]

Calibration depends a lot on the architecture family. There are huge differences even between ConvNet-variants.

Remark: The decrease in ECE values for recent NN-variants *could* also be attributed to them being trained on more data. However, the authors of [140] find that "Model size, pretraining duration, and pretraining dataset size cannot fully explain differences in calibration properties between model families." (Well-calibratedness has a lot to do with overfitting. Increasing the number of training samples could result in better ECE on its own. However, this is apparently not the deciding factor.)

"The poor calibration of past models can often be remedied by post-hoc recalibration such as temperature scaling (Guo et al., 2017), which raises the question of whether a difference between models remains after recalibration. We find that the most recent architectures are better calibrated than past models even after temperature scaling." [140]

4.8.3   Easy Fix for Better ECE: Temperature Scaling

Let us discuss Temperature Scaling [69]. For DNN classifiers, one could fix their calibration via post-processing on the softmax outputs. Suppose that the model output $f(x)$ is the result of a softmax operation over logits $g(x)$:

$$f(x) = \text{softmax}(g(x)) \in \mathbb{R}^K.$$

Softmax converts the logits to parameters of a categorical distribution. We define temperature scaling with the temperature $T > 0$ as follows:

$$f(x; T) = \text{softmax}(g(x)/T).$$

In words, we divide each logit value by T.

When $T \downarrow 0$, the elements of the argument of the softmax explode to infinity, the differences between the $\arg\max$ and the other elements increase more and more. Thus, the output of softmax, $f(x; T)$, becomes a one-hot vector. (As the difference grows, we are stressing the argmax value more and more.)

When $T \to \infty$, the elements of the argument of the softmax go to 0. The differences between the elements decrease more and more. Thus, the output of softmax, $f(x; T)$, becomes uniform.

One can find the $T > 0$ that returns the best ECE score over a validation set. We let the model's

predictive confidence be

$$\left\{ \max_k f_k(x_i; T) \right\}_{i=1,\dots,N_{\text{val}}}$$

over the validation set and search for the $T > 0$ that minimizes the ECE. We can perform a grid search over different T values and find the one that works best.

Temperature scaling improves calibration quite dramatically. Results are shown in Table 4.2. $T = 1$ usually results in suboptimal ECE results; the models are not well-calibrated. $T = T_{\text{val}}^*$ (after performing the search over the val set) results in sub-2% ECE values in general, whereas before the average was around $8\text{-}10\%$. This is a nice and easy fix.[22]

Table 4.2: Comparison of temperature scaling with an untuned baseline. Temperature scaling can lead to a drastic improvement in calibration. Table adapted from [69].

Dataset	Model	Uncalibrated ($T = 1$)	Temp. Scaling ($T = T_{\text{val}}^*$)
Birds	ResNet 50	9.19%	**1.85%**
Cars	ResNet 50	4.3%	2.35%
CIFAR-10	ResNet 110	4.6%	0.83%
CIFAR-10	ResNet 110 (SD)	4.12%	**0.6%**
CIFAR-10	Wide ResNet 32	4.52%	**0.54%**
CIFAR-10	DenseNet 40	3.28%	**0.33%**
CIFAR-10	LeNet 5	3.02%	**0.93%**
CIFAR-100	ResNet 110	16.53%	**1.26%**
CIFAR-100	ResNet 110 (SD)	12.67%	0.96%
CIFAR-100	Wide ResNet 32	15.0%	**2.32%**
CIFAR-100	DenseNet 40	10.37%	1.18%
CIFAR-100	LeNet 5	4.85%	**2.02%**
ImageNet	DenseNet 161	6.28%	**1.99%**
ImageNet	ResNet 152	5.48%	**1.86%**
SVHN	ResNet 152 (SD)	0.44%	0.17%
20 News	DAN 3	8.02%	4.11%
Reuters	DAN 3	0.85%	0.91%
SST Binary	TreeLSTM	6.63%	1.84%
SST Fine Grained	TreeLSTM	6.71%	2.56%

4.9 Do we really need proper scoring?

4.9.1 Ranking Condition

The previous proper scoring rules for the correctness of prediction demanded that $c(x) = P(L = 1 \mid X = x)$ be their optimal value, i.e., that confidences directly give the probabilities of correctness. Calibration followed a similar principle. Let us now consider slightly weaker ranking conditions.

$$\text{If } P(L = 1 \mid x_1) > P(L = 1 \mid x_2) \text{ then } c(x_1) > c(x_2).$$

That is, we want to have the confidence values in the right order. Instead of requiring $c(x)$ to be equal to the actual probability, we only require that the ranking is preserved. If this condition holds, there exists a monotonic calibration function $g \colon \mathbb{R} \to \mathbb{R}$ such that $g(c(X)) =$

---

[22]ECE here is calculated with 15 bins. We can already see that $M = 10$ is not consistently applied through papers, though it is a popular choice.

$P(L = 1 \mid X)$ for input variable X. That is, the ranking condition is almost the same as the calibration condition, up to a monotonic transformation. (We, of course, would have to find this g as a post-processing step if we wanted truthful predictive uncertainties.) This is more approachable than requiring DNNs to be outputting the true confidence values. And it is, in fact, sufficient for many applications, such as when we filter out too-uncertain examples via a threshold.

Based on this intuition, people have produced different metrics for quantifying the ranking condition. Essentially, we have two ingredients:

1. **Confidence estimates.** $c_i := c(x_i) \in \mathbb{R}$ is the *unnormalized* confidence value for test sample x_i.
2. **Correctness of prediction.** $L_i := \mathbf{1}(\arg\max_k f_k(x_i) = y_i) \in \{0, 1\}$ for test sample x_i.

Instead of trying to estimate the true predictive uncertainty p_i from L_i and comparing ranking (we can do this with binning), one may use the raw binary L_i to benchmark the c_i estimates. In ECE, we binned the confidence values (restricted to $[0, 1]$) and took the average of the L_is in the bin, which was our estimate of p_i (very coarse). Now we simply use the raw binary values and benchmark how predictive the confidence estimates are for the L_i values per sample.

We turn the task into a binary detection task for L_i, where the only feature is c_i. The question is: Can c_i tell us anything about the prediction correctness?

4.9.2 Binary Detection Metrics

Given features c_i and target binary labels L_i as well as a threshold $t \in \mathbb{R}$, we predict 1 ("correct") when $c_i \geq t$ and 0 when $c_i < t$. This lets us define the following index sets:

$$\text{True positives: } \mathrm{TP}(t) = \{i : L_i = 1 \wedge c_i \geq t\}$$
$$\text{False positives: } \mathrm{FP}(t) = \{i : L_i = 0 \wedge c_i \geq t\}$$
$$\text{False negatives: } \mathrm{FN}(t) = \{i : L_i = 1 \wedge c_i < t\}$$
$$\text{True negatives: } \mathrm{TN}(t) = \{i : L_i = 0 \wedge c_i < t\}$$
$$\mathrm{Precision}(t) = \frac{|\mathrm{TP}(t)|}{|\mathrm{TP}(t)| + |\mathrm{FP}(t)|}$$
$$\mathrm{Recall}(t) = \frac{|\mathrm{TP}(t)|}{|\mathrm{TP}(t)| + |\mathrm{FN}(t)|}.$$

Informally, precision tells us how pure our positive predictions are at threshold t. Out of the positively predicted samples, how many were correct? Similarly, recall tells us how many of the actual positive samples in the dataset are recalled (predicted positive) at threshold t.

One can draw a curve for $\mathrm{Precision}(t)$ and $\mathrm{Recall}(t)$ for all possible thresholds t from $-\infty$ to $+\infty$ or, for a probability c_i, from 0 to 1. This is the *precision-recall curve*, shown in Figure 4.21.

As we go on the recall axis from left to right, we observe the following values for precision and recall. First, we predict all samples as negative. In this case, precision is undefined. Then we recall the sample with the highest c_i that is actually positive. Recall is almost 0, and precision is 1. We continue..., and at the last point, we recall all actual positive samples (i.e., the recall is one). As we predict everything to be positive, the precision is the fraction of true positive
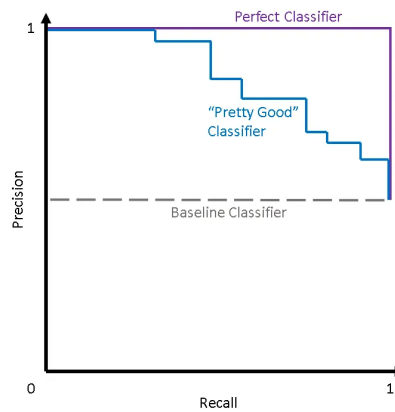
Figure 4.21: Example precision-recall curve that showcases a random classifier, a perfect one, and one in between. Figure taken from [189].

samples. This point is always on the line of the random detector.

To summarize this curve, we can compute the area under the precision-recall curve (AUPR). This is a metric for how well we are predicting (how correct our predictions are based on c_i values). For the perfect detector, $\mathrm{AUPR} = 1$. While we recall all the actual positive samples, we also never recall actual negative samples. For a random detector, $\mathrm{AUPR} = P(L = 1)$ where $P(L = 1)$ is the ratio of positive samples in the dataset. AUPR can be calculated in two ways: AUPR-Success is the method we discussed above. In AUPR-Error, we use errors ($L = 0$) as the positive class. Both are often reported together for predictive uncertainty evaluation.

A drawback of the AUPR is that the random classifier's performance depends on $P(L = 1)$. For example, if $P(L = 1) = 0.99$ (i.e., the test set is severely imbalanced), then AUPR is already 99% for a random detector. It lacks the resolution to see the improvement above the random detector baseline.

The Receiver Operating Characteristic (ROC) curve fixes this. It compares the following quantities:

$$\mathrm{TPR}(t) = \mathrm{Recall}(t) = \frac{|\mathrm{TP}(t)|}{|\mathrm{TP}(t)| + |\mathrm{FN}(t)|} = \frac{|\mathrm{TP}(t)|}{|\mathrm{P}|}$$

$$\mathrm{FPR}(t) = \frac{|\mathrm{FP}(t)|}{|\mathrm{FP}(t)| + |\mathrm{TN}(t)|} = \frac{|\mathrm{FP}(t)|}{|\mathrm{N}|}.$$

Here, FPR tells us how many of the actual negative samples in the dataset are recalled (predicted positive) at threshold t. This is "1 - the recall for the negative samples."

Similarly to the Precision-Recall curve, one can draw a curve of $\mathrm{TPR}(t)$ and $\mathrm{FPR}(t)$ for all t from $-\infty$ to $+\infty$ or, for a probability c_i, from 0 to 1. This is the *ROC curve*, shown in Figure 4.22.

As we go on the FPR axis from left to right, the FPR and TPR values change as follows. First, we predict all samples as negative. There, TPR is 0, and FPR is 0. We continue until the last point, where we predict all samples as positive. There, TPR is 1, and FPR is 1.
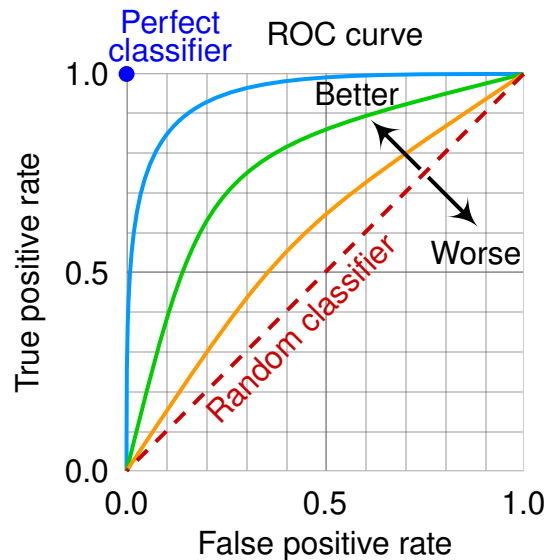
Figure 4.22: Example ROC curve showing results for a perfect classifier, a random one, and ones in between. Figure taken from [41].

The area under the ROC curve (AUROC) can be computed as a summary metric. The AUROC has a nice interpretation: It gives the probability that a correct sample ($L = 1$) has a higher certainty $c(x)$ than an incorrect one. This very much captures our ranking goal. For the perfect ordering AUROC $= 1$. And, interestingly, for a random order, AUROC $= 0.5$, regardless of $P(L = 1)$. *This makes AUROC the recommended metric over AUPR, especially on unbalanced datasets.*

4.10 $c(x)$ **as Non-Predictive Uncertainty**

So far, we have expected $c(x)$ to be an estimate of the predictive (un)certainty – whether the model is going to get the answer right or wrong. $c(x)$, the *confidence estimate*, was required to be a good representation of the likelihood of getting the answer right ($L = 1$).

However, we have discussed two more equally important uncertainties: the epistemic and the aleatoric components. We can design benchmarks for each of these sources separately, i.e., measure the quality of a particular $c(x)$ as the predictor for other factors (i.e., not predictive uncertainty anymore). Here, we impose no restrictions on the estimator $c(x)$ we might use, only that it returns a probability $\in [0, 1]$ for a binary prediction task. Possibilities for non-predictive uncertainty benchmarks are listed below.

Is the sample x an OOD sample? In this case, we can treat $c(x)$ is an OOD detector. This is not perfectly aligned with predictive uncertainty. Even if a sample is OOD, the model might get the answer confidently right, and even if it is ID, the model can be unconfident. It is rather a measure of epistemic uncertainty.

Is the sample x severely corrupted? Corruption is related to predictive uncertainty, but they are not perfectly aligned: The level of corruption in an input sample is only one source of uncertainty. This aspect has close ties to aleatoric uncertainty.

Does the sample x admit multiple answers? This is – by definition – aleatoric uncertainty.

As can be seen, these questions are more closely related to identifying particular aspects of uncertainties tied to either epistemic or aleatoric sources.

4.10.1 $c(x)$ as an OOD Detector

We write Y for the binary variable indicating

- $Y = 1$ if x is from outside the training distribution.
- $Y = 0$ if x is from inside the training distribution.

Then, we would expect high $P(Y = 1)$ for higher uncertainty values $1 - c(x)$. That is, $c(x)$ shall be a good estimator for epistemic uncertainty. $c(x)$ can be, again, treated as a feature for the binary prediction of OOD-ness. We may then evaluate $c(x)$ for its OOD detection performance with AUPR or AUROC. These are evaluation metrics we already know from predictive uncertainty that are generally used for *ranking* uncertainties. In the literature for OOD detection or general uncertainty estimation, we often see OOD detection performances reported in terms of area under curve metrics.

4.10.2 $c(x)$ as a Multiplicity Detector

This is a much less popular choice. We write Y for the binary variable indicating

- $Y = 1$ if the true label for x has multiple possibilities, maybe because of inherent ambiguity in the task or due to corruption.
- $Y = 0$ if there exists a unique label for x.

Then, we would expect high $P(Y = 1)$ for higher uncertainty values $1 - c(x)$. That is, $c(x)$ shall be a good estimator for aleatoric uncertainty (whether the sample accommodates more than one answer). $c(x)$ can be, again, treated as a feature for the binary prediction of aleatoric uncertainty. We may then evaluate $c(x)$ for its multiplicity detection performance with AUPR or AUROC.

4.10.3 Summary of Evaluation Methods so far for Uncertainty

For *predictive uncertainty* (whether the model is going to get the prediction right), we have seen (1) proper scoring rules such as log probability and Brier score, (2) metrics based on model calibration such as ECE, MCE, and reliability diagrams (that are more intuitive metrics), and (3) ranking (or "weak calibration") using AUROC or AUPR. The third approach uses different thresholds for the retrieval of correctly predicted samples.

If we only care about *epistemic uncertainty*, it makes sense to consider the downstream proxy task of OOD detection to measure the quality of our uncertainty estimates. We can measure OOD detection performance using AUROC or AUPR. Plotting the ROC or precision-recall curves can also be insightful.

For *aleatoric uncertainty*, one might want to look at multiplicity/corruption detection. Detection performance can be, again, measured by AUROC or AUPR.

4.11   Estimating Epistemic Uncertainty

As we have seen, epistemic uncertainty means we are unsure about our prediction because several models could fit the training data (because we have not experienced enough training data to distinguish the correct from the incorrect model.) There are two possibilities:

1. The size of our training set is too small, and so the variance of our estimator is too high.
2. The training data distribution does not cover some meaningful regions in the input space; there are some underexplored areas.

Epistemic uncertainty has a close connection with Bayesian machine learning. A great tool for dealing with multiple possibilities in maths is probability theory:

$$P(\theta \mid \mathcal{D}) \propto P(\theta)P(\mathcal{D} \mid \theta) = P(\theta)\prod_{i=1}^{N} P(x_i \mid \theta).$$

A posterior distribution over the parameter space is the Bayesian way of saying, "This space accommodates multiple possible solutions after observing the training set and taking our prior beliefs into consideration." A "wider" distribution means higher uncertainty regarding the true model. (The one that "generated" the dataset.)

It can be instructive to consider the "input space point of view": We are adding more and more observations to underexplored regions of the input space. These give more and more supervision: We are narrowing down the possible range of θs based on the observations. This should ideally be happening with Bayesian ML as we observe more data.

4.11.1   Space of Model Parameters θ

This space is at the center of our attention in Bayesian ML. The notion of parameters θ is often interchangeably used with weights w and, sadly, also with functions, models, or hypotheses h. Using Bayesian inference

$$P(\theta \mid \mathcal{D}) \propto P(\theta)P(\mathcal{D} \mid \theta) = P(\theta)\prod_{i=1}^{N} P(x_i \mid \theta),$$

we are narrowing down our hypothesis space from the wide prior space by observing more and more data until we arrive at the final posterior. We hope this distribution contains the true model (the one that actually "generated" the dataset) with high probability. Figure 4.23(b) is the ideal visualization of what should happen with Bayesian ML.

4.11.2   Approximate Posterior Distribution Families

In the previous section, we discussed why a posterior over models is a great way to represent multiple possibilities. In most cases, however, the true posterior (i.e., the one given by Bayes' rule) over our weights/models is intractable. (The prior specification is also often left implicit.) Therefore, we have to make some approximations to our true posterior. We need to define the distributional format of our posterior approximations; in other words, the approximate posterior distribution family. The posterior $P(\theta \mid \mathcal{D})$ can be thought of as an infinite set of models (using sensible priors). We denote our approximation by $Q_\phi(\theta)$, where ϕ are
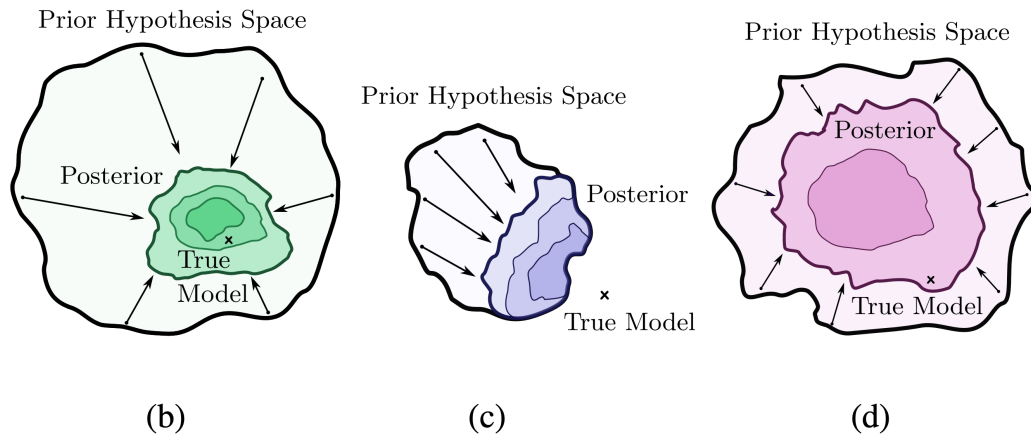
Figure 4.23: Different scenarios for optimization in the hypothesis space. "(b) By representing a large hypothesis space, a model can contract around a true solution, which in the real world is often very sophisticated. (c) With truncated support, a model will converge to an erroneous solution. (d) Even if the hypothesis space contains the truth, a model will not efficiently contract unless it also has reasonable inductive biases." [211] Figure taken from [211].

the parameters of this parametric distribution. The posterior is often approximated without explicitly specifying the prior.

Definition 4.16: Dirac Delta Measure

The Dirac delta is a generalized function over the real numbers whose value is zero everywhere except at zero. For our purposes, it represents the fact that we only have one possible parameter configuration θ in our posterior. Formally, it is a measure. Without going too much into Lebesgue integration theory, the gist is that it acts like a Kronecker delta.

Note: This definition only acts as an intuitive description of the Dirac measure. Interested readers should refer to measure theoretical treatments of the notion.

$Q_\phi(\theta)$ can be, e.g., ...

- **...a generic multimodal distribution.** For example, it can be a Mixture of Gaussians (MoG), but any other distribution can be chosen. A MoG with an appropriate number of modes is enough to cover any continuous distributions if we allow an arbitrary number of modes.
- **...a uni-modal Gaussian distribution.** Many people like to use this for computational simplicity and tractability.
- **...a sum of Dirac delta distributions.** Some people use such semi-deterministic $Q_\phi(\theta)$s.
- **...a single Dirac delta distribution.** This takes us back to deterministic ML. A deterministic posterior approximation means a single point estimate for θ (MLE, MAP).

Of course, under any sensible prior belief and problem setup, the true posterior $P(\theta \mid \mathcal{D})$ will never be a sum of Dirac deltas. Nevertheless, we might use it as an *approximation* to the true posterior. In this section, we will always *approximate* the true posterior $P(\theta \mid \mathcal{D})$ w.r.t. either

an implicit or explicit prior distribution.

Deterministic vs. Bayesian ML

There is a whole spectrum between probabilistic Bayesian ML and deterministic ML. We may also recover the original deterministic ML formulation by choosing our approximate posterior family to be the family of Dirac deltas. Thus, the Bayesian framework is a generalization of deterministic ML. One could express various forms of posterior uncertainty by considering different approximate posterior distribution families.

Deterministic ML first optimizes a single model (parameter set) over the training set, $\theta^*(\mathcal{D})$. Then, for a test sample, it predicts the label as

$$P(y \mid x, \mathcal{D}) = P(y \mid x, \theta^*(\mathcal{D})).$$

We use only this single model to produce the output for the input of interest. From the Bayesian perspective, this is equivalent to having a Dirac posterior. As epistemic uncertainty arises from the existence of multiple plausible models, but we only consider a single one in deterministic ML, we cannot represent epistemic uncertainty using deterministic ML (i.e., we treat is as 0).

Bayesian ML finds a distribution of models, $Q_\phi(\theta \mid \mathcal{D})$, the approximate posterior over the models after observing the training data. Think of Bayesian ML as training an infinite number of models simultaneously (whenever our approximate posterior does not only accommodate a finite set of models).

Quantifying Epistemic Uncertainty

Now, we have the most important ingredient to represent epistemic uncertainty: a set of models. However, measuring the diversity of this set directly is hard. Therefore, people usually look at the averaged prediction of the models, formalized as follows. For a test sample, Bayesian ML predicts the label using Bayesian Model Averaging (BMA)/marginalization:

$$P(y \mid x, \mathcal{D}) = \int P(y \mid x, \theta) \underbrace{Q_\phi(\theta)}_{\approx P(\theta \mid \mathcal{D})} d\theta = \mathbb{E}_{Q_\phi(\theta)}\left[P(y \mid x, \theta)\right].$$

Thus, we take the average prediction from the approximate posterior distribution (the voting from an "infinite number of models") at test time. This can be further approximated as

$$P(y \mid x, \mathcal{D}) \overset{\text{MC}}{\approx} \frac{1}{M} \sum_{i=1}^{M} P(y \mid x, \theta^{(i)}), \qquad \theta^{(i)} \sim Q_\phi(\theta).$$

The entropy $\mathbb{H}(P(y \mid x, \mathcal{D}))$ or the max-prob for classification $\max_k P(Y = k \mid x, \mathcal{D})$ are popular choices to quantify epistemic uncertainty.

Intuition of BMA. We expect the outputs of all models in the posterior to be similar on the training data, as we explicitly train the models on the training set. When we have a test sample in the training data region, we expect $P(y \mid x, \theta)$ (i.e., the vector of probabilities in classification) to be similar across the models, as the sample will probably lie on the same

side of the decision boundaries of the models (which gets tricky to think about in multi-class classification). The models will also be confident in the predictions (up to aleatoric uncertainty), having been trained on similar samples. Therefore, the BMA output $P(y \mid x, \mathcal{D})$ will show high confidence (e.g., it will have max-prob = 99%). When we have a test sample in an underexplored region, we expect the individual $P(y \mid x, \theta)$s to be divergent, as nothing forces the models' decision boundaries to agree in these regions (as we have not trained on samples from these regions).[23] Therefore, the models give divergent answers (i.e., the max-prob indices are different). The BMA output will show low confidence: e.g. max-prob = 59%. Averaging/integrating gives us a mixture, and the maximal value of the mixture will be more smoothed out. Even if the individual models are overconfident, the average output will not be. By averaging, the arg max can even become different from all individual arg maxes. For example, in the case of a discrete set of models,

$$\mathrm{avg}\left((0.51, 0.01, 0.48), (0.01, 0.51, 0.48)\right) = (0.26, 0.26, 0.48).$$

To provide further intuition for why the BMA can represent epistemic uncertainty: Models are sure about different things; when we average their outputs, it makes the ensemble more unsure. Thus, we get better epistemic uncertainty prediction. **Note**: The BMA output still contains the aleatoric uncertainty represented by the individual models – the BMA represents predictive uncertainty (both epistemic and aleatoric uncertainty) in the most precise sense.

4.11.3 Ensembling

Since ensembling is hugely successful in practice, we will focus a bit more on it in the next sections. Ensemble learning is usually done as follows (popularized by Balaji et al. [116]).

1. Select M different random seeds. These are different starting points for the optimization in the parameter space.
2. Train the M models regularly, using either a bagged dataset for each model (where we sample with replacement from the original training set) or the original training set.

As the loss landscape is highly non-convex, we usually end up with different local minima depending on where we start. Therefore, we usually get a diverse set of models. Random seeds also control the noise on the objective function (loss landscape) itself, not only the starting points on "the" landscape. The seeds influence …

- **… the formation of batches of training samples for SGD.** If we change the seed, we change the batching, as the reshuffling of the dataset is seeded differently.
- **… the random components of the data augmentation process.** Therefore, the actual loss landscape is also changed. We almost always perform data augmentation.
- **… the random network components, such as Dropout, DropConnect, or Stochastic Depth.** In DropConnect (2013) [202], instead of dropping out activations (neurons), we drop connections between neurons in subsequent layers. Stochastic Depth (2016) [82] shrinks the network's depth during training, keeping it unchanged during testing. It randomly drops entire ResBlocks during training:

$$H_l = \mathrm{ReLU}(b_l f_l(H_{l-1}) + \mathrm{Id}(H_{l-1}))$$

---

[23]Here, we also need the model posterior we obtain to represent a diverse set of plausible models.

where b_l is a binary random variable.

Changing the random seed, therefore, changes many things, which usually encourages enough diversity in our ensemble. Using bagging to obtain separate training sets for each model further encourages diversity.

BMA with Ensembles

In *model ensembling*, we train several deterministic models on the same (or subsampled) data simultaneously. In this case, our posterior approximation becomes a mixture of M Dirac deltas, where M is the number of models in our ensemble. We claim that this is Bayesian.

$$Q_\phi(\theta) = Q_{\theta^{(1)},\ldots,\theta^{(M)}}(\theta) = \frac{1}{M} \sum_{m=1}^{M} \delta(\theta - \theta^{(m)}).$$

After training the M models, BMA boils down to taking the average over the ensemble members' predictions

$$
\begin{aligned}
P(y \mid x, \mathcal{D}) &= \int P(y \mid x, \theta) P(\theta \mid \mathcal{D}) \, d\theta \\
&\approx \int P(y \mid x, \theta) Q_\phi(\theta) \, d\theta \\
&= \int P(y \mid x, \theta) \frac{1}{M} \sum_{m=1}^{M} \delta(\theta - \theta^{(m)}) \, d\theta \\
&= \frac{1}{M} \sum_{m=1}^{M} \int P(y \mid x, \theta) \delta(\theta - \theta^{(m)}) \, d\theta \\
&= \frac{1}{M} \sum_{m=1}^{M} P(y \mid x, \theta^{(m)}).
\end{aligned}
$$

If the reader is not well versed in measure theory, the last equality can be considered a part of the Dirac measure's definition.[24] This corresponds to averaging the predictions of individual models. y can be a scalar value in regression, one particular class in a classification setting, or even the whole class distribution.

Previously, we have discussed an intuitive explanation for why the BMA can represent epistemic uncertainty. On the side, ensembles also often provide better accuracy. The intuition here is that single predictors make different mistakes and overfit differently. This noise cancels out by averaging, and we get a better test accuracy. This phenomenon is formalized and widely used in statistics: Readers might find the various techniques for bootstrap aggregation (or bagging) interesting.

Let us collect the pros and cons of ensembles.

- **Pro**:
    - Conceptually simple – run the training algorithm M times and average outputs.
    - Applicable to a wide range of models – from linear regression to ChatGPT.

---

[24]When considering the Dirac measure, one should write $\int P(y \mid x, \theta) d\delta(\theta - \theta^{(m)})$, which is a rigorous form of Lebesgue integration.

- Parallelizable – if we have a lot of computational resources, we can train multiple models simultaneously on different cluster nodes (GPUs).
    - Performant – ensembles are not only able to represent epistemic uncertainty but are also often more accurate.
- **Contra**:
    - Ensembles do not realize the full potential of Bayesian ML (no infinite number of models, no connectivity between the models).
    - Space and time complexities scale linearly with M. If we have a limited number of GPUs, we must wait until the previous model finishes training (the same holds even for evaluation). Compute scales linearly even if we parallelize, time might not. To summarize, this does not scale nicely. However, we often share some weights to increase the number of models we include in the ensemble (e.g., to infinity). We will discuss several approaches to training an "infinite number of models" below.

Finally, we note that ensembling roughly approximates the true posterior that is given w.r.t. the weight initialization scheme, which is our implicit prior. In other methods, we have no such connections, and the (implicit) prior remains undisclosed.

4.11.4   Dropout

Having a combinatorial number of models during training sounds familiar. We have used dropout for model training for quite some time. When using dropout [188], we sample the dropout masks in every iteration, so a different model is being trained at every iteration. The models are, of course, very correlated. Every time we are training our net with different neurons missing. This is an ensemble of many models. We train each of them for just a couple of steps, but they are so similar that optimizing one model translates over to improving the other models too. On the spectrum of Bayesian methods, dropout is between the sum of Diracs (training a few models) and the variational approach (that trains an infinite number of models). Some people say dropout is Bayesian.

The dropout objective is

$$\frac{1}{N} \sum_{n=1}^{N} \log P(y_n \mid x_n, s \odot \theta).$$

This is a simple CE loss over the training dataset, but we turn on/off each weight dimension randomly in each iteration: $s^{(i)} \sim \mathrm{Bern}(s \mid p)$. This is very similar to the BBB data term. However, we draw $\theta(s) = s \odot \theta$ from a huge discrete, categorical distribution, not a Gaussian.

What makes this interesting to create ensemble predictions is that we can also use dropout at inference time, as introduced in the paper "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning" [54]. Eventually, any configuration of turning on/off parameters may fit the given training data well (resulting in low NLL loss). This does not have to be the case for non-training data: there will be disagreements between models for OOD samples. The method is good for detecting such OOD samples. Although we have always trained many models simultaneously using dropout, we have not taken advantage of that during inference. To apply dropout at inference time (test time), we do BMA across

different Bernoulli mask choices (different weight samples) and average the predictions:

$$P(y \mid x, \mathcal{D}) = \int P(y \mid x, \theta) P(\theta \mid \mathcal{D}) \, d\theta$$

$$\approx \int P(y \mid x, \theta) Q_\phi(\theta) \, d\theta$$

$$\overset{\text{MC}}{\approx} \frac{1}{K} \sum_{k=1}^{K} P(y \mid x, \theta^{(k)}) \qquad s^{(k)} \sim \text{Bern}(s \mid p), \theta^{(k)} = s^{(k)} \odot \theta.$$

4.11.5   Evaluation of Ensembling and Dropout in Practice

Let us discuss a paper on evaluating ensembling and dropout in practice [116].

Results of Ensembling

We start with Figure 4.24. In distribution, the spread of the output categorical distribution is minimal, no matter how many models we have in the ensemble. The prediction is nearly always close to a one-hot vector. The spread here is measured by the entropy of the distribution. Entropy 0 means the categorically distributed random variable is constant, p is a true one-hot vector. High entropy means p is close to being uniform. Out of distribution, a single model still produces close to 0 entropy values. However, the ensemble over more and more models has increasing entropy on the OOD samples.
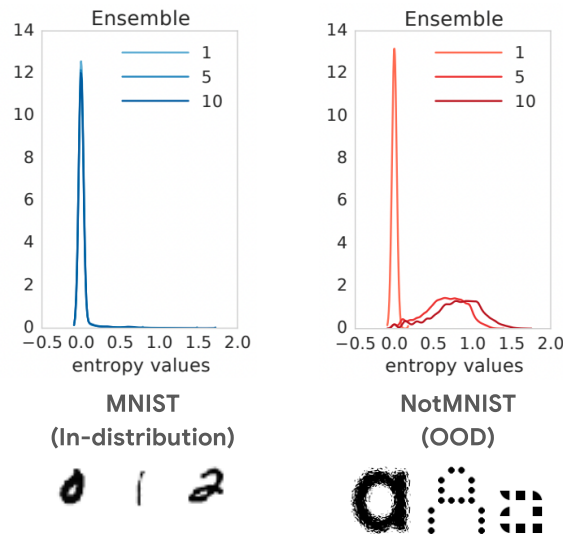
Figure 4.24: Entropy values on ID and OOD datasets with a varying number of models in the ensemble. Ensembling results in higher entropy values on OOD samples. Base figure taken from [116].

Finally, let us consider the quantitative results of Table 4.3 from an ImageNet experiment. Ensembling also works at the ImageNet scale. Adding more members to the ensemble decreases error and increases accuracy. (Training on NLL also improves accuracy, not just uncertainty estimates.) Test NLL and Brier scores also improve by increasing the number of models. One could conclude that we obtain better predictive and aleatoric uncertainties.

Table 4.3: Quantitative results of ensembling on ImageNet. All considered metrics improve with more models. Both the NLL and Brier scores correlate calibration with accuracy. Table taken from [116].

M	Top-1 error %	Top-5 error %	NLL	Brier Score $\times 10^{-3}$
1	22.166	6.129	0.959	0.317
2	20.462	5.274	0.867	0.294
3	19.709	4.955	0.836	0.286
4	19.334	4.723	0.818	0.282
5	19.104	4.637	0.809	0.280
6	18.986	4.532	0.803	0.278
7	18.860	4.485	0.797	0.277
8	18.771	4.430	0.794	0.276
9	18.728	4.373	0.791	0.276
10	18.675	4.364	0.789	0.275

However, it could also be the case that the improvements in these scores are just due to the higher accuracy. Drawing conclusions from proper scoring rule values is, therefore, tricky.

Comparison of Ensembling and Dropout

Ensembling and dropout seem to be plausible ways to represent epistemic uncertainty. Let us now focus on the top part of Figure 4.25. We take the NLL and the Brier Score of the true label.

Ensembling. As we add more and more nets to the ensemble, we see a decrease in the classification error (or, equivalently, an increase in accuracy). This is not surprising, as everyone is doing ensembling to get better accuracies. Ensembling with more models also seems to produce better aleatoric and predictive uncertainty estimation. We can conclude this because, for multi-class classification, the log probability scoring rule and the multi-class Brier score are strictly proper scoring rules for both *predictive* and *aleatoric* uncertainty estimation using max-prob. Therefore, by measuring the log-likelihood and the multi-class Brier score, we are also measuring how far away we are from perfect aleatoric uncertainty prediction. **Note**: By training on NLL, we encourage each model to give correct predictive uncertainties on the training set, and we also ensemble to get correct epistemic uncertainties. The models usually generalize better by ensembling, and we also get better predictive uncertainties on the test samples. Ensembling seems to work for a small dataset and a simple neural network.

Dropout. Sampling more and more nets from dropout seems to plateau quite early and at notably worse values than what we can achieve by ensembling. These days, MC dropout is treated as a method that does not really work. Many people are critical of it.

Note: Aleatoric uncertainty cannot be reduced by ensembling or using dropout: It is completely independent of the model. However, the model posterior might become better at *modeling* the aleatoric uncertainty.

Let us now turn to the bottom part of Figure 4.25 that evaluates a VGG-style ConvNet on SVHN (street view house numbers). Ensembling is also scalable to large models and "large" datasets. We can use ensembling for any model: We simply have to average the outputs.
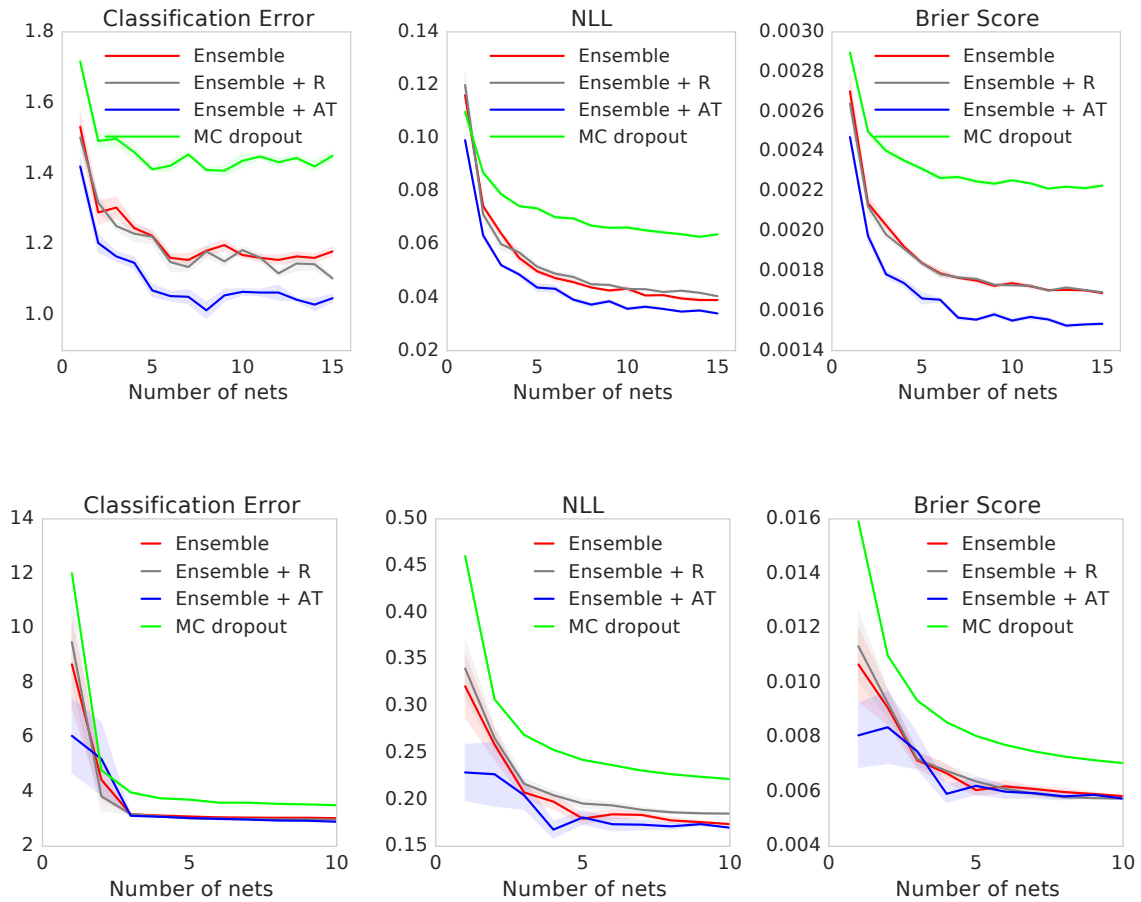
Figure 4.25: *Top.* Evaluation of epistemic uncertainty estimation methods on the MNIST dataset using a 3-layer MLP. *Bottom.* Evaluation on the SVHN dataset using a VGG-style convnet. In both cases, ensembling improves both accuracy and proper scoring metrics. Dropout plateaus earlier and gives suboptimal results. AT: Adversarial training added. R: Random signed vector added (baseline, no difference). Figure taken from [116].

4.11.6 Training an Infinite Number of Models – Bayes By Backprop

Now, we consider a method for training an infinite number of models, called Bayes By Backprop. Training an infinite number of models is possible when the approximate posterior is a continuous distribution (e.g., Gaussian). Expressing infinite possibilities with a finite number of parameters can be easily achieved using parameterized probability distributions. This work explicitly models $P(\theta)$ and approximates the true posterior w.r.t. this prior and the training likelihood by

$$P(\theta \mid \mathcal{D}) \approx \mathcal{N}\left(\theta \mid \mu^*(\mathcal{D}), \Sigma^*(\mathcal{D})\right) =: Q_\phi(\theta),$$

where * denotes the μ and Σ values attained by training on dataset $\mathcal{D}$. We simply model the mean and variance, assuming the posterior is approximately Gaussian. Instead of training θ directly, we are training μ and Σ on $\mathcal{D}$. θs are just samples from the Gaussian. One can choose $P(\theta)$ arbitrarily. However, to keep things closed-form, one usually also chooses a Gaussian.

Of course, we know that the true posterior w.r.t. the chosen prior is likely not a Gaussian. It is usually much more complex. Nevertheless, we may still search for the best Gaussian describing the posterior. This is called *variational approximation/inference*. We minimize the

"distance" between our true posterior and the Gaussian approximation:

$$\min_{\mu,\Sigma} d\left(\mathcal{N}\left(\mu(\mathcal{D}), \Sigma(\mathcal{D})\right), P(\theta \mid \mathcal{D})\right).$$

A popular choice for measuring the divergence (not distance!) between two distributions is the Kullback-Leibler (KL) divergence. With that choice, our problem becomes

$$\min_{\mu,\Sigma} \mathrm{KL}\left(\mathcal{N}\left(\mu(\mathcal{D}), \Sigma(\mathcal{D})\right) \parallel P(\theta \mid \mathcal{D})\right).$$

Training this directly is impossible, as we do not know the true posterior. However, we can still derive an equivalent optimization problem that does not require us to calculate the exact posterior. Figure 4.26 illustrates this optimization problem.
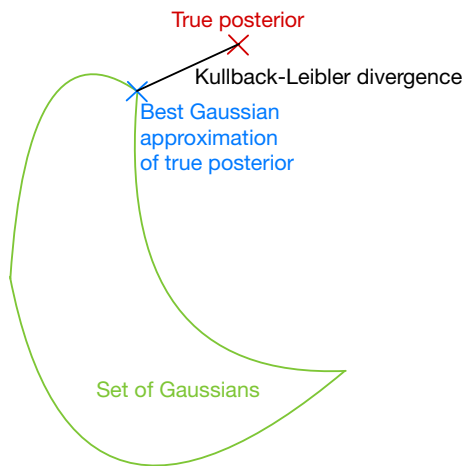
Figure 4.26: Informal illustration of the Bayes By Backprop optimization problem. The procedure aims to find the best Gaussian approximation of the true posterior. The use divergence between the two is the KL divergence.

Using the fact that

$$
\begin{aligned}
\log \frac{1}{P(\theta \mid \mathcal{D})} &= -\log P(\theta \mid \mathcal{D}) \\
&= -\log \frac{P(\mathcal{D} \mid \theta)P(\theta)}{P(\mathcal{D})} \\
&= -\log P(\mathcal{D} \mid \theta)P(\theta) + \log P(\mathcal{D}) \\
&= \log \frac{1}{P(\mathcal{D} \mid \theta)P(\theta)} + C
\end{aligned}
$$

and

$$\log P(\mathcal{D} \mid \theta) \overset{\text{IID}}{=} \log \prod_{n=1}^{N} P(x_n, y_n \mid \theta)$$

$$= \log \prod_{n=1}^{N} \left(P(y_n \mid x_n, \theta) P(x_n \mid \theta) \right)$$

$$= \sum_{n=1}^{N} \left(\log P(y_n \mid x_n, \theta) + \log P(x_n) \right) \qquad (x_n \perp\!\!\!\perp \theta)$$

$$= \sum_{n=1}^{N} \log P(y_n \mid x_n, \theta) + \underbrace{\sum_{n=1}^{N} \log P(x_n)}_{C'},$$

we rewrite our training objective as

$$\mathrm{KL}\left(\mathcal{N}\left(\mu(\mathcal{D}), \Sigma(\mathcal{D})\right) \,\|\, P(\theta \mid \mathcal{D})\right)$$

$$= \int \mathcal{N}\left(\theta \mid \mu, \Sigma\right) \log \frac{\mathcal{N}\left(\theta \mid \mu, \Sigma\right)}{P(\theta \mid \mathcal{D})} \, d\theta$$

$$= \int \mathcal{N}\left(\theta \mid \mu, \Sigma\right) \log \frac{\mathcal{N}\left(\theta \mid \mu, \Sigma\right)}{P(\mathcal{D} \mid \theta) P(\theta)} \, d\theta + C$$

$$= \int \mathcal{N}\left(\theta \mid \mu, \Sigma\right) \log \frac{\mathcal{N}\left(\theta \mid \mu, \Sigma\right)}{P(\theta)} \, d\theta - \int \mathcal{N}\left(\theta \mid \mu, \Sigma\right) \log P(\mathcal{D} \mid \theta) \, d\theta + C$$

$$= \mathrm{KL}\left(\mathcal{N}(\mu, \Sigma) \,\|\, P(\theta)\right) - \underbrace{\mathbb{E}_{\theta \sim \mathcal{N}(\mu, \Sigma)} \log P(\mathcal{D} \mid \theta)}_{\text{CE}} + C$$

$$= \mathrm{KL}\left(\mathcal{N}(\mu, \Sigma) \,\|\, P(\theta)\right) - \sum_{n} \mathbb{E}_{\theta \sim \mathcal{N}(\mu, \Sigma)} \log P(y_n \mid x_n, \theta) + C$$

$$\overset{\text{MC}}{\approx} \mathrm{KL}\left(\mathcal{N}(\mu, \Sigma) \,\|\, P(\theta)\right) - \frac{1}{K} \sum_{n=1}^{N} \sum_{k=1}^{K} \log P(y_n \mid x_n, \theta^{(k)}) + C$$

where $\theta^{(k)} \sim \mathcal{N}(\mu, \Sigma)$ and we collapse all terms into C that do not contain μ and Σ, the parameters we optimize. In the MC sampling, usually, we take $K = 1$ for training. This is usually fine because we are MC estimating the expected gradient anyway, with a small batch size (SGD). This expectation approximation can also be made coarse, as noise in SGD was shown to be a regularizer and promote better generalization [185].

Our final optimization problem is thus

$$\min_{\mu, \Sigma} \mathrm{KL}\left(\mathcal{N}(\mu, \Sigma) \,\|\, P(\theta)\right) - \frac{1}{K} \sum_{n} \sum_{k} \log P(y_n \mid x_n, \theta^{(k)}) \qquad \theta^{(k)} \sim \mathcal{N}(\mu, \Sigma).$$

The first term is the prior term, the regularizer. The second term is the data term, the likelihood. We took conceptual, rigorous steps to justify what we are deriving, but this equation makes sense on its own as well.

This is already a convenient loss function, but we want to make it *more DNN-friendly*. We have complete freedom to choose the prior for the KL term. We only need to encode our beliefs through our prior, which can be anything. (This, of course, influences the true posterior but not the true model that generated the data. We want the true model to have high density in the true posterior.) In the parameter space, there are many symmetries; equivalent solutions

are spread across the entire space. Regardless of which part of the space we choose, it is very likely that we will find a suitable solution locally. This might serve as a weak justification of the choice of a standard normal distribution as the prior:[25]

$$P(\theta) := \mathcal{N}\left(\theta \mid 0, I\right).$$

We also restrict our posterior to Gaussians with diagonal covariance matrices:

$$\Sigma = \operatorname{diag}(\sigma^2).$$

(The full covariance matrix with full degrees of freedom would introduce many computational problems.) Thus, we approximate $P(\theta \mid \mathcal{D})$ with a heteroscedastic diagonal Gaussian. Then the KL divergence can be given in closed form, as it is between two multivariate Gaussians:

$$\mathrm{KL}\left(\mathcal{N}(\mu, \operatorname{diag}(\sigma^2)) \parallel \mathcal{N}(\theta \mid 0, I)\right) = \frac{1}{2} \sum_i \left[\mu_i^2 + \sigma_i^2 - \log \sigma_i^2 - 1\right].$$

The only remaining problem why the loss is not DNN-friendly is that the loss does not depend on μ and Σ straightforwardly. We have to sample from a distribution parameterized by μ, Σ, which is not differentiable in the naive way w.r.t. μ, Σ. The reparameterization trick is used here to detach μ and Σ from the randomness in the approximate posterior. We compute the model parameter via

$$\theta = \mu + \sigma \odot \epsilon,$$

where $\odot$ means pointwise multiplication and $\epsilon \sim \mathcal{N}(0, I)$. We only have to sample ϵs (the random part which does not depend on μ, Σ) and push it through the above transformation to obtain the θ values. This separates the randomness and backpropagation.

Lastly, we need to ensure that the σ vector is always positive. It cannot be just an unbounded parameter, like usual. We counteract this by parameterizing ρ instead (which is a normal `nn.Parameter`), which may take on negative values too, and setting

$$\sigma := \operatorname{softplus}(\rho) = \log(1 + \exp(\rho)) > 0$$

where all operations are element-wise. The actual softplus function also has a hyperparameter β – we keep everything minimal here. Therefore, we obtain a closed-form, differentiable loss for μ, σ without any constraints. An example PyTorch code for BBB in a network with a single linear layer is given in Listing 4.1. We consider multi-class logistic regression in the BBB formulation. It can be trained with backpropagation and SGD.

Variational approximation (which justifies what we are doing theoretically) consists of a prior KL term and a likelihood term. For the likelihood, we sample a parameter θ from the infinite possibilities of models at every iteration. This is the "secret sauce" for training an infinite number of models simultaneously while sharing weights (μ, Σ) and saving computation. To separate the sampling operation from BP (i.e., to have gradient flow to the parameters of the approximate posterior), we use the reparameterization trick. Frequently, we need to clip parameter values to a certain range – we use softplus to ensure $\sigma > 0$.

---

[25]The fact that Gaussianity makes integrals more tractable and the $L-2$ regularization a Gaussian prior imposes is widely known to work well are more convincing arguments

```python
class BBBLinear(nn.Module):
    def __init__(self, input_dim, output_dim):
        super().__init__()
        self.mu = nn.Parameter(
            torch.tensor(input_dim, output_dim).uniform_(-0.1, 0.1)
        )
        self.rho = nn.Parameter(
            torch.tensor(input_dim, output_dim).uniform_(-3, 2)
        )
        # Sizes: number of weights in the model.

    def forward(self, x):
        eps = torch.randn_like(self.mu)  # requires_grad is not propagated, K =
 1
        sigma = F.softplus(self.rho)
        theta = self.mu + sigma * eps
        return x @ theta  # logits

    def compute_loss(self, logits, targets):
        # K = 1, negative sum of log-probs
        neg_log_likelihood = F.cross_entropy(
            logits, targets, reduction="sum"
        )
        sigma = F.softplus(self.rho)
        kl_prior = 0.5 * (
            self.mu ** 2 + sigma ** 2 - torch.log(sigma ** 2) - 1
        ).sum()
        return kl_prior + neg_log_likelihood
```

Listing 4.1: PyTorch code for a linear BBB classifier. The reparameterization trick and the special loss can be implemented in just a few lines of code.

After training the model with the given formulation, we obtain the optimal parameters μ^*, Σ^* for our Gaussian approximation. We then compute the BMA based on the learned approximate posterior as

$$
\begin{aligned}
P(y \mid x, \mathcal{D}) &= \int P(y \mid x, \theta) P(\theta \mid \mathcal{D}) \, d\theta \\
&\approx \int P(y \mid x, \theta) Q_{\mu^*, \Sigma^*}(\theta) \, d\theta \\
&= \int P(y \mid x, \theta) \mathcal{N}(\theta \mid \mu^*, \Sigma^*) \, d\theta \\
&\stackrel{\text{MC}}{\approx} \frac{1}{K} \sum_{k=1}^{K} P(y \mid x, \theta^{(k)})
\end{aligned}
$$

where $\theta^{(k)} \sim \mathcal{N}(\mu^*, \Sigma^*)$.

Note: All models with high mass in our posterior make sense. This is a huge statement, as we can sample infinitely many models, meaning we have an entire nice *region* of models in the parameter space. At test time, BBB works the same as ensembles. However, we always draw a new set of θs, and we truly integrate over all θs of our approximate posterior. In contrast, the θs in ensembles are fixed.

Gaussian posterior approximations are restrictive... Why is this better than ensembles?

Gaussian posterior approximation is a different way to model the posterior than the sum of Diracs. The training procedure and the form of the approximation (meaning of the posterior space) are different.

- **Pro**: We can think about confidence intervals for the approximate posterior (in such a high-D space still; one number for each weight). It is also meaningful to ensure that a certain area around μ is always a solution. If we care about getting a region in the parameter space where everything is a solution, then this has more edge. This is a huge volume (a subset of a million-D space) compared to ensembles (that have no volume, as they are just points).
- **Contra**: We have to specify an explicit prior with which the problem remains tractable.

In general, this is not a *better* solution than the ensemble method. The ensemble method does not give a variational approximation and basically samples from the true posterior w.r.t. the weight initialization prior.

An overview of training an infinite number of models

One possible recipe for training an infinite number of models is as follows.

1. Sample $\theta^{(k)} \sim \mathcal{N}(\mu, \Sigma)$ and train that model (μ, Σ) with the likelihood and the KL term. This is training an infinite number of models represented by $\mathcal{N}(\mu, \Sigma)$ at once.
2. After training, we trust $\mathcal{N}(\mu^*, \Sigma^*)$ (the approximate posterior) to represent a good set of plausible models (of infinite cardinality) that generally work well for the training data. If those models disagree on some sample x (i.e., $K^{-1} \sum_k P(y \mid x, \theta^{(k)})$ has high uncertainty), then x is likely to be alien to these models. (The epistemic uncertainty is high, the sample is likely to be OOD and from an unseen region.)

Additional Information: Regularization Term

Why is there no λ term in the BBB objective formulation to balance the effect of the two terms? We can do it, it is a more general formulation. Here we did not augment the derived formula with any hyperparameters. However, this effect is already controlled by the prior variance to some extent (not exactly, check the effect in the KL term). This is a nice gain from the probabilistic formulation.

Additional Information: Choosing a Diagonal Σ in BBB

Choosing a diagonal covariance matrix for our Gaussian approximation can be questionable if our true posterior is elongated in some directions. This is illustrated in Figure 4.27. We generally cannot know whether this will happen in advance without extensive investigation in random directions. We force our variational posterior under the true posterior because we use the *reverse KL divergence* as our objective (true posterior is the second argument of KL), which "squishes" our approximate posterior into regions of the true posterior with high density. With the *forward KL divergence* (approximate posterior is the second argument of KL), the exact opposite happens: we want to have high density with our approximate posterior wherever the true posterior has high density.
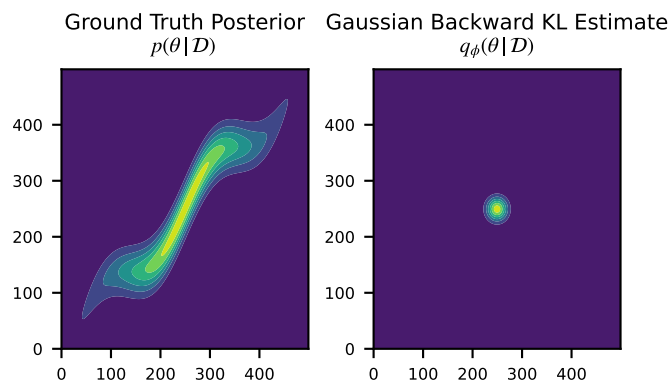
Figure 4.27: Gaussian posterior approximation using the reverse Kullback-Leibler divergence. The resulting approximate posterior only fits a small high-density region of the true posterior.

Additional Information: BBB looks just like VAEs. What is different?

These are called variational because they use variational inference. They consider a complex posterior (over an unobserved (not a training sample) variable θ or z, intermediate latent values) and use a tractable family Q to approximate it. In VAEs we want to maximize $P(x \mid \theta)$, and we approximate the intractable $P(z \mid x, \theta)$ with $Q_\phi(z \mid x) = \mathcal{N}(z \mid \eta_\phi(x), \Lambda_\phi(x))$. The general idea is the same: estimate the posterior over the latent variable given the training data. Build a KL distance between the true posterior and the approximate posterior. (Eventually, we get a KL term against the prior plus a data term.

We train the parameters of the approximate posterior. In VAEs [100], we also train the decoder θ, not just the encoder ϕ.) However, in VAEs, we optimize the ELBO (evidence lower bound); here, we optimize an MC approximation of the true objective. Another difference: In BBB, we are performing variational inference over the parameters we have to train, but in VAEs, we are doing inference over the intermediate outputs z, not the parameters. In VAEs, we use MLE to learn the parameters θ and ϕ. In BBB, our entire problem is about the variational approximation of the true posterior.

Variational approximations happen in many contexts. As an analogy, we can train a DNN for some problem, and it is always the same story…Well, it is, but several things are different.

VAEs are from ICLR 2014. BBB is from ICML 2015. BBB actually cites the VAE paper.

4.11.7   Weight Space

We recommend looking at loss landscape visualizations. Imagining these when discussing Bayesian ML and loss landscapes, in general, makes the topics a lot easier to interpret. It is nice to share these visualizations in our heads. In particular, we refer to two videos, The Loss landscape and Loss Landscape Explorer 1.1. We can see a visualization of traveling on the loss landscape to a local minimum. The latter video shows the loss landscape explorer 1.1, which can explore the loss landscape live on real data.

The way these visualizations are created is discussed in the FAQ session of the webpage of the authors.

"How to deal with so many dimensions? It is very challenging to visualize a very large number of dimensions. If we want to understand the shape of the loss landscapes, somehow we need to reduce the number of dimensions. One of the ways in which we can do that is by using a couple of random directions in space, random vectors that have the same size of our weight vectors. Those two random directions compose a plane. And that plane slices through the multidimensional space to reveal its structure in 2 dimensions. If we then add a 3rd vertical dimension, the loss value at each point in that plane, we can then visualize the structure of the landscape in our familiar 3 dimensions. (Visualizing the Loss Landscape of Neural Nets, Li et al.)" [126]

It is easy to make such landscape visualizations look nice by "cheating": One can pick random directions until they get something that is visually appealing, then report only these as cherry-picked results. Of course, this is academic malpractice, but its possibility should always be considered, especially when reviewing novel works.

4.11.8   Training a Curve of an Infinite Number of Models

Based on loss landscape visualizations, we can find creative new ways to train an infinite number of models. For example, we can parameterize a curve [57] between two trained models in the parameter space. This is illustrated in Figure 4.28. Previously, we fit a Gaussian around a point in the parameter space (BBB). We might also think about training more global connections between two, possibly faraway points in the parameter space. We can get an x-y
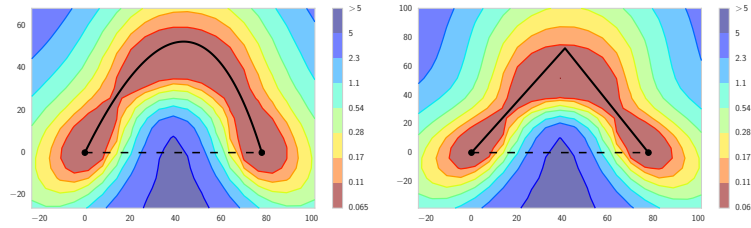
Figure 4.28: Training loss surface of a Resnet-164 model on CIFAR-100. *All* models on the obtained curves have low training loss. The training loss is an L_2-regularized CE loss. Two different parametric curves are shown after optimization. Figure taken from [57].
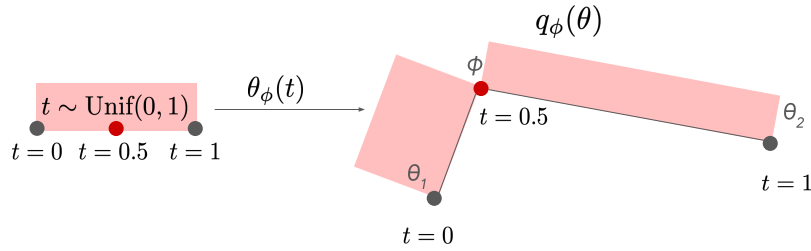
Figure 4.29: Piecewise uniform distribution over a piecewise linear curve, treated as the approximate posterior $Q_\phi(\theta)$.

cut of the parameter space, where the training loss values are indicated by colors. Here, x and y indicate two selected axes from the parameter space. They are determined by the third trained point of the curve.

Here, one trains not just a single θ, but a continuous set of θs along the curve. [26] On the plots, we see an infinite number of models that perform well on the training set. Along the curves, the training loss is always very low. Perhaps all points along the curve are good solutions for the training set. This is also Bayesian, as we are training an infinite number of models according to a learned parametric approximate posterior.

We train the curve above as follows. (This is the same story as before.)

1. Train two independent models θ_1 and θ_2 on different seeds. They are fixed throughout and treated as constants. They are the endpoints of our parametric curve.
2. We parameterize a curve via a third model ϕ. We define the curve via the line segments $\theta_1 - \phi$ and $\phi - \theta_2$.

$$\theta_\phi(t) = \begin{cases} 2(t\phi + (0.5 - t)\theta_1) & \text{if } t \in [0, 0.5) \\ 2((t - 0.5)\theta_2 + (1 - t)\phi) & \text{if } t \in [0.5, 1] \end{cases},$$

which is a bijection between $\theta_\phi(t)$ and t (if $\theta_1 \neq \theta_2$).
3. We model a piecewise uniform distribution over our parametric curve embedded in a high-D space. This is shown in Figure 4.29.[27]
4. At each iteration, sample one parameter from the piecewise uniform distribution over the curve at a time. We sample $t \sim \text{Unif}[0, 1]$; then, $\theta_\phi(t)$ is a sample of the approximate posterior $Q_\phi(\theta)$.

---

[26]This is also of measure 0, just like the ensemble posterior approximation.
[27]The two line segments are equally probable. Therefore, the piecewise density values differ when ϕ is not equidistant to the two parameters.

5. ϕ is optimized such that any model on the curve has low training loss. The trained curve is supposed to be a subset of the solution set of the training loss function. We use the reparameterization trick to separate randomness and backprop to the parameters of the distribution that describe the curve.

The optimization problem is as follows (which is typical BNN training).

$$\min_{\theta} \mathbb{E}_{\theta \sim Q_\phi(\theta)} \left[-\frac{1}{N} \sum_n \log P(y_n \mid x_n, \theta) \right].$$

The objective function can be rewritten as follows using the reparameterization trick:

$$\mathbb{E}_{t \sim \mathrm{Unif}[0,1]} \left[-\frac{1}{N} \sum_n \log P(y_n \mid x_n, \theta_\phi(t)) \right]$$

with the curve defined above. The curve is piecewise linear w.r.t. t (not differentiable at $t = 0.5$) but is entirely linear w.r.t. ϕ (t is just a fixed parameter then), so it is differentiable everywhere.

It is also easy to see that the procedure is differentiable after selecting t, as that is the only source of randomness. Sampling t and obtaining the actual parameter $\theta_\phi(t)$ are well separated by design. We do not have to use the reparameterization trick, it is "already used".
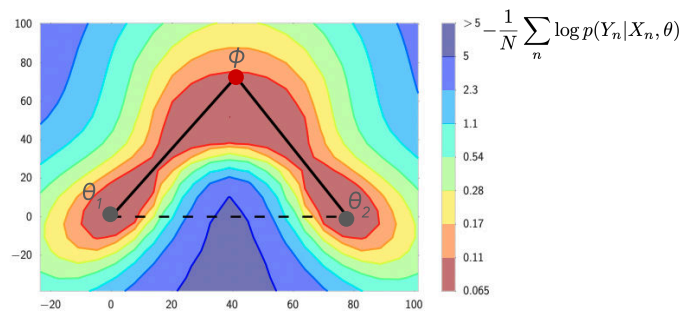
Figure 4.30: Sampling uniformly from the curve during training ensures that all models on the curve have a low training loss. The 2D training loss surface slice is plotted in which the parameterized curve resides. The authors argue that the parameters in the middle of the curve tend to *generalize* better than the endpoints (i.e., their test loss is lower than those of the endpoints, for being embedded in the middle of a wider basin of the loss surface). Base figure taken from [57].

The training ensures that every point θ on the surface has low training loss. This is empirically verified in Figure 4.30, where the training loss values are plotted. All losses are below ≈ 0.11 on the curve. A general observation is that almost all pairs of independently trained models (θ_1, θ_2) for DNNs are connected through a third point ϕ in a low-loss "highway" that we can easily find. This gives an interesting intuition for the loss landscape: Most solutions in the DL landscape are connected by some piecewise linear curve. This is not so surprising: We have millions/billions of dimensions to choose from. We can likely find a 2D cut of the loss in which there exists a parametric curve parameterized by ϕ that connects the two endpoints with a low training loss.

The NN has a vast capacity (many dimensions) to accommodate an infinite number of solutions globally rather than around a certain point. Previously, we have shown that it is possible

to train an infinite set of models around a specific point locally (Gaussian posterior). Here, we are expanding that idea to global traversal of the parameter space. This was the first work that showed that it is possible.

After training the model with the given formulation, we compute the BMA based on the learned approximate posterior as

$$P(y \mid x, \mathcal{D}) = \int P(y \mid x, \theta) P(\theta \mid \mathcal{D}) \, d\theta$$

$$\approx \int P(y \mid x, \theta) Q_\phi(\theta) \, d\theta$$

$$\overset{\text{MC}}{\approx} \frac{1}{K} \sum_{k=1}^{K} P(y \mid x, \theta_\phi(t^{(k)}))$$

where $t \sim \mathrm{Unif}[0, 1]$.

Figure 4.31: Mode connectivity visualization. The direct path between the two local minima contains high-loss parameter configurations as well. However, we can find a line connecting the two where *all* configurations result in low loss. Figure taken from [126].

A visualization of mode connectivity is given in Figure 4.31. We have two solutions that can be connected by some curve in the parameter space. On the curve, *test accuracy is also nearly constant*. This has strong implications for generalization.

Additional Information: Is this method Bayesian?

This is not a purely Bayesian approach: The true posterior $P(\theta \mid \mathcal{D})$ is approximated by the density $Q_\phi(\theta)$, which is obtained by maximum likelihood (we maximize the likelihood of the dataset w.r.t. ϕ). So, we do not consider any prior beliefs over the parameters, and indeed it is probably unlikely that the posterior would be anything close to being a

$$-\log p(Y_{\mathrm{train}}|X_{\mathrm{train}},\theta) \qquad -\log p(Y_{\mathrm{color}}|X_{\mathrm{test}},\theta) \qquad -\log p(Y_{\mathrm{orient}}|X_{\mathrm{test}},\theta)$$
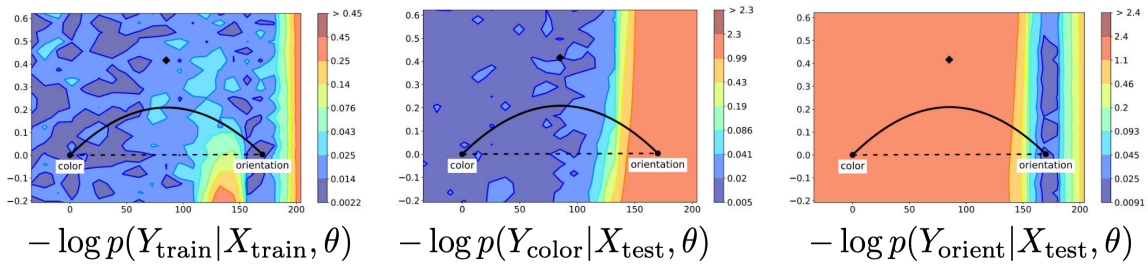
Figure 4.32: Negative log-likelihood of the diagonal training set and unbiased test set with different labels. *Left.* On the training set, the found curve of models has a low loss. One of the endpoints is a color-biased model, the other is orientation-biased. Therefore, one can obtain a curve of models that interpolates between two models with different biases. *Middle.* When considering a test set with color labels, the color-biased endpoint performs much better, as expected. However, there are many models on the curve that also perform well. There is a relatively quick shift between color-biased and orientation-biased models on the curve. There are also more color-biased models (as the blue area is larger). *Right.* On the same test set with orientation labels, the orientation-biased endpoint performs well, and also a small region of models on the curve (corresponding to the blue region). Base figure taken from [173].

curve if we chose our prior as something like a Gaussian. This work does not care about the prior. It samples some initial models, trains them, fits the curve, and treats it as a posterior approximation. It still has many nice properties and allows interesting insights into the parameter space and the loss surface.

However, Bayesian in this book refers to training infinitely many models, not performing Bayesian inference using the prior + likelihood formulation. For a true Bayesian, the prior matters a lot. For the purpose of this book, it does not. We also see ensembling as a Bayesian method. All we are doing is approximating the otherwise intractable true posterior in various ways, sometimes taking an explicit prior into account, sometimes not. This is a common interpretation in the field, and is hard to connect it to any rigorous Bayesian theory.

We are **not** using a variational approximation of the true posterior: We do not have an explicit prior, and the training objective also does not take any prior into account, as we are performing maximum likelihood estimation over the third parameter of the curve.

Additional Information: Further Surprise

We can find further surprises [173] when considering models biased to different cues. This is shown in Figure 4.32. Even heterogeneous pairs of models (θ_1, θ_2) can be connected with some curve, where heterogeneous means that the two solutions are biased to different attributes.

In the training data, color and orientation labels coincide; we have a diagonal dataset (Section 2.7.1). We can use either of the cues to get low training loss. Here, θ_{color} refers to a model biased to color (i.e., the usual solution we get), and $\theta_{\mathrm{orientation}}$ corresponds to a model biased to orientation (which is an unusual solution). **Note**: We have two sets

of inputs: X_{train} and X_{test}. However, for X_{test}, we consider two labeling schemes: one w.r.t. color and one using orientation as the task cue. Therefore, the loss landscapes are different for these three datasets in total. The parameter x-y cut is shared across these datasets.

The axes are chosen as follows. The starting point is two models with different biases. We train a third model using the formulation above. We obtain three points in a million-D space that determine a unique plane (2D subspace) that contains all three models. The other dimensions are hidden in the plots. The negative log-likelihood is plotted for all models (parameter configurations) in this plane.

On the left, it is possible to connect these very different solutions with a curve on the training set landscape: The loss for the training data is very low for the entire curve of models, as for the training set, it does not matter which cue our model chooses.

In the middle, when we consider the color test set and the *same* curve, we have many models with a low loss, i.e., many models on the curve are biased towards color (as the blue region is rather large). However, an entire region of models that had low training loss suddenly has high test loss (right, orange part of the middle plot): This shows that these models learned spurious correlations w.r.t. the color labeling scheme.

On the right, when we consider the orientation test set (i.e., we only change the labels compared to the middle plot) and the *same* curve, we have a lot fewer models on the curve with a low loss, i.e., models that are biased towards orientation: The blue (low-loss) region is rather small. The yellow region shows a transition from color-biased models to orientation-biased models, and all color-based models have a high loss (red region) on the orientation task.

It is nice to see that the space of color-biased solutions is much larger than that of the orientation-biased solutions. This probably explains why if we simply train a model, it is more likely to get a color-biased solution than solutions biased to other cues. This is a volumetric POV for explaining why color is a more favored bias for the models than other cues.

Another example: Frogs being the foreground cue and swamps being the background cue. The training samples consist of frogs in swamps. The middle plot would then correspond to pictures of swamps that do not necessarily contain a frog (unbiased dataset). Here, looking at the foreground does not solve the problem. Many more models are biased toward the background than the foreground.

4.11.9 Stochastic Weight Averaging

We can exploit the randomness in SGD. This is another cheap source of Bayesian ML. The method is called Stochastic Weight Averaging [132] (SWA). An informal overview is given in Figure 4.33. In SGD, we usually use an LR schedule. When the LR is sufficiently reduced, solutions are not moving too much around a certain point in space. We treat the set of points (models) towards the end of training (i.e., when the model roughly converged and the models are indeed plausible under the data) as samples from some Gaussian (see Figure 4.34). This

is the whole idea behind SWA with a Gaussian (SWAG).
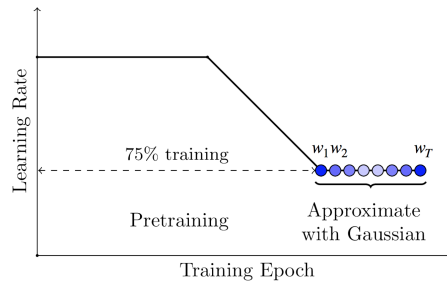
Figure 4.33: Informal overview of Stochastic Weight Averaging. We give an approximate posterior by considering parameter configurations from the later 25% of the epochs. Not all of these models have necessarily converged. Figure taken from [210].

SGD thus inherently trains a large number of models. The SGD trajectory is noisy because of the small batches. The training's final few iterations (epochs) can be treated as samples from the approximate posterior distribution. This is the MCMC point of view of training and sampling, first introduced in "Bayesian Learning via Stochastic Gradient Langevin Dynamics" [205], published way before this paper, in 2011. SGLD is an MCMC method to train and sample a posterior. The intuition is the same as what we discuss here. We treat the current procedure as if we were MCMC sampling the posterior (because of the noise from SGD) that is determined by the loss landscape. (If we do not regularize, we only have an uninformative prior, and the loss landscape is the negative log-likelihood.)

A visualization of SWAG is shown in Figure 4.34. Based on the mean and variance of the models of the last couple of epochs, we give a Gaussian approximation. Now, we do not use a variational approximation to get this Gaussian: We do not minimize KL divergences.
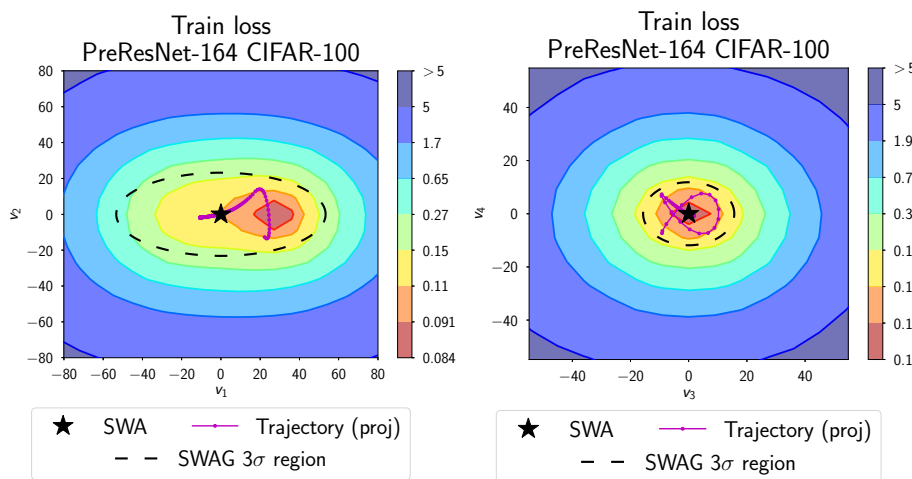
Figure 4.34: "**[Left]:** Posterior joint density surface in the plane spanned by eigenvectors of SWAG covariance matrix corresponding to the first and second largest eigenvalues and **Right:** the third and fourth largest eigenvalues. All plots are produced using PreResNet-164 on CIFAR-100. The SWAG distribution projected onto these directions fits the geometry of the posterior density remarkably well." [132] Figure taken from [132].

The method's assumption is that the posterior is approximately Gaussian:[28]

$$Q(\theta) \approx P(\theta \mid \mathcal{D})$$
$$Q(\theta) = \mathcal{N}(\theta \mid \mu(\mathcal{D}), \Sigma(\mathcal{D})).$$

The Gaussian parameters are computed from the parameters of the last L epochs (= iterations): $\theta_1 \ldots, \theta_L$.

$$\mu(\mathcal{D}) = \frac{1}{L} \sum_l \theta_l$$

$$\Sigma(\mathcal{D}) = \frac{1}{L} \sum_l \theta_l \theta_l^\top - \left(\frac{1}{L} \sum_l \theta_l \right) \left(\frac{1}{L} \sum_l \theta_l \right)^\top.$$

SWAG is not so scalable.

The problem with the above formulation is the full empirical covariance matrix: For a mid-sized network with a few million parameters, computing and storing this matrix becomes infeasible.

SWAG-Diag uses a diagonal approximation of SWAG. The only difference is how the covariance matrix is approximated. As expected from its name, SWAG-Diag uses a diagonal approximation:

$$\Sigma(\mathcal{D}) = \operatorname{diag}\left(\frac{1}{L} \sum_l \theta_l^2 - \left(\frac{1}{L} \sum_l \theta_l \right)^2 \right),$$

where the squaring operations are element-wise.

When using SWAG-Diag, we do not need to train M different models (like in an ensemble setup), nor do we need to calculate a full covariance approximation (like vanilla SWAG). We only need to train normally and give a Gaussian approximation based on the last few epochs. This is very easy and comes at almost no cost. We can, e.g., do it on ImageNet, or could do it for ChatGPT too. A comparison of SAWG and SWA with other methods is shown in Figure 4.35. SGD corresponds to the standard training of a single model. Unlike a reliability diagram, the plots directly show the deviation from the line of perfect calibration. Therefore, closer to 0 is better. SWAG-Diag is sadly very similar to SGD regarding the reliability diagram – SWAG is way better than SWAG-diag, even on ImageNet. It seems that SWAG-diag only scales better computationally, but the results do not follow.

4.11.10   On the Principledness of Bayesian Approaches

Bayesian approaches look principled. They *are* principled, given that lots of assumptions are actually true:

- We have a sensible prior that does not make learning infeasible (e.g., the true model (parameter configuration) is outside the support of the prior) or inefficient (e.g., the true model is in the tail, so we need a huge dataset to have high mass at the true model in the approximate posterior).

---

[28]The corresponding priors (of multiple experiments) are specified in [132], e.g., L_2 regularization.
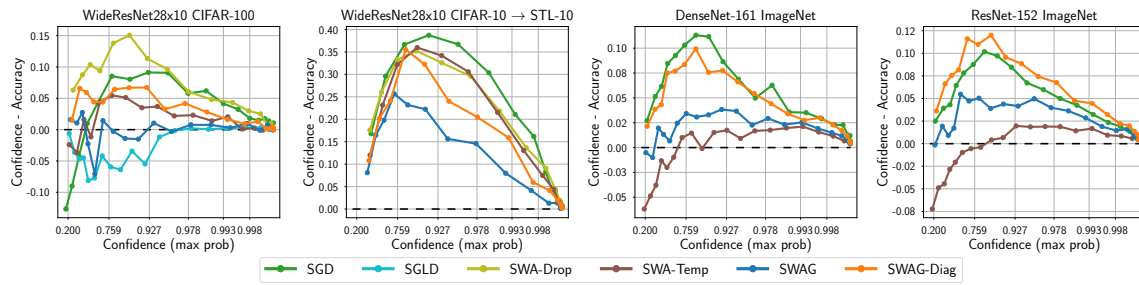
Figure 4.35: "Reliability diagrams for WideResNet28x10 on CIFAR-100 and transfer task; ResNet-152 and DenseNet-161 on ImageNet. Confidence is the value of the max softmax output. [...] SWAG is able to substantially improve calibration over standard training (SGD), as well as SWA. Additionally, SWAG significantly outperforms temperature scaling for transfer learning (CIFAR-10 to STL), where the target data are not from the same distribution as the training data." [132]. Figure taken from [132].

- The posterior follows the assumed distribution (e.g., a Gaussian). Of course, the posterior will seldom be truly Gaussian. This is a huge assumption.[29]

In high-dimensional parameter spaces (millions/billions), it is challenging to guarantee those criteria. To ensure that our posterior is concentrated around the true model, we need many samples (which is a foundational problem, not a shortcoming of approximations). To recover the true posterior, we need it to be in the approximate family. Even to *verify* correctness, we would need many samples from the true posterior (an exponentially scaling number in the number of dimensions), especially for complex distributions. This is infeasible for deep learning.

4.12 Non-Bayesian Approaches to Epistemic Uncertainty: Measuring Distances in the Feature Space

We have seen that we can give epistemic uncertainty estimates (by measuring, e.g., the variance of the predictions) when training an infinite (or large) number of models. In principle, however, we do not require training an infinite number of models. Let us remember our basic requirement for epistemic uncertainty: $c(x)$ is expected to be low when x is away from seen examples (OOD). Hence, we can also try to estimate epistemic uncertainty by measuring the distance between test sample x and training samples in the feature space.

4.12.1 Mahalanobis Distance

Let us discuss "A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks" [119]. We want to measure the closeness of a test sample x to one of the classes in the feature space for OOD detection. To give a feature representation to each class, we consider the training samples in the feature space (e.g., the penultimate layer of DNNs, with dimensionality ≈ 1000). This is illustrated in Figure 4.36.

In Figure 4.36, we computed the distance of our test sample to all training samples. As we

---

[29]Bayesians usually claim that at least they are open with their assumptions. Frequentists *also use priors*, but implicitly, which makes them less principled in the Bayesian sense.
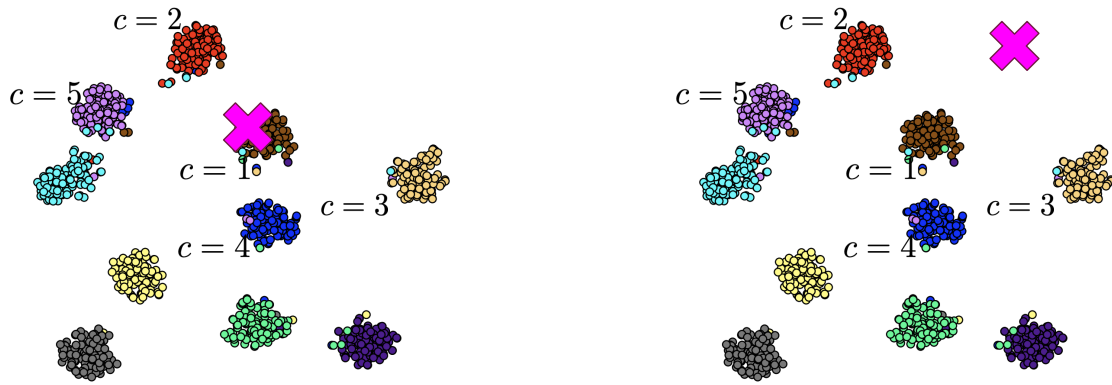
Figure 4.36: Feature space representation of training samples, where classes are encoded by color *Left cross.* Test features are close to training sample features of class 1 (one of the clusters). This is an in-distribution (ID) test sample. *Right cross.* Test features are not close to training sample features of any class (neither of the clusters). This is an OOD test sample. Base figure taken from [119].

do not want to keep all training sample features for future reference, we compute the mean and covariance for each class in the feature space based on all samples in the training set, then approximate each class by a heteroscedastic, non-diagonal Gaussian:

$$\mu_k = \frac{1}{N_k} \sum_{i:y_i=k} f(x_i)$$

$$\Sigma_k = \frac{1}{N_k} \sum_{i:y_i=k} (f(x_i) - \mu_k)(f(x_i) - \mu_k)^\top.$$

The authors of [119] also consider "tied cov" to simplify computations by unifying the covariance across classes:

$$\Sigma = \frac{1}{N} \sum_k N_k \Sigma_k$$

$$= \frac{1}{N} \sum_k \sum_{i:y_i=k} (f(x_i) - \mu_k)(f(x_i) - \mu_k)^\top.$$

This is *not* the same as calculating the covariance matrix for the entire dataset, as the individual class means are preserved. This is the weighted average of all covariance matrices where the weights are N_k/N. Every class has a different number of samples. We then measure the **Mahalanobis distance** between a test sample x and the Gaussian for class k as

$$M(x, k) = (f(x) - \mu_k)^\top \Sigma^{-1} (f(x) - \mu_k).$$

Additional Information: Interpretations of the Mahalanobis Distance

There are two ways to think about the Mahalanobis distance. (1) It is roughly the NLL of the test sample given class k (up to a constant). (2) It is the L_2 distance between sample

---

> x and the class mean, weighting every dimension by the precision (inverse covariance) matrix. This is a distorted L_2 distance that weights directions with large precision (small variance) more. Both interpretations are useful to keep in mind.

Then we define the confidence measure $c(x)$ based on the smallest Mahalanobis distance to a Gaussian:

$$c(x) := -\min_c M(x, c).$$

According to this definition, $c(x)$ is low when x is OOD (i.e., $\min_k M(x, k)$ is high) and analogously $c(x)$ is high when x is ID (i.e., $\min_k M(x, k)$ is low).

Note: This method was designed for detecting OOD samples. It did not consider aleatoric uncertainty, which is also present in the estimates. The field is now aware that it can also influence predictive uncertainty. One could, e.g., measure aleatoric uncertainty as the ratio of distances of the two closest distributions. When it is approximately one, we have high aleatoric uncertainty, because we are split between the classes.

Results of Mahalanobis Distance

To discuss the Mahalanobis distance's ability to detect OOD samples, we consider Table 4.4. This showcases the detection accuracy of OOD samples under various metrics. AUPR in and out correspond to whether the ID or OOD class is considered positive. The max-prob confidence measure is not suitable for OOD detection as much. The performance in detecting OOD samples is considerably better for the Mahalanobis distance in the feature space. This is, of course, just an example. One should not conclude that the Mahalanobis distance is always a better confidence measure than max-prob, just for this particular setup.

Table 4.4: Comparison of max-probability and Mahalanobis distance for distinguishing OOD samples. The Mahalanobis distance gives better results on OOD detection. Detection performance is measured using a ResNet trained on CIFAR-10. The OOD dataset is SVHN. For the Mahalanobis distance method, no feature ensembling and input pre-processing is used. Table adapted from [119].

$c(x)$	TNR at TPR 95%	AUROC	Detection accuracy	AUPR in	AUPR out
Max-probability	32.47	89.88	85.06	85.40	93.96
Mahalanobis	54.51	93.92	89.13	91.56	95.95

We also consider a more fine-grained evaluation of the Mahalanobis distance's ability to detect OOD samples in Table 4.5. Here, various distribution changes are considered where the confidence measure correctly has to tell ID and OOD apart. The baseline is, again, max-prob. The superiority of the Mahalanobis distance in this OOD detection setup is clear. **Note**: Generally, one needs a good set of in-dist datasets and a diverse set of OOD datasets to measure OOD detection performance.

Table 4.5: Accuracies expressed in percentages on distinguishing OOD data under various distribution changes. The Mahalanobis distance consistently outperforms the max-prob confidence measure. The detection accuracies are reported for the OOD samples. Table adapted from [119].

In-dist (model)	OOD	Validation on OOD samples		
		TNR at TPR 95%	AUROC	Detection acc.
		Max-prob / Mahalanobis		
CIFAR-10 (DenseNet)	SVHN	40.2 / **90.8**	89.9 / **98.1**	83.2 / **93.9**
	TinyImageNet	58.9 / **95.0**	94.1 / **98.8**	88.5 / **95.0**
	LSUN	66.6 / **97.2**	95.4 / **99.3**	90.3 / **96.3**
CIFAR-100 (DenseNet)	SVHN	26.7 / **82.5**	82.7 / **97.2**	75.6 / **91.5**
	TinyImageNet	17.6 / **86.6**	71.7 / **97.4**	65.7 / **92.2**
	LSUN	16.7 / **91.4**	70.8 / **98.0**	64.9 / **93.9**
SVHN (DenseNet)	CIFAR-10	69.3 / **96.8**	91.9 / **98.9**	86.6 / **95.9**
	TinyImageNet	79.8 / **99.9**	94.8 / **99.9**	90.2 / **98.9**
	LSUN	77.1 / **100**	94.1 / **99.9**	89.1 / **99.3**
CIFAR-10 (ResNet)	SVHN	32.5 / **96.4**	89.9 / **99.1**	85.1 / **95.8**
	TinyImageNet	44.7 / **97.1**	91.0 / **99.5**	85.1 / **96.3**
	LSUN	45.4 / **98.9**	91.0 / **99.7**	85.3 / **97.7**
CIFAR-100 (ResNet)	SVHN	20.3 / **91.9**	79.5 / **98.4**	73.2 / **93.7**
	TinyImageNet	20.4 / **90.9**	77.2 / **98.2**	70.8 / **93.3**
	LSUN	18.8 / **90.9**	75.8 / **98.2**	69.9 / **93.5**
SVHN (ResNet)	CIFAR-10	78.3 / **98.4**	92.9 / **99.3**	90.0 / **96.9**
	TinyImageNet	79.0 / **99.9**	93.5 / **99.9**	90.4 / **99.1**
	LSUN	74.3 / **99.9**	91.6 / **99.9**	89.0 / **99.5**

4.12.2 Other types of distances than Mahalanobis: RBF kernel

We discuss the work "Uncertainty Estimation Using a Single Deep Deterministic Neural Network" [9] to highlight another distance-based uncertainty estimator.

Here, instead of computing the Mahalanobis distance w.r.t. the class Gaussians, we first compute the L_2 distance between test sample and centroid of class k, $\mu_k = \frac{1}{N_k} \sum_{i:y_i=k} f(x_i)$:

$$d(x,k) = \|f(x) - \mu_k\|_2^2.$$

Then, we compute the RBF kernel value for class k as

$$K_k(f(x), \mu_k) = \exp\left(-\frac{d(x,k)}{2\sigma^2}\right)$$

where σ is a hyperparameter. This is a special case of the Mahalanobis distance where the covariance is isotropic (hence the name "radial" basis function), and we take the squared L_2 norm.

The kernel value has a nice property:

$$K_k(f(x), \mu_k) \in (0,1]$$

where higher values indicate greater similarity. This is more interpretable than the Mahalanobis distance, and it also has a nice interpretation as a probability. All σ does is to control the temperature of this distribution. Finally, we define our confidence level as

$$c(x) := \max_k K_k(f(x), \mu_k).$$

Low confidence indicates an OOD sample: $c(x)$ can be interpreted as the probability of x *not* being OOD. Conveniently, $c(x) \in (0, 1]$, thus, we can apply a proper scoring rule and train using the resulting criterion. In the derivation below, we exclude the case of $K_k(f(x), \mu_k) = 1$ and also simplify notation to just K_k. The negative log probability scoring rule for the max-RBF similarity is given by

$$\mathcal{L} = \begin{cases} -\log \max_k K_k & \text{if } Y_{\arg\max_k K_k} = 1 \\ -\log(1 - \max_k K_k) & \text{if } Y_{\arg\max_k K_k} = 0 \end{cases} \tag{4.1}$$

where Y is a one-hot (random) vector for the GT class. As Y is a one-hot vector, $Y_{\arg\max_k K_k} = 1$ means that the prediction is *correct*, whereas $Y_{\arg\max_k K_k} = 0$ shows an *incorrect* prediction.

When the prediction is correct, we gain $\log c(x)$ reward (or lose $-\log c(x)$ reward). If we were very confident, we would gain the most. This encourages the network to have a high $c(x)$, i.e., make the feature representation of X even closer to the current centroid. We are optimizing correct predictive uncertainty estimation by $c(x)$. When the prediction is incorrect, we gain $\log(1 - c(x))$ reward. We repel the current centroid.

We upper bound Equation 4.1 with a familiar loss function, BCE. When $Y_{\arg\max_c K_c} = 1$ (upper branch), we write

$$-\log\max_k K_k = -\sum_k Y_k \log K_k$$

because Y is a one-hot vector. We also have

$$-\log(1 - \max_k K_k) = -\log\min_k(1 - K_k) = \max_k -\log(1 - K_k)$$

where we used for the last equality that $\log$ is monotonically increasing. When $Y_{\arg\max_k K_k} = 0$ (lower branch), this can be bounded from above as

$$\max_k -\log(\underbrace{\underbrace{1 - K_k}_{\in(0,1)}}_{\in(0,+\infty)}) \leq \sum_{k:Y_k=0} -\log(1 - K_k)$$

$$= \sum_k -(1 - Y_k)\log(1 - K_k).$$

Thus, we finally have that

$$\mathcal{L} \leq \begin{cases} -\sum_k \overbrace{Y_k \log K_k}^{>0} & \text{if } Y_{\arg\max_k K_k} = 1 \\ \underbrace{-\sum_k (1 - Y_k)\log(1 - K_k)}_{>0} & \text{if } Y_{\arg\max_k K_k} = 0 \end{cases}$$

$$\leq -\sum_k \left(Y_k \log K_k + (1 - Y_k)\log(1 - K_k)\right).$$

The authors of [9] optimize this upper-bound proxy loss on a finite (deterministic) dataset $\{(x_i, y_i)\}_{i=1}^N$. The loss is the sum of BCEs of one-vs-rest classifications where K_k is our predicted probability of membership of class k for sample x. This advocates the use of this form

of BCE for optimizing our classifiers. This encourages correct predictive uncertainty ($L = 1$) reports for $c(x) := \max_k K_k(f(x), \mu_k)$.

A remaining problem is that the class centroids $\mu_k = \frac{1}{N_k} \sum_{i:y_i=k} f(x_i)$ are continuously updated during training. These are needed for all K_k and for all x. Suppose we recompute centroids every time we update our parameters. In that case, we will have a very noisy training procedure, as the targets (centroid means) are constantly moving, and we are trying to chase after them for the right class for each sample. Also, recalculating these for the entire dataset after every network update is infeasible. To solve both, we use a moving average for more stable centroid estimation at each iteration and more stable training:

$$N_k \leftarrow \gamma N_k + (1 - \gamma)n_k$$
$$m_k \leftarrow \gamma m_k + (1 - \gamma) \sum_{i \in \text{minibatch}: y_i = k} f(x_i)$$
$$\mu_k \leftarrow \frac{m_k}{N_k}.$$

where

- N_k is the "soft" number of samples per class in mini-batch: It is the moving average of the number of samples per class in mini-batch. This changes over iterations; we also need to smooth this out.
- m_k is the moving average of the sum of class k sample features.
- μ_k is the average feature location (centroid) for class k.
- n_k is the number of samples per class in the current mini-batch.
- $\gamma \in [0.99, 0.999]$ corresponds to the momentum term in the moving average. To make learning stable, it is chosen to be quite high.

Note: When a training sample has high aleatoric uncertainty, it will be positioned between likely centroids at the end of training. When a training sample has low aleatoric uncertainty, it will be very closely clustered to the correct class. When an OOD sample comes, it will have low confidence. However, we can also get low confidence for samples with high aleatoric uncertainty. We cannot distinguish these two cases based on the confidence value. This work just attributes low confidence to epistemic uncertainty.

Results of RBF Kernel

We first discuss Figure 4.37 that showcases qualitative results. The confidence estimate successfully distinguishes the two sources of data (ID, OOD). In distribution, the maximal kernel similarity[30] is very high, and samples are well clustered in the feature space. Out of distribution, samples tend to have different maximal kernel similarities than one. We qualitatively conclude that $c(x) = \max_c K_c(f(x), \mu_c)$ is a good indicator of how to separate OOD samples from ID samples after training the network. One can find the best separating threshold (or just report AUROC or AUPR).

We also discuss quantitative results shown in Table 4.6. Results are quantified using the AUROC score. ("Can we separate ID from OOD based on $c(x)$ predictions?") DUQ corresponds to deterministic uncertainty quantification using the confidence score $c(x) = \max_k K_k(f(x), \mu_k)$.

---

[30] The x-axis label reads "kernel distance" but is actually the kernel similarity. Distance is low when similarity is high.
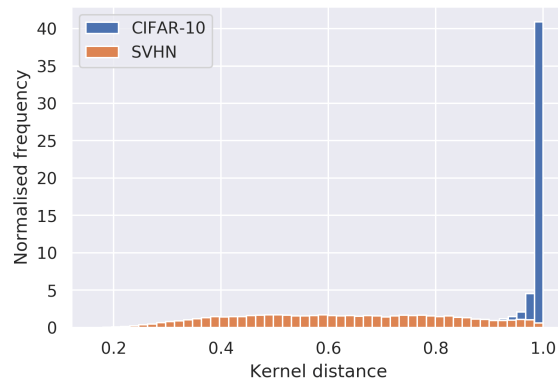
Figure 4.37: Results of RBF kernel confidence estimation. In distribution, the kernel similarities are all high, i.e., the mapped samples are concentrated in high-density regions of the feature space. Out of distribution, the similarities are mixed, meaning there are a variety of samples "the model is not familiar with." ID dataset is CIFAR-10, OOD dataset is SVHN. Figure taken from [9].

The name highlights that they do not have to stochastically train multiple (or even an infinite number of) models to obtain epistemic uncertainty estimates. LL ratio is a method we do not discuss. 'Single model' denotes DUQ trained with softmax-cross-entropy. It uses the same $c(x)$ formulation but is trained with the usual softmax-cross-entropy loss. As shown in the Table, the method gives good results after training with the proposed objective (DUQ).

Table 4.6: AUROC results on FashionMNIST, with MNIST being the OOD set. DUQ (using $c(x) = \max_k K_k(f(x), \mu_k)$) outperforms most methods."Deep Ensembles is by Lakshminarayanan et al. (2017), Mahalanobis Distance by Lee et al. (2018), LL ratio by Ren et al. (2019). Results marked by (ll) are obtained from Ren et al. (2019), (ours) is implemented using our architecture. Single model is our architecture, but trained with softmax/cross entropy." [9] Table taken from [9].

Method	AUROC
DUQ	0.955
LL ratio (generative model)	0.994
Single model	0.843
5 - Deep Ensembles (ours)	0.861
5 - Deep Ensembles (ll)	0.839
Mahalanobis Distance (ll)	0.942

4.12.3 Summary of Modeling Epistemic Uncertainty

We have seen two general ways of modeling epistemic uncertainty. In Bayesian ML, we train a set of models simultaneously. We measure their disagreement during inference through Bayesian model averaging (BMA). We can also choose to measure distances in the feature space. In particular, we can compute the distance to the closest class centroid in the feature space to get a sense of how surprising an input sample is. Both have been successfully applied to the problem of OOD detection, which is a proxy task for epistemic uncertainty.

4.13 Modeling Aleatoric Uncertainty

As we have seen, aleatoric uncertainty refers to "I do not know because there are multiple plausible answers." This happens when true label y is not a deterministic function of input x, as multiple possibilities could be an answer for input x.

Below, we give a high-level overview of the ingredients we will use to represent aleatoric uncertainty.

4.13.1 Roadmap to Representing Aleatoric Uncertainty

There are *two ingredients* that are used together for the recipe of representing aleatoric uncertainty.

Architecture. Formulate a model architecture that accommodates multiple possible outputs. We should prepare, e.g., a probabilistic output where our model outputs the parameters of this output distribution rather than a single prediction.

Loss function. Taking a proper scoring rule for matching the predicted output distribution to the one dictated by the dataset (examples are discussed in Sections 4.13.2 and 4.13.4) is often sufficient.

Let us follow our recipe and extend proper scoring rules to more generic distributions, as this will allow us to recover truthful aleatoric uncertainty estimates. We start with matching output distributions in classification.

4.13.2 Aleatoric Uncertainty In Classification

As discussed previously, aleatoric uncertainty refers to the inherent variability of the labels, i.e., the non-deterministic nature of the data generating process. To *represent* aleatoric uncertainty, we would like our model to output a *distribution* which is faithful to $P(Y \mid X = x)$.

Proper Scoring Rules to the Rescue, Again

So far, our discussion centered around binary distributions, where we tried to match a confidence value $c(x)$ to the true probability of an event, such as $P(L = 1)$, where L represents the correctness of prediction. Here, both $c(x)$ and $P(L = 1)$ corresponded to the parameters of respective Bernoulli distributions. By matching the Bernoulli parameters, we were also matching the Bernoulli distributions. To achieve this, we leveraged (strictly) proper scoring rules.

We now extend the notion of proper scoring rules to general discrete distributions. In particular, we want to match a distribution Q (a categorical distribution encoded by a vector of probabilities) to the true discrete distribution P. Let y be a sample of distribution $P(Y)$ – for example, the GT class index. We then define the scoring rule as a function $S(Q, y)$. Arguments are Q (the predicted distribution) and y (a sample from true distribution $P(Y)$). This scoring rule is *strictly proper* when the expected score $\mathbb{E}_{P(Y)}S(Q, Y)$ is maximized iff $Q \equiv P$ (i.e., when the distributions match). S may also be described as a function of Q's parameters (e.g., the

parameter vector of a categorical distribution or the parameter of a Bernoulli distribution) rather than Q itself.

If we want, we can then further compress Q into a scalar. The aleatoric confidence can be, e.g., given by the max-prob $\max_k Q(Y = k)$, or by the entropy of the predicted distribution, $\mathbb{H}(Q)$.

Let us discuss some popular proper scoring rules for matching predicted categorical distributions, encoded by softmax outputs $f(x)$, to the GT distributions with GT probabilities $P(Y = y \mid X = x) \; \forall y \in \mathcal{Y}$ from which we can only sample.

Log Probability Scoring Rule

The log probability scoring rule (negative CE) for categorical distributions is defined as

$$S(f, y) = \sum_k y_k \log f_k(x) = \log f_y(x),$$

where y is the true class.[31] It can be shown that S defined this way is a strictly proper scoring rule, i.e.

$$\mathbb{E}_{P(Y)} S(f, Y)$$

is maximal iff

$$f_k(x) = P(Y = k \mid x) \; \forall k \in \{1, \dots, C\}.$$

This is great news! Many DNNs already minimize a NLL loss of the form

$$\mathcal{L} = -\sum_k y_k \log f_k(x) = -\log f_y(x),$$

which means they are already matching their predictions to the aleatoric uncertainty of a data source. Since it is a proper scoring rule, in the expectation of Y, we encourage our DNN to predict $f(x)$ that correctly represents the spread of $P(Y \mid X = x)$ in the training set.

Multi-Class Brier Scoring Rule

To match a probability vector encoding a categorical distribution to the true distribution $P(Y \mid X = x)$, we can also use the Brier scoring rule. Consider a predicted probability vector $f(x) \in [0,1]^K$ with $\sum_{k=1}^K f_k(x) = 1$ and a categorical random variable $Y \in \{1, \dots, K\}$. The multi-class Brier scoring rule is defined as

$$S(f(x), y) = -(1 - f_y(x))^2 + f_y(x)^2 - \sum_{k=1}^K f_k(x)^2.$$

Claim 4.5. *The multi-class Brier score is a strictly proper scoring rule for aleatoric uncertainty.*

---

[31] y is often used to denote both a one-hot vector of a class and the class label. This is just an abuse of notation.

Proof. First, we rewrite $\mathbb{E}_{P(Y|X=x)}S(f,Y)$ as

$$\mathbb{E}_{P(Y|X=x)}S(f,Y) = \sum_{k=1}^{K} P(Y=k)\left[-(1-f_k(x))^2 + f_k(x)^2 - \sum_{l=1}^{K} f_l(x)^2\right]$$

$$= \sum_{k=1}^{K} P(Y=k)\left[-f_k(x)^2 + 2f_k - 1 + f_k(x)^2 - \sum_{l=1}^{K} f_l(x)^2\right]$$

$$= -\sum_{k=1}^{K}\left[P(Y=k)(1-2f_k(x)) + \sum_{l=1}^{K} P(Y=k)f_l(x)^2\right]$$

$$= -\sum_{k=1}^{K} P(Y=k)(1-2f_k(x)) - \sum_{l=1}^{K} f_l(x)^2$$

$$= -\sum_{k=1}^{K}\left[P(Y=k)(1-2f_k(x)) + f_k(x)^2\right]$$

for all $f \in \Delta^K$ which is the $(K-1)$-dimensional probability simplex.

A necessary condition for the maximizer of the Brier scoring rule in expectation is as follows. $\forall r \in \{1,\ldots,K\}, f \in \Delta^K$:

$$\frac{\partial}{\partial f_r}\left(-\sum_{k=1}^{K}\left[P(Y=k)(1-f_k(x)) + f_k(x)^2\right]\right) = -\sum_{k=1}^{K}\frac{\partial}{\partial f_r}\left[P(Y=k)(1-2f_k(x)) + f_k(x)^2\right]$$

$$= -\frac{\partial}{\partial f_r}\left[P(Y=r)(1-2f_r(x)) + f_r(x)^2\right]$$

$$= -(-2P(Y=r) + 2f_r(x)) \stackrel{!}{=} 0$$

$$\iff f_r(x) = P(Y=r).$$

As $\frac{\partial}{\partial f_r}\left(-[-2P(Y=r) + 2f_r(x)]\right) = -(0+2) = -2 < 0 \; \forall r \in \{1,\ldots,K\}, f \in \Delta^K$, the above, $f(x) \equiv P(Y \mid X = x)$ is the unique maximizer of the multi-class Brier scoring rule's expectation. Therefore, it is strictly proper. $\qquad\square$

Using softmax with the NLL Loss

Let us discuss the most popular setup for classification that uses the steps introduced in Section 4.13.1.

Using a softmax output with the NLL loss is by far the most common activation and loss function for (multi-class) classification problems. Luckily, it is also designed to handle the aleatoric uncertainty in the true $P(y \mid x)$ distribution, which is potentially multimodal (according to humans). **Note**: NLL loss = CE loss = softmax CE loss = log-likelihood loss = negative log probability for classification.

Ingredient 1. The softmax output $f(x)$ for input image x has the right dimensionality (number of classes) to represent any $P(y \mid x)$. $f(x)$ outputs the parameter vector p of the output categorical distribution. Therefore, the architectural condition is satisfied. The model is ready to represent aleatoric uncertainty.

Ingredient 2. Is the method also *encouraged* to represent the *true* aleatoric uncertainty? We consider the loss function $-\log f_Y(x)$. We have seen that, in expectation of Y, it guides the

model to produce the GT distribution $P(y \mid x)$.

Toy experiment with the NLL loss

Importantly, our DNN is *not* encouraged to be overconfident (nearly one-hot) for $\arg\max_k P(Y = k \mid X = x)$ when using the NLL loss, as many people suggest. The loss encourages the model to produce distributions $f(x)$ with variance when the true distribution also has a non-zero variance. This is, of course, considering infinite data. For finite datasets where the model usually does not see two labelings for a single data point, it can arbitrarily overfit to the given labeling for each datum (given sufficient expressivity). This makes the model *not* predict the true aleatoric uncertainty for a sample (only the empirical probability), which can result in the model being extremely overconfident in one of the possible answers. This can be possibly mitigated by adversarial training (increase region of class y prediction) or by regularizing the model based on how many data points we have.
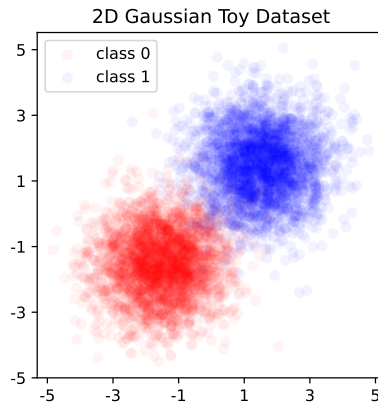
Figure 4.38: Homoscedastic 2D class-wise Gaussian dataset in a binary classification setting, used for the experiment in the notebook.

We provide a notebook to clean up the possible source of the misconception of the NLL encouraging overconfidence. The dataset is generated from two homoscedastic 2D Gaussians with a small overlap near $(0,0)$ (Figure 4.38). The task is binary classification. Since we know the Gaussians that generate each class, we can calculate the true probability that a sample x is of class 0:

$$
\begin{aligned}
P(Y = 0 \mid X = x) &= \frac{P(X = x \mid Y = 0)P(Y = 0)}{P(X = x \mid Y = 0)P(Y = 0) + P(X = x \mid Y = 1)P(Y = 1)} \\
&= \frac{P(X = x \mid Y = 0)}{P(X = x \mid Y = 0) + P(X = x \mid Y = 1)} \\
&= \frac{\mathcal{N}(x \mid \mu_0, \Sigma_0)}{\mathcal{N}(x \mid \mu_0, \Sigma_0) + \mathcal{N}(x \mid \mu_1, \Sigma_1)}
\end{aligned}
$$

where we assumed a uniform label prior. This is just the ratio of the likelihood of x being a part of class 0 and the total likelihood of it being a part of any of the two. We visualize the predicted ($f_0(x)$) and GT ($P(y \mid x)$) probabilities in Figure 4.39.

We have very high GT label certainty for the lower and upper triangles. On the diagonal, the GT label certainties are close to 0.5, signaling high aleatoric uncertainty. The question is: If
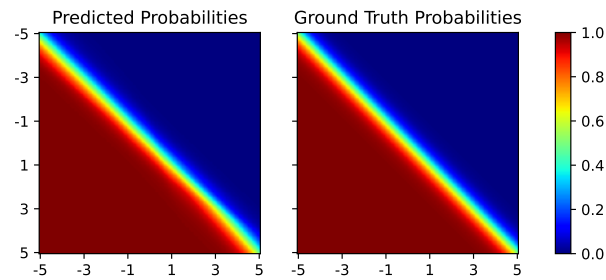
Figure 4.39: Predicted and ground truth probabilities from the experiment in the notebook. The two probability maps are almost indistinguishable.

we train with this data, does a 2-layer DNN predict something close to this after applying sigmoid? The model will observe mixed supervision near the class boundary $x_1 + x_2 = 0$. (It is also not expressive enough to overfit to the training set and produce incorrect aleatoric uncertainties. We have enough data points.) Such mixed supervision and the NLL objective result in the correct estimation of $P(y \mid x)$.

The model outputs closely resemble the true $P(Y = 0 \mid X = x)$ at nearly all x values. The model *can* learn correct aleatoric uncertainty estimation (as supported by the theory of proper scoring rules). We can see the pointwise difference between $f_0(x)$ and $P(Y = 0 \mid X = x)$ in Figure 4.40.
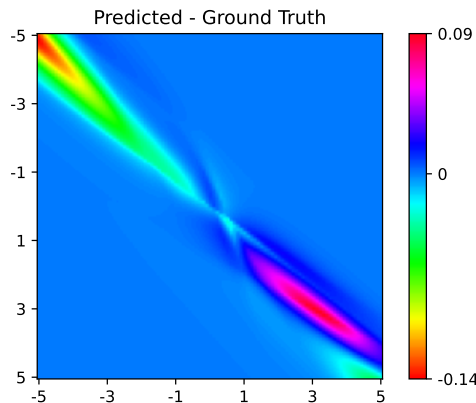
Figure 4.40: Difference of predicted and ground truth probabilities from the experiment in the notebook. The differences are mostly negligible, but there are slight deviations around the corners (where data was more sparse).

The pointwise differences are minor (note the scale). The popular softmax + CE design is already capable of handling aleatoric uncertainty. We have differences in regions with a few training samples. Here the model struggles a bit to find the correct values. This also has close ties with epistemic uncertainty.

4.13.3 Detour: Proper Scoring Rules for Aleatoric and Predictive Uncertainty

Let us now refer back to Section 4.5.6, where we established the fact that the negative cross-entropy loss is a lower bound to a strictly proper scoring rule for the correctness of prediction. We also briefly mentioned that it was, in fact, not only a lower bound: The negative cross-entropy loss also guarantees the exact recovery of the true probability of correctness, $P(L = 1)$. As the relationship between proper scoring rules for aleatoric uncertainty and predictive uncertainty is quite subtle, in this section, we wish to prove the aforementioned result and formalize this relationship.

We can just use an aleatoric uncertainty proper scoring rule for the correctness of prediction.

Informally, we will show that any strictly proper scoring rule for aleatoric uncertainty acts as a strictly proper scoring rule for the correctness of predictions as well when using max-prob confidence estimates.

Claim 4.6. *If $f(x) = \arg\max_{p \in \Delta^K} \mathbb{E}_{P(Y|X=x)} S(p, Y)$ for a strictly proper scoring rule S, then $c(x) = \max_k f_k(x) = P(L = 1)$ is the unique maximizer of any expected strictly proper scoring rule for $P(L = 1)$ for all $x \in \mathcal{X}$.*

Proof. As $S(f(x), y)$ is strictly proper for aleatoric uncertainty, the unique maximizer of the expected score is $f(x) = P(Y \mid X = x) \in \mathbb{R}^K$. We need to show that $c(x) = \max_{k \in \{1,\ldots,K\}} P(Y = k \mid X = x)$ maximizes any expected strictly proper scoring rule for the correctness of prediction, i.e., $c(x) = P(L = 1)$. To see this, one can observe that

$$P(L = 1) = P\left(Y = \arg\max_{k \in \{1,\ldots,K\}} P(Y = k \mid X = x)\right) = \max_{k \in \{1,\ldots,K\}} P(Y = k \mid X = x).$$

This means that the $c(x)$ we obtained is also a maximizer for any expected strictly proper scoring rule for the correctness of prediction, concluding our proof. □

In the above proof, we did not assume any particular strictly proper scoring rule. One could, e.g., choose the log probability scoring rule. In that case, training a model with the CE loss (NLL) minimization objective encourages both the correct prediction of the classes (original role) and the truthful report of predictive uncertainty using the max-probability confidence estimates $c(x) = \max_k f_k(x)$, and we do not only maximize a lower bound: we have strict properness. However, we could just as well choose any other strictly proper scoring rule, and the relationship would be preserved.

Sometimes, we can also use a predictive uncertainty proper scoring rule for aleatoric uncertainty.

Below, we show that the relationship also works backward in binary classification but not in multi-class classification.

Claim 4.7. *In binary classification, if $c(x) = \max(f(x), 1 - f(x))$ is the unique maximizer of any expected strictly proper scoring rule for $P(L = 1)$ (i.e., $c(x) = P(L = 1)$), then*

$$f(x) = \arg\max_{p \in [0,1]} \mathbb{E}_{P(Y|X=x)} S(p, Y)$$

for any (strictly) proper scoring rule S for aleatoric uncertainty for all $x \in \mathcal{X}$.

Proof. Let x be arbitrary. $c(x)$ being the unique maximizer of a strictly proper scoring rule for $P(L = 1 \mid X = x)$ implies $c(x) = P(L = 1 \mid X = x)$. Let us denote

$$\hat{y} = \begin{cases} 1 & \text{if } f(x) \geq 0.5 \\ 0 & \text{otherwise} \end{cases}.$$

One can observe that

$$P(L = 1 \mid X = x) = P(Y = \hat{y} \mid X = x),$$

therefore, we write

$$c(x) = \max(f(x), 1 - f(x)) = P(Y = \hat{y} \mid X = x). \tag{4.2}$$

Suppose that $\hat{y} = 1$. Then, $\max(f(x), 1 - f(x)) = f(x)$ and $P(Y = \hat{y} \mid X = x) = P(Y = 1 \mid X = x)$. Therefore, Equation 4.2 becomes

$$f(x) = P(Y = 1 \mid X = x),$$

i.e., we recover the GT probability of class 1. The GT probability of class 0 is also recovered as $1 - f(x) = 1 - P(Y = 1 \mid X = x) = P(Y = 0 \mid X = x)$.

Now, suppose that $\hat{y} = 0$. Then, $\max(f(x), 1 - f(x)) = 1 - f(x)$ and $P(Y = \hat{y} \mid X = x) = P(Y = 0 \mid X = x)$. Here, Equation 4.2 simplifies to

$$1 - f(x) = P(Y = 0 \mid X = x),$$

i.e., we recover the GT probability of class 0. The GT probability of class 1 is recovered analogously to the previous case.

Therefore, for any (strictly) proper scoring rule S for aleatoric uncertainty, we indeed obtain

$$f(x) = \arg\max_{p \in [0,1]} \mathbb{E}_{P(Y|X=x)} S(p, Y)$$

by definition. □

This does not hold for the multi-class case. Here, we can denote

$$\hat{y} = \arg\max_{k \in \{1, \ldots, K\}} f_k(x).$$

Then, as $c(x) = \max_{k \in \{1, \ldots, K\}} f_k(x)$ and $P(L = 1) = P(Y = \hat{y})$, the necessary and sufficient

criterion $c(x) = P(L = 1)$ for the maximizer of the expected score can be rewritten as

$$\max_{k \in \{1, \dots, K\}} f_k(x) = P\left(Y = \hat{y}\right). \tag{4.3}$$

Here, the probability vector $f(x)$ has K elements, but only the maximal element is matched to the GT probability of the predicted class. This can have surprising consequences that were not observed in the binary case because by matching one probability there, we matched the other one as well.

Predicting a suboptimal class (one with non-maximal GT probability) *does not* lead to a contradiction of Equation 4.3. One can convince themselves that $f(x) := (0.32, 0.33, 0.35) \in \Delta^3$ with $P(Y \mid X = x) = (0.5, 0.15, 0.35) \in \Delta^3$ satisfies the above constraint, meaning it is a maximizer of the expected score, even though the resulting prediction is incorrect. To understand this seemingly strange behavior, we note that *the maximizer* $\arg\max_{k \in \{1, \dots, K\}} f_k(x)$ *can change over the course of optimization*. In fact, under any sensible loss (e.g., cross-entropy) and a network of enough capacity, the maximizer usually becomes the GT class (at least on ID samples). Optimizing not only over the maximum probability ($\max_{k \in \{1, \dots, K\}} f_k(x)$) but also the maximizer of it ($\arg\max_{k \in \{1, \dots, K\}} f_k(x)$) results in the desired behavior of our network predicting the correct class and recovering the true $\max_{k \in \{1, \dots, K\}} P(Y = k \mid X = x)$ as the corresponding predicted probability.

Still, in the multi-class case, even though we can match the probabilities corresponding to the maximizer class of $P(Y \mid X = x)$, we get no guarantees that the remaining probabilities will also be matched. For example, $f(x) := (0.5, 0.15, 0.35) \in \Delta^3$ with $P(Y \mid X = x) = (0.5, 0.3, 0.2) \in \Delta^3$ (1) satisfy the constraint, (2) predict the correct class, but (3) the probabilities are not matched. Therefore, the equivalence of max-prob proper scoring rules for predictive uncertainty and aleatoric uncertainty only holds in the binary case. In the multi-class case, we cannot establish such a correspondence.

Wait, what about epistemic uncertainty?

One intuitive explanation for the above results is that proper scoring rules only give statements about the recovery of the true probabilities *in expectation*, or in other words, in the limit of infinitely many labels for a given input x. In this limit, epistemic uncertainty naturally vanishes for this particular input x, as we have full information about the (conditional) data generating process $P(Y \mid X = x)$. Therefore, it makes sense that proper scoring for predictive uncertainty reduces to that for aleatoric uncertainty.

4.13.4 Aleatoric Uncertainty in Regression

In regression, our true distribution $P(y \mid x)$ is continuous and has variance in all sensible settings (think of measurement noise, to begin with). Thus, Y is a continuous random variable. We now assume that the true distribution follows a normal distribution:

$$P(y \mid x) = \mathcal{N}(y \mid \mu(x), \sigma(x)^2 I).$$

The usual setting is that σ^2 does not depend on x, i.e., it is *homoscedastic*. However, that would make aleatoric uncertainty prediction uninteresting, since it would return constant

values for all x. Thus, we allow $\sigma(x)$ to vary with x, i.e., it is *heteroscedastic*. We will see that heteroscedastic regression can model this aleatoric uncertainty (when trained right).

Heteroscedastic regression can be nicely visualized in 1D as shown in Figure 4.41. Each x-slice corresponds to a Gaussian distribution with a different σ. $\sigma^2(x)$ is the variance (level of noise) of $P(y \mid x)$ that depends on x. Our goal is to correctly represent not only $\mu(x)$ (what we usually do) but also $\sigma^2(x)$ at every x.
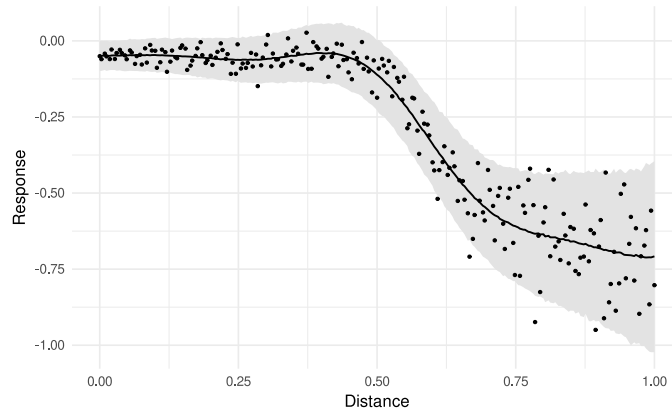
Figure 4.41: Example heteroscedastic regression dataset. Figure taken from [199].

Example of connecting epistemic and aleatoric uncertainty in heteroscedastic regression

We briefly discuss an illustration from the work "Leveraging Graph and Deep Learning Uncertainties to Detect Anomalous Trajectories" [183], shown in Figure 4.42. On the left, we see no training data, resulting in high epistemic uncertainty. To the right of this region, we see aleatoric uncertainty with heteroscedasticity. Here, we have more training data which results in less epistemic uncertainty. However, label noise leads to large aleatoric uncertainty. (In underexplored regions, we, of course, might also have high aleatoric uncertainty, but we have no observations to estimate that.)
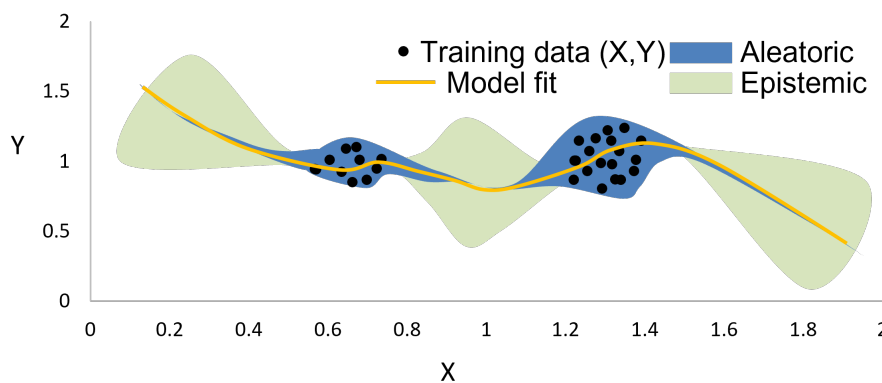
Figure 4.42: Connection of epistemic and aleatoric uncertainty. In distribution, we can still have high aleatoric uncertainty as the generating process can be arbitrarily noisy. Out of distribution, epistemic uncertainty increases. Figure taken from [183].

Another example application of heteroscedastic regression: monocular depth estimation

A sample from the NYU Depth Dataset V2 [181] is shown in Figure 4.43. A single image does not contain all information about which point is at what distance (perspective projection loses information about how far away the projected point was). A black box that does not reflect any light means that there is *no way* we can figure out its depth in the image. Thus, we have an inherent ambiguity of depth corresponding to a single image: we are greeted by aleatoric uncertainty. Also, the "deeper" the object, the more noise there is in the true depth value (noisy label). Furthermore, depending on the position in the image, we have different levels of aleatoric uncertainty – we need per-pixel estimates. This leads to heteroscedastic regression with aleatoric uncertainty: We want to predict a depth map (regression problem) and represent the corresponding aleatoric uncertainties. One could, e.g., do this by using a Gaussian noise model where the aleatoric uncertainty is represented by the standard deviation of the Gaussian at a particular pixel.
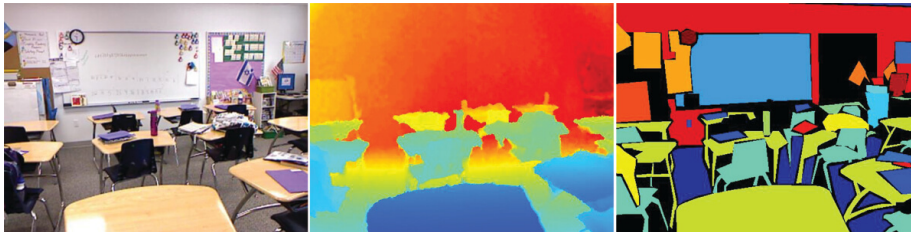
Figure 4.43: Depth estimation is a field that can benefit from uncertainty estimates. On the left, we see an image from an RGB camera. In the middle, we see a preprocessed (metric) depth map where objects farther away are represented by red color. On the right, per-pixel labels are given for the image using *semantic segmentation*. Figure taken from [181].

Log Probability Scoring Rule

To follow our recipe from Section 4.13.1 in a regression setting, let us extend the definition of scoring rules to matching a continuous distribution $Q(Y)$ to the true continuous distribution $P(Y)$. Let y be a sample of $P(Y)$. We define the scoring rule as a function $S(Q, y)$. It is strictly proper when its expected score $\mathbb{E}_{P(Y)} S(Q, Y)$ is maximized iff $Q \equiv P$, i.e., the aleatoric uncertainty estimation is correct. A reasonable expectation at this point is that if we train with some proper scoring rule for the correct output distribution, we will be able to "solve" the problem of aleatoric uncertainty.

The log probability scoring rule for continuous distributions is defined by $S(Q, y) = \log Q(y)$, the log density of the sample. And indeed, one can show that it is strictly proper w.r.t. P.

4.13.5 Using Proper Scoring Rules to Recover $\mathcal{N}(\mu(x), \sigma^2(x)I)$

Let us assume a regression setting where our actual conditional distribution is a *heteroscedastic* (depends on x), *isotropic* (same variance in all directions, defined by a single number $\sigma^2(x)$) Gaussian: $P(y \mid x) = \mathcal{N}(y \mid \mu(x), \sigma^2(x)I)$.

We wish to train a model with output $f(x) = (\hat{\mu}(x), \hat{\sigma}(x))$. Our model has two heads, one for predicting the mean, and one for predicting the standard deviation in each direction, at each

point x. Our predicted distribution is then $Q(y \mid x) = \mathcal{N}\left(y \mid \hat{\mu}(x), \hat{\sigma}^2(x)I\right)$. As can be seen, our approximate distribution is also a heteroscedastic, isotropic Gaussian, just like the true conditional distribution. Therefore, in principle, we can recover the true distributions.

For a fixed input x, the log probability scoring rule yields

$$-S(q, y) = -\log Q(y \mid x) = -\log\left(\frac{1}{\sqrt{(2\pi)^d(\hat{\sigma}^2(x))^d}}\exp\left[-\frac{1}{2\hat{\sigma}^2(x)}\|y - \hat{\mu}(x)\|_2^2\right]\right)$$

where $y \sim P(Y \mid x) = \mathcal{N}(Y \mid \mu(x), \sigma^2(x)I)$ is a sample from the true conditional distribution. This leads to the equation

$$-S(q, y) = \frac{1}{2\hat{\sigma}^2(x)}\|y - \hat{\mu}(x)\|_2^2 + \frac{d}{2}\log\hat{\sigma}^2(x) + C.$$

Minimizing the above quantity has a very intuitive interpretation. We use a scaled L_2 loss for fitting $\hat{\mu}$ to the GT label (i.e., we perform MLE on the mean). The L_2 loss is weighted by the sample-wise weight $\frac{1}{2\hat{\sigma}^2(x)}$. When the model says x is ambiguous (it is having a hard time with it, uncertain about the prediction), i.e., $\hat{\sigma}^2$ is high, the loss penalizes less for the wrong $\hat{\mu}$, as the weight is small. The $\log\hat{\sigma}^2(x)$ term prevents the model from saying that every sample is hard (ambiguous). This is the balancing term (regularizer) that gives a large penalty when the model is always uncertain about the correct value for $\hat{\mu}(x)$. Otherwise, the model could be extremely uncertain for all samples, and that could result in an arbitrarily small objective value. This forces the model only to be uncertain (high variance in predicted conditional) when it has to be (because of target ambiguities). In such cases, the first term does not penalize much, allowing deviations from the (disagreeing) labels. The model learns to say, "I don't know." when the true conditional has high variance, too (conflicting labels). Here, we are still performing MLE, but now both over the means and standard deviations per sample.

Results of Proper Scoring on a Toy Dataset

In this section, we aim to verify that this training gives us good aleatoric uncertainty predictions in practice. To this end, a three-layer FCN is fit to the dataset in Figure 4.44. Aleatoric uncertainty is present in the dataset. The uncertainty is also heteroscedastic, i.e., we have varying degrees of mixed supervision during training. For the "same" x, we will see many different values of y. This can confuse the model, as it often gets marginally different GT labels for very similar inputs. However, the model has a way to express its per-sample confusion: via the predicted standard deviation $\hat{\sigma}(x)$. The model does not have enough capacity to overfit to every single sample. There are also enough data points to counteract this.

The ground truth and predicted conditionals (x-slices) are shown in Figures 4.45 and 4.46, respectively. The predictions were attained by minimizing the loss introduced above, i.e., performing MLE over means and homoscedastic standard deviations. Training with the introduced proper scoring rule really enables the recovery of the correct conditional distribution *given enough data*. The predicted means and standard deviations match well with their GT counterparts: The model represents faithful aleatoric uncertainty estimates. Whenever we have to represent aleatoric uncertainty in our regression problem, we recommend using this
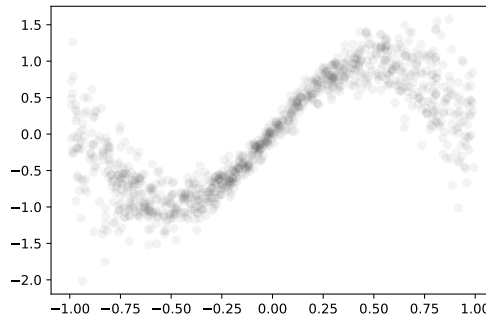
Figure 4.44: Toy regression dataset with heteroscedastic aleatoric uncertainty from the note-book.

loss function or similar probabilistic formulations.
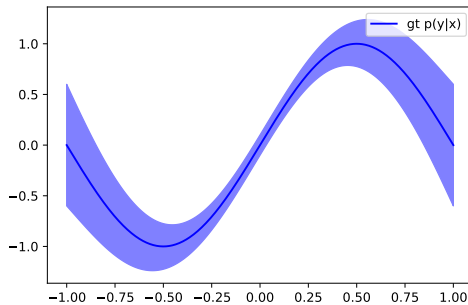
Figure 4.45: Ground truth aleatoric uncer-tainties from the notebook.
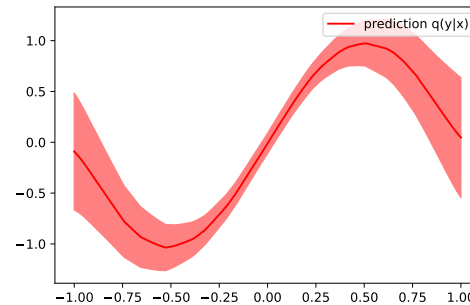
Figure 4.46: Predicted aleatoric uncertain-ties from the notebook.

4.13.6 Aleatoric Uncertainty for Complex Output: Multimodal Future Prediction

So far, we have only seen a 1D regression problem and a 2D classification problem. What about more complex outputs and more complex conditional distributions $P(Y \mid X = x)$? We briefly discuss modeling aleatoric uncertainty for multimodal future prediction [135], illus-trated in Figure 4.47. The task is future prediction, which also matters a lot for self-driving cars: Where will another car go in the next 5 seconds? Where will a pedestrian move in the next 10 seconds? The task inherently has aleatoric uncertainty, as there is no single correct answer for future position prediction in most cases. In regression, we observed the same thing. However, here, the task also inherently has a multimodal true conditional distribution $P(Y \mid X = x)$ over future positions for many image inputs x, as vehicles can have multiple equally plausible trajectories at intersections, roundabouts, and even on highways. In the training set, we will see many cases of mixed supervision. There will be many cases where a vehicle will take turns (1), (2), or (3) as well in Figure 4.47, coming from the same direction. Therefore, we need something better than a single Gaussian to approximate the true multi-modal conditional distribution. We have to accommodate all the complexity that can happen in the future.

Figure 4.47: Presence of aleatoric uncertainty in future prediction. "Given the past images, the past positions of an object (red boxes), and the experience from the training data, the approach predicts a multimodal distribution over future states of that object (visualized by the overlaid heatmap). The vehicle is most likely to move straight (1), but could also continue on the roundabout (2) or turn right (3)." [135] Figure taken from [135].

4.13.7 Aleatoric Uncertainty for Complex Output: Image Inpainting

Now, we turn to model aleatoric uncertainty for image inpainting [222], as shown in Figure 4.48. The task is inpainting. It inherently has aleatoric uncertainty, as there is no single correct answer for missing parts of images.

However, we have thousands of dimensions now, as we do not model a 2D position stochastically; rather, we model an entire missing patch in the image stochastically. (I.e., we have a *joint* distribution of pixels in the missing patch, conditioned on the current input.) We have a large degree of freedom for the missing image patch based on the data manifold. It is complicated to fit a faithful parametric distribution in thousand-dimensional spaces. Usually, we assume a simple parametric distribution. However, such simplifications harm the faithfulness of the distribution we fit. For example, a Gaussian centered at a plausible fill does not explain all possible variations but also has high density for implausible window contents. The true conditional distribution is highly multimodal and complex.

In Figure 4.48, we see samples according to the density $P(y \mid x)$ for two images, which shows us that this distribution is indeed highly multimodal, as all inpainted regions are plausible. We can also think of this as a regression problem with per-pixel aleatoric uncertainty predictions, but we still have to consider thousands of dimensions. By doing so, we assume independence between the output variables (missing pixels in the patch), which is a notable simplification.

4.13.8 Aleatoric Uncertainty for Complex Output: 2D to 3D

Let us now consider the task of 2D to 3D [51] "projection". An intuitive illustration is shown in Figure 4.49. Here, the task is providing 3D reconstruction from a 2D projection. This is highly

Figure 4.48: Presence of aleatoric uncertainty in image inpainting. There are multiple plausible reconstructions of the gray areas. Figure taken from [222].

aleatoric: many different shapes give the same 2D projection. Predicting the 3D shape from a 2D projection is ill-posed. Different experts predict different latent 3D shapes that could have resulted in the given 2D projection. We, again, have "thousands of degrees of freedom" (actually, we have infinitely many) in the shapes that give the same 2D projection. If we consider thousands of dimensions of variation, it is complicated to fit a parametric distribution. Therefore, the space of all possible solutions of 3D shapes is an ultra-high-dimensional space with multimodality. We need to handle such complexity as well.
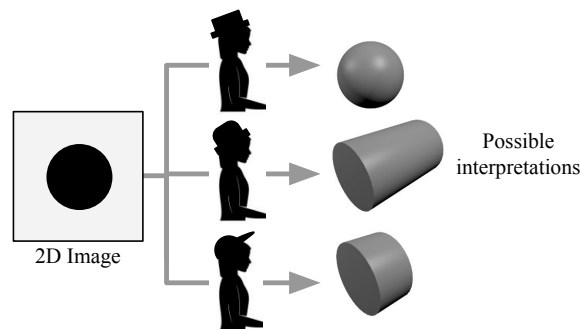
Figure 4.49: Presence of aleatoric uncertainty in 2D to 3D "projection". Infinitely many 3D objects can give rise to the same 2D projection. Figure taken from [51].

4.13.9 Common Challenges in Aleatoric Uncertainty

The presence of aleatoric uncertainty signals that there is no single good answer. We have mixed supervision during training: $Y \mid X = x$ has a non-zero variance. Moreover, *the distribution of possible answers is not simple*. Considering continuous outputs, it is challenging to estimate the true density $P(y \mid x)$ with a Gaussian $Q(y \mid x)$ outside the realm of simple problems. Usually, we have a multimodal true density $P(y \mid x)$, and we often also have a high-dimensional output variable $Y \mid X = x$. Predicting multimodal distributions over high-dimensional output spaces (without simplifications arising from independence assumptions) is extremely hard.

4.13.10 MoGs with Fixed Variance for Aleatoric Uncertainty Estimation in Regression Problems

When the true conditional distribution $P(Y \mid X = x)$ is multimodal, a natural choice for approximating the density $P(y \mid x)$ is to use a Mixture of Gaussians (MoG). They have much capacity to model a lot of different distributions. To simplify derivations and the distributions we model, we assume a fixed variance (homoscedasticity) over all inputs and Gaussians in the mixture. We also only model an isotropic mixture, meaning that we encode the variance of each Gaussian by the same, single fixed number. Isotropic MoGs with fixed variance can be written as

$$Q(y \mid x) = \frac{1}{M} \sum_m Q_m(y \mid x)$$

where

$$Q_m(y \mid x) = \mathcal{N}(y \mid f_m(x), \sigma^2 I).$$

$Q(y \mid x)$ is the *estimated* multimodal conditional distribution. MoGs are "universal approximators" as $\sigma^2 \downarrow 0$ and $M \to \infty$. We can approximate almost any density in the space with this approximator.

> **Additional Information: Alternative Interpretation of MoGs**
>
> MoGs have close connections to kernel regression in non-parametric statistics (in particular, the RBF kernel). It also has connections to GPs: the individual mixture components can be GPs. However, in kernel regression, M means the number of data points in the dataset: As $M \to \infty$, we can recover the true function (the Bayes regressor). Here we are not talking about this interpretation.

An illustration of MoGs used for our purposes is shown in Figure 4.50.
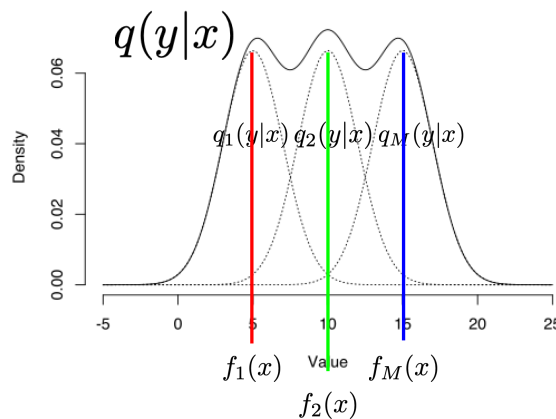
Figure 4.50: Gaussian mixture with three modes corresponding to three model outputs.

For practical purposes to model MoGs, it makes sense to consider the "experts" interpretation and a multi-head DNN, as shown in Figure 4.51. We have a common DNN trunk and M experts (heads) on top of the common DNN trunk. ("Experts" is the word used in the field.) The MoG is described by the M experts, with each component centered around the output of the corresponding head. We also assign parameters separately to different experts – each expert

reports their own hypothesis on the correct output.

Training MoGs for Aleatoric Uncertainty Estimation

How should we train the model to capture the diversity of $P(y \mid x)$ that we will see in the training data? The architecture clearly allows us to do so, but we need to ensure that we encourage this behavior during training.
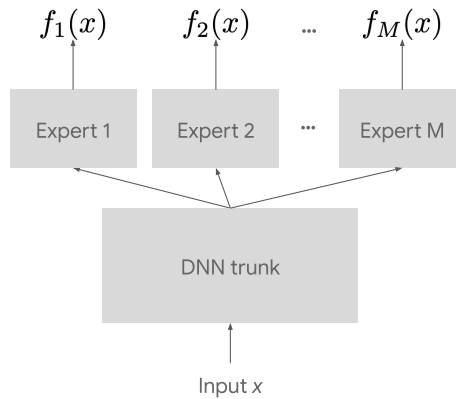
Figure 4.51: Illustration of separately parameterized expert heads on a common DNN trunk.

We have seen that the log probability is a proper scoring rule for capturing a continuous $P(y \mid x)$. We plug our approximate distribution into the negative log probability (NLL) to obtain

$$\mathcal{L} := -\log Q(y \mid x) = -\log \left(\frac{1}{M} \sum_m Q_m(y \mid x) \right).$$

This is the negative log of the MoG PDF evaluated at GT sample y. We aim to write down the gradient of this loss function w.r.t. the weights of individual experts in terms of the losses of the individual heads. The motivation for this will become clearer later. The individual gradients tell us what the signal is we are sending to the individual experts for training their parameters. We will soon see that the total loss $\mathcal{L}$ of a single sample decomposes very nicely into a weighted sum of individual losses of the experts. Thus, the total gradient also has a nice interpretation. We will also see that the K-diverse loss has many connections to this MoG NLL formulation.

Decomposing the Total Gradient

For the shared backbone, the gradient calculation is simple and coincides with how we regularly train DNNs: We backpropagate to the backbone parameters and calculate

$$\frac{\partial \mathcal{L}}{\partial \theta_{\text{backbone}}},$$

then perform a gradient update as usual.

However, the gradients of the experts give rise to insightful intuitions we wish to discuss in detail. To this end, let us analyze the MoG loss in terms of the losses of *individual experts*:

$$\mathcal{L}_m := -\log Q_m(y \mid x).$$

Consider expert $l \in \{1, \ldots, M\}$ that has weights θ_l that are separated from the DNN trunk. By noting that

$$Q_m(y \mid x) = \frac{1}{\sqrt{(2\pi)^d (\sigma^2)^d}} \exp\left(-\frac{1}{2\sigma^2} \|y - f_m(x)\|_2^2\right),$$

we write

$$
\begin{aligned}
\frac{\partial}{\partial \theta_l} \mathcal{L} &= -\frac{\partial}{\partial \theta_l} \log\left(\frac{1}{M} \sum_m Q_m(y \mid x)\right) \\
&= -\frac{\frac{\partial}{\partial \theta_l} \frac{1}{M} \sum_m Q_m(y \mid x)}{\frac{1}{M} \sum_m Q_m(y \mid x)} \\
&= -\frac{\sum_m \frac{\partial}{\partial \theta_l} Q_m(y \mid x)}{\sum_m Q_m(y \mid x)} \\
&= -\frac{\sum_m \frac{\partial}{\partial \theta_l} \exp\left(-\frac{1}{2\sigma^2} \|y - f_m(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2} \|y - f_m(x)\|_2^2\right)} \\
&= \frac{\exp\left(-\frac{1}{2\sigma^2} \|y - f_l(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2} \|y - f_m(x)\|_2^2\right)} \cdot \frac{\partial}{\partial \theta_l} \frac{1}{2\sigma^2} \|y - f_l(x)\|_2^2 \\
&= w_l \cdot \frac{\partial}{\partial \theta_l}\left(-\log \exp\left(-\frac{1}{2\sigma^2} \|y - f_l(x)\|_2^2\right)\right) \\
&= w_l \cdot \frac{\partial}{\partial \theta_l}\left(-\log\left(Q_l(y \mid x) \cdot \sqrt{(2\pi)^d (\sigma^2)^d}\right)\right) \\
&= w_l \cdot \frac{\partial}{\partial \theta_l}\left(-\log Q_l(y \mid x)\right) \\
&= w_l \cdot \frac{\partial}{\partial \theta_l} \mathcal{L}_l,
\end{aligned}
$$

where we defined

$$w_l := \frac{\exp\left(-\frac{1}{2\sigma^2} \|y - f_l(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2} \|y - f_m(x)\|_2^2\right)}$$

the softmax over individual expert outputs. It is important to note here that in any implementation, we need a stopgrad operation for calculating w_l, as we are doing inference (using the model parameters) to calculate the weight itself. To break w_l further apart, we can see that it is of the form

$$
\begin{aligned}
w_l &= \frac{\exp\left(-\frac{1}{2\sigma^2} \|y - f_l(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2} \|y - f_m(x)\|_2^2\right)} \\
&= \frac{\exp\left(\log Q_l(y \mid x) + \log\left(\sqrt{(2\pi)^d (\sigma^2)^d}\right)\right)}{\sum_m \exp\left(\log Q_m(y \mid x) + \log\left(\sqrt{(2\pi)^d (\sigma^2)^d}\right)\right)} \\
&= \frac{\exp(-\mathcal{L}_l)}{\sum_m \exp(-\mathcal{L}_m)} \\
&= \mathrm{softmax}\left(\left[-\mathcal{L}_m\right]_{m=1}^M\right)_l.
\end{aligned}
$$

Therefore, the weights are given by a softmax over the individual negative losses. What is this encoding precisely? Softmax gives an approximation (a soft version) of the $\arg\max$ operation.[32] When $-\mathcal{L}$ is maximal, that entry in the loss vector gets a high probability. The other

---

[32] Many people consider the name a misnomer because of this and propose to use the term "softargmax" instead.

entries get exponentially lower values. One can also write

$$\mathcal{L} = \sum_l \mathrm{stopgrad}(w_l) \cdot \mathcal{L}_l + C$$

which is derived below. When an expert has a small loss on a sample x, it gets a large weight and, thus, a large gradient update. If an expert performs poorly on some sample, they will not learn it later either (with high probability). This allows the best-performing model to perfect its prediction the most on the sample in question. The models, therefore, become more and more diverse as the training progresses. As the models are randomly initialized, they will probably perform well on different samples to begin with, and their diversity is reinforced through the loss formulation throughout the training procedure.

Effect of the Fixed Variance σ^2

Let us discuss the effect of the fixed variance σ^2 on *diversity*. We can also write the weight w_l as

$$w_l = \frac{\exp\left(\mathrm{logit}_l / \sigma^2\right)}{\sum_m \exp\left(\mathrm{logit}_l / \sigma^2\right)}$$

where $\mathrm{logit}_l = -\frac{1}{2}\|y - f_l(x)\|_2^2$ (w.r.t. the softmax). The σ^2 parameter acts as a *temperature* parameter for the softmax operation to calculate the weights. We divide the logit values by this before applying $\exp$. As $\sigma^2 \to \infty$, the softmax operation gives a uniform distribution. We are training every expert equally. This is almost like training an ensemble but with a shared backbone. We only have minor differences between experts. When $\sigma^2 = 1$, we have no temperature scaling in the softmax. As $\sigma^2 \downarrow 0$, the softmax result ends up being a one-hot vector indicating the argmax among negative losses. Here we are only training the most correct expert in this case, and we completely stop gradients to other experts. This results in a very diverse set of experts.

To summarize, for high-dimensional and structural, possibly multimodal $P(y \mid x)$, $Q = \mathrm{MoG}$ (with different heads of the network) and proper scoring lead to K-diverse and smooth K-diverse loss functions. The loss favors more correct solutions among the set of experts. It encourages the models to produce more diverse solutions. The choice of σ^2 controls how diverse we want the experts to be. It is a hyperparameter.

> **Additional Information: Derivation of the Alternative Formulation of the Log-Likelihood MoG Loss $\mathcal{L}$**
>
> The loss we start out from is
>
> $$\mathcal{L} = -\log\left(\frac{1}{M}\sum_l Q_l(y \mid x)\right)$$
> $$= -\log\left(\frac{1}{M}\sum_l \frac{1}{\sqrt{(2\pi\sigma^2)^d}}\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)\right).$$

We want to show that

$$\mathcal{L} = \sum_l \text{stopgrad}(w_l) \cdot \mathcal{L}_l + c,$$

where

$$w_l = \frac{\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2}\|y - f_m(x)\|_2^2\right)}$$

and

$$\mathcal{L}_l = -\log Q_l(y \mid x).$$

Let us calculate the gradient of both formulations (LHS and RHS) w.r.t. the entire network's weights:

$$
\begin{aligned}
\frac{\partial}{\partial\theta}\mathcal{L} &= -\frac{\frac{1}{M}\sum_l \frac{1}{\sqrt{(2\pi\sigma^2)^d}}\frac{\partial}{\partial\theta}\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)}{\frac{1}{M}\sum_l \frac{1}{\sqrt{(2\pi\sigma^2)^d}}\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)} \\
&= -\frac{\sum_l \frac{\partial}{\partial\theta}\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)}{\sum_l \exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)} \\
&= -\sum_l \frac{\frac{\partial}{\partial\theta}\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2}\|y - f_m(x)\|_2^2\right)} \\
&= -\sum_l \frac{\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)}{\sum_m \exp\left(-\frac{1}{2\sigma^2}\|y - f_m(x)\|_2^2\right)}\frac{\partial}{\partial\theta}\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right) \\
&= -\sum_l w_l \frac{\partial}{\partial\theta}\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right) \\
&= \sum_l w_l \frac{\partial}{\partial\theta}\left(\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right),
\end{aligned}
$$

and similarly,

$$
\begin{aligned}
&\frac{\partial}{\partial\theta}\left(\sum_l \text{stopgrad}(w_l) \cdot \mathcal{L}_l + c\right) \\
&= \sum_l w_l \frac{\partial}{\partial\theta}\mathcal{L}_l \\
&= \sum_l \frac{\partial}{\partial\theta}\left(-\log\frac{1}{\sqrt{(2\pi\sigma^2)^d}}\exp\left(-\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right)\right) \\
&= \sum_l \frac{\partial}{\partial\theta}\left(\log\sqrt{(2\pi\sigma^2)^d} + \frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right) \\
&= \sum_l w_l \frac{\partial}{\partial\theta}\left(\frac{1}{2\sigma^2}\|y - f_l(x)\|_2^2\right).
\end{aligned}
$$

This necessarily means that the two expressions only differ in constants (w.r.t. the weights). However, they *do* differ in constants, as we saw that we canceled some zero-gradient terms in the second expression but did not cancel any in the first.

4.13.11 K-Diverse Loss

The special case of MoG NLL training when $\sigma^2 \downarrow 0$ is called the *k-diverse loss* or *min-loss* in literature.

$$\mathcal{L} = \sum_l \text{stopgrad}(w_l) \cdot \mathcal{L}_l \xrightarrow{\sigma^2 \downarrow 0} \sum_l \mathbf{1}(l = \arg\max_p w_p) \cdot \mathcal{L}_l = \min_l \mathcal{L}_l,$$

as $\arg\max_p w_p = \arg\min_p \mathcal{L}_p$. The total loss is the minimal loss over the set of experts in this case. The gradient is only backpropagated through the most correct expert for each sample. That is, only the expert corresponding to the best answer gets updated. For a fixed input sample, we nearly always train the same expert. This explains why we call the heads "experts": Every head becomes an expert for a subset of samples. The model is incentivized to give diverse answers to every question. This is because only one of the predictions from the heads has to be correct to achieve a very low loss value. This is, of course, a very cheap shortcut for diversity. However, the loss achieves its aim: It diversifies the output predictions.

We see different forms of the min-loss in the literature, hidden under different notations. We will discuss some of these in the following sections.

DiverseNet

The DiverseNet [51] method uses the K-diverse loss:

$$\mathcal{L}_{\text{div}} = \sum_{(x,\mathcal{Y}) \in \mathcal{D}} \sum_{y \in \mathcal{Y}} \min_{c \in \mathcal{C}} \ell\left(f(c,x), y\right)$$

where ℓ is the loss function, c is the expert (hypothesis/head) index, and y corresponds to multiple GT labels for each x, sampled according to $P(y \mid x)$.

This is the same K-diverse loss we saw before; the difference is that we always have multiple labels too, that are sampled from $P(y \mid x)$.

The authors of [51] want DiverseNet to output diverse solutions for each input x. However, they also want each output to be sensible to avoid shortcuts like predicting a constant digit with each expert on MNIST with ten experts. To this end, they also employ a catch-up loss

$$\mathcal{L}_{\text{catchup}} = \frac{1}{|\mathcal{C}|} \sum_{(x,\mathcal{Y}) \in \mathcal{D}} \max_{c \in \mathcal{C}} \min_{y \in \mathcal{Y}} \ell(f(c,x), y)$$

that encourages the worst-performing model on the sample to improve as well.

If we train with the final loss

$$\mathcal{L} = \mathcal{L}_{\text{div}} + \beta \mathcal{L}_{\text{catchup}}$$

with a tuned β parameter, we perform well on their proposed MNIST outpainting task, as illustrated in Figure 4.52. The input here is the upper right fraction of an MNIST image. The output is the completed MNIST image. The authors consider many different loss functions and argue that only theirs is giving a notably diverse set of solutions. Their proposed loss is a good choice when such a high aleatoric is present in most samples.

We also discuss Figure 4.53, where the authors apply their model to a traffic direction prediction task. Given an image, the task is to predict the possible passage of traffic in the future

Figure 4.52: Predictions of various models on the MNIST outpainting task. "[...] For each model, [...], we make eight predictions on each image. The ground truth row shows the $\mathcal{Y}$ values in MNIST$_{OCC}$, found from the test set using nearest neighbor lookup – the left-most value in $\mathcal{Y}$ is the image from which x was generated. [...]" [51] Figure taken from [51].
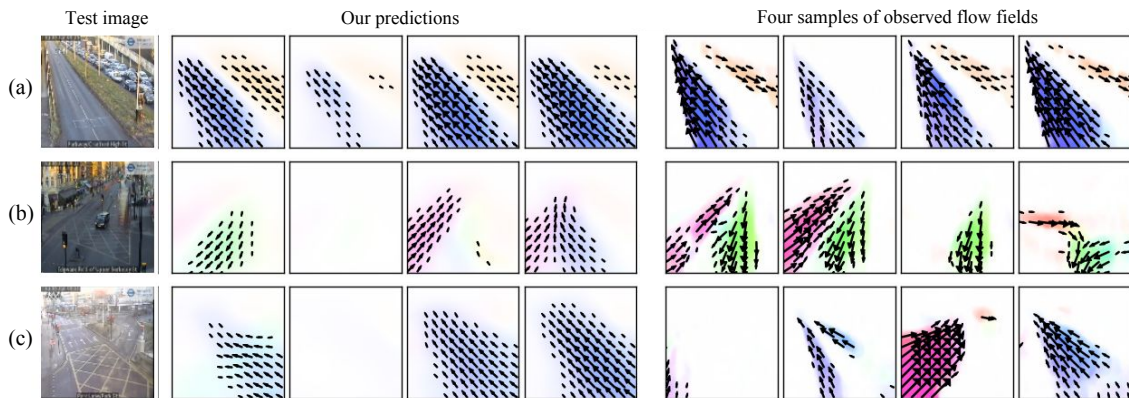
Figure 4.53: Results on the traffic flow prediction task, which is inherently aleatoric. The GT flow maps are notably diverse. The model makes very accurate aleatoric uncertainty predictions in some cases, but there are also several incorrect predictions. Figure taken from [51].

from the same view. The GT $P(y \mid x)$ flow fields from the same view show that the task has large inherent aleatoric uncertainty. The authors argue that their method gives a diverse set of predictions for the possible flow of traffic that are also truthful to the observations.

Finally, we consider the 2D to 3D reconstruction task, illustrated in Figure 4.54. The input is a silhouette: a 2D projection with removed texture. The task is to reconstruct the 3D shape from the 2D silhouette, which we have seen to induce many ambiguities previously. The method produces a diverse set of possibilities that are all quite plausible.

Mixture Density Networks

The min-loss is also hidden in the work on improved mixture density networks [135]. They apply the *Winner-Takes-All (WTA) [71] loss* (which is the same as our previous K-diverse loss) to
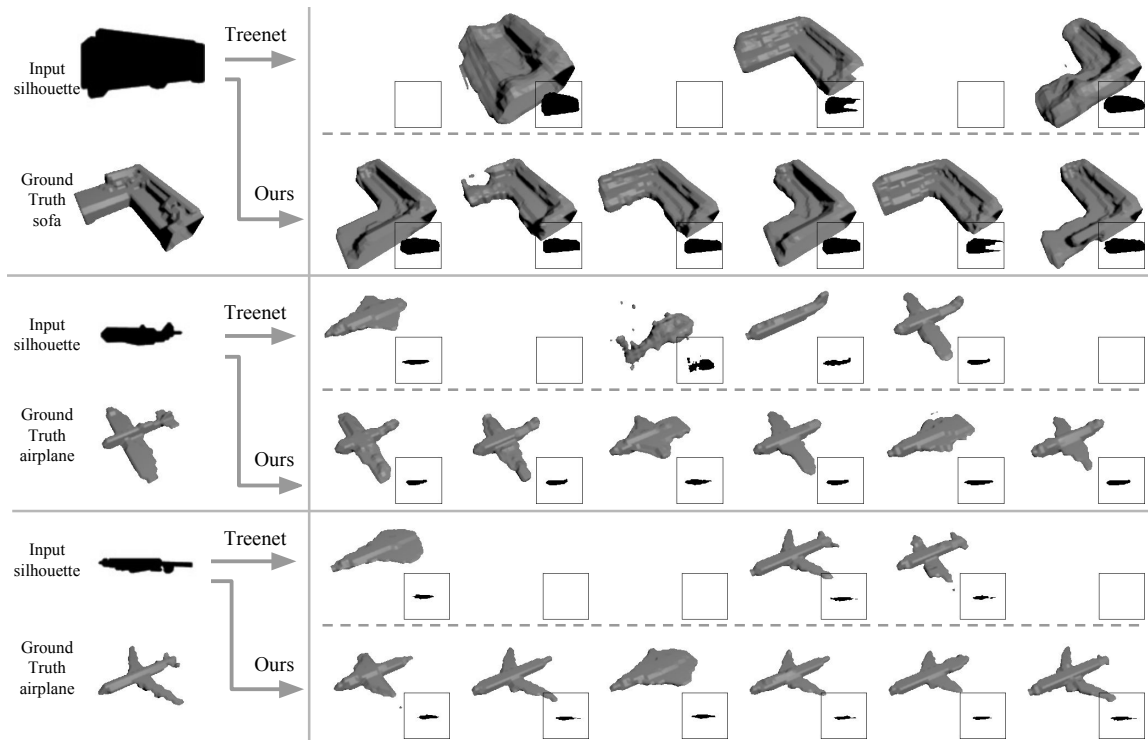
Figure 4.54: Results on the 2D to 3D reconstruction task. The reconstructed objects are quite diverse in appearance and reasonably plausible. Figure taken from [51].

their novel problem of future prediction:

$$\mathcal{L}_{\mathrm{WTA}} = \sum_{k=1}^{K} w_k \ell(h_k, \hat{y})$$

$$w_i = \mathbf{1}\left(i = \arg\min_k \|\mu_k - \hat{y}\| \right).$$

where k is the expert (hypothesis) index, μ_k is the prediction of expert k, $\hat{y}$ is the GT (!) location, and ℓ is the loss function. The weight term is a Kronecker delta indicating the most correct solution out of the set of experts (i.e. when the predicted mean is closest to the GT location). The winner is the expert with minimal prediction distance. That expert is the only one trained on that sample; it takes the full penalty.

Let us discuss Figure 4.55. As shown by the authors, training with the WTA loss results in more diversified and more multimodal solutions. The red boxes correspond to the current and previous frame locations. The magenta boxes denote the future frame GT location. The isoprobability contours correspond to the predicted densities $Q(y \mid x)$ over the future frame location by the methods. Ideally, the prediction should enclose the actual location and be diverse (multimodal). The method of [135] seems to be superior to the baselines. The min-loss is good for enforcing diversity in ill-posed tasks.

So far, we have considered the prediction of higher-dimensional outputs in regression tasks using DNNs while modeling aleatoric uncertainty. In the last section, we will see use cases of aleatoric uncertainty in representation learning.
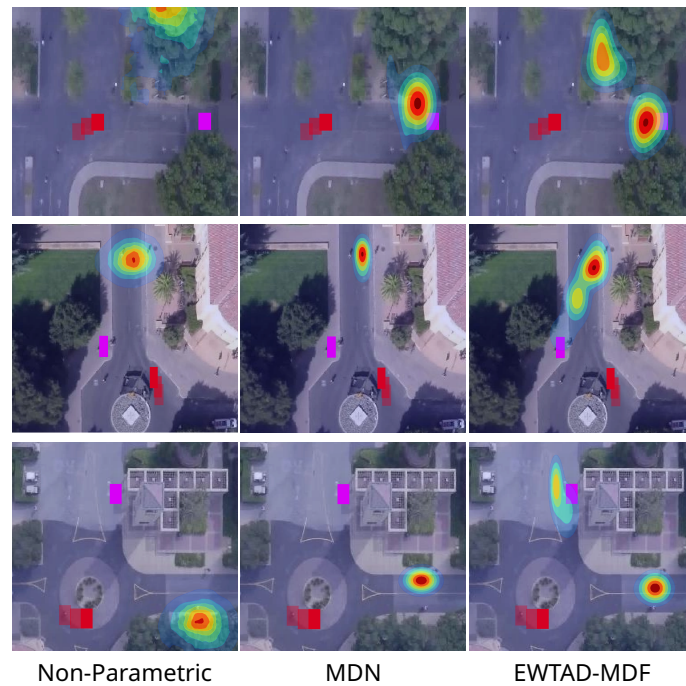
Figure 4.55: Qualitative evaluation of different probabilistic methods on future location prediction. Non-parametric and MDN are the baselines, and EWTAD-MDF is the proposed method. "Given three past locations of the target object (red boxes), the task is to predict possible future locations. A heatmap overlay is used to show the predicted distribution over future locations, while the ground truth location is indicated with a magenta box. Both variants of the proposed method capture the multimodality better, while MDN and non-parametric methods reveal overfitting and mode-collapse." [135] Figure taken from [135].

4.14 Aleatoric Uncertainty in Representation Learning

Aleatoric uncertainty also has close ties to representation learning. Let us first consider what we mean by a representation.

Definition 4.17: Representation

A representation is a mapping of input x to some space of latent variables z. The semantic distance between two inputs (e.g., images), $d(x_1, x_2)$, is encoded through more straightforward metrics in the latent space $\mathcal{Z}$. An example is the Euclidean distance:

$$d(x_1, x_2) = \|z_1 - z_2\|_2.$$

This is possible in principle because the mapping from the input space to the space of latent variables can be highly non-linear, making it potentially expressive enough to extract semantic meaning through the embedding procedure.

4.14.1   Prelude: Representation Learning

> **Definition 4.18: Retrieval Task**
>
> The retrieval task in representation learning refers to returning inputs x from a dataset that are semantically similar to a test input x_{test}, as defined by the task. One can perform this efficiently if they possess a performant feature embedding f, as they only have to compute the feature embeddings of the dataset samples once. Then they only have to search for k-NNs in the latent space. Finally, e.g., using a simple hash table, one can return the original dataset samples corresponding to the k-NN feature embeddings. An example of a system solving a retrieval task is Google Images.

For learned representations, we can consider (1) the retrieval task detailed above, but also (2) the downstream fine-tuning of our feature embedding on a classification problem. These both become much more manageable and computationally more efficient if we already have a robust embedding.

For retrieval, instead of having to iterate through all images x_i in the dataset, one can index them by passing them through the feature embedding and saving the feature vectors. For comparing semantic distances, one just has to use cosine similarity or L_2 distance in the latent space, which is very efficient to compute. Obtaining k-NNs can be done very efficiently nowadays, e.g., using `faiss` [89].
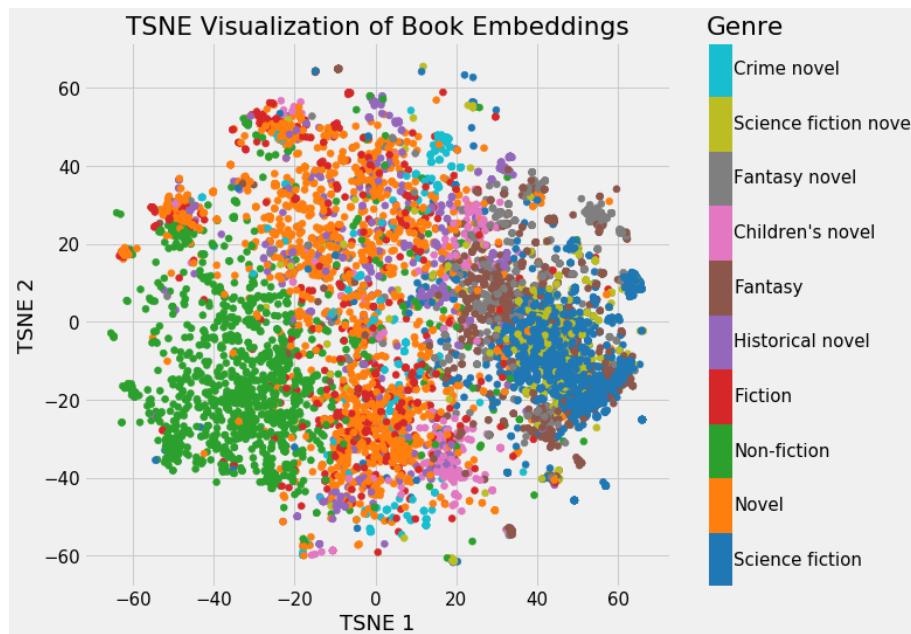
Figure 4.56: Illustration of t-SNE word embeddings for books. Books of the same genre are often clustered together, although no label supervision is used in the method. Figure taken from [45].

Suppose we are only interested in a low-dimensional, unsupervised, interpretable representation of our book dataset samples (Figure 4.56). In that case, we can obtain a 2D t-SNE representation of different books where similarities in the original space are aimed to be

preserved in the embedding space. Instead of saving all characters in every book, we embed them in a 2D space where the (distance-based) similarities are as preserved as possible. The original samples do not have to be saved. As can be seen, this representation gives a coarse clustering w.r.t. the genre, even though the method did not have access to genre labels. Thus, the 2D location encodes the genre of the book to some extent.

Training a representation

Our goal is to estimate a deterministic mapping $z = f(x)$ (construct the latent space) such that the simple distance metric (e.g., L_2 distance or cosine similarity) in the latent space reflects semantic distances (class membership) of the corresponding inputs. The ingredients (i.e., the training data) usually come in two forms.

- **Pairs.** $(x_1, x_2, y = 1)$ when x_1, x_2 are from the same category or $(x_1, x_2, y = 0)$ when they are not.
- **Triplets.** (x_a, x_p, x_n), where x_a is the anchor image, x_p is a positive match to it (should be mapped closer), and x_n is a negative match to it (should be mapped farther). Here we do not have to consider categories at all (discrete); we just have to specify relative similarities (relative distances in the latent space). As we are only supervising with relative distances to train the embedding, all we know is that x_a and x_p should be closer in the latent space than x_a, x_n. This can, however, encode more subtle relationships than the previous supervision technique: we cannot only specify which samples belong to the same group, but we can also specify which groups are more similar and which are less, and which samples are more similar in the same group.[33]

In either case, we only get a very coarse description of where each sample should be mapped in the latent space. We have to construct the latent space where, e.g., the Euclidean distance should encode semantic similarity using this very coarse training signal. The latent embedding z is never observed (not part of supervision); it has to be inferred.

4.14.2   Aleatoric Uncertainty in Representation (Learning)

Aleatoric uncertainty can be present in representation learning.

In a typical embedding space of 2-digit MNIST images, very clear samples of classes 17 and 47 with no ambiguity are mapped to two quite different points in the embedding space. When we observe a new sample of class 17 with no ambiguity, it should be ideally mapped close to the previous 17 sample (maybe not very close if the handwriting is quite different). It is quite hard to confuse 17 with 47, so even if the handwriting differs between the 17 samples, they should still be reasonably far away from 47. The same should happen with a new sample of class 47 with no ambiguity, just with the classes flipped in the previous scenario.

Now, consider a sample where a crucial part of the first digit is deleted. Both 17 and 47 are plausible true classes, i.e., we have aleatoric uncertainty. It is now unclear where this image should be put in the latent embedding space. Putting it close to 47 makes sense;

---

[33]This is a somewhat idealistic description of triplet supervision where the negative mining is proportional to the semantic similarity. In practice, people consider two extremes: either using a classification dataset to provide positive/negative samples (very coarse) or using self-supervised learning where only the image itself is considered positive to itself (too fine-grained). This still works in practice, though.

this could have been a 47. However, putting it close to 17 makes sense as well for the same reason. Putting it in the middle is not really a good solution: it does not signal closeness to either of the two plausible classes, even though it is plausibly very close to *both* of them. This is impossible to communicate with deterministic embeddings. Whenever the inverse generative process $P(Z \mid X)$ is noisy, deterministic embeddings cannot faithfully represent the true relationship of an embedding and an image.[34]

We argue that a good solution is not to consider point embeddings (deterministic functions of x) but rather an entire distribution over latent variables given an input x whose distribution is multimodal (e.g., a MoG). Such a probabilistic embedding can put two different components in two different parts of the space if our latent is uncertain (torn between two categories). For, e.g., a Gaussian noise image, we would ideally get a probabilistic embedding spread across the space, not just torn between two distributions. This represents what the underlying latent generator of X – i.e., the latent variable Z – *could* have plausibly been, whereas a point embedding in the middle of plausible embedding regions represents what latent code the noisy image was *surely* generated from.[35]
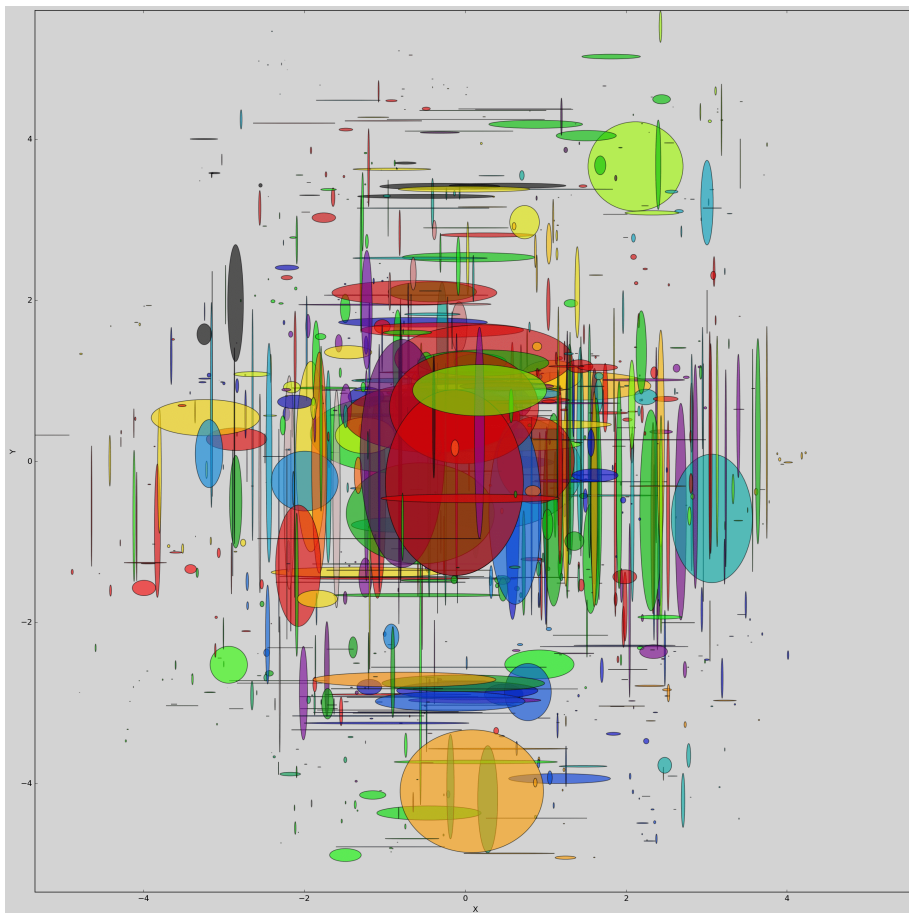
Figure 4.57: Estimated latent heteroscedastic Gaussian distributions $P(Z \mid X)$. Figure taken from [150].

---

[34] $P(X \mid Z)$ might even be deterministic, but $P(Z \mid X)$ can still be stochastic, and vice versa.
[35] This does not model the noise in the inverse generative process. Rather, it tries to assign noisy images to single points in the embedding space as well, which is not a reasonable probabilistic interpretation.

Problem statement for representation of aleatoric uncertainty in representation learning

We discuss "Modeling Uncertainty with Hedged Instance Embedding" [150]. The goal is to estimate latent distributions $P(Z \mid X = x)$. Z is a latent variable here; it could be, e.g., estimated by a heteroscedastic diagonal Gaussian. Then, probabilistic embeddings look like in Figure 4.57. These embeddings are 2D Gaussians for corrupted test images from 2-digit MNIST. It is desired that corrupted test images have embeddings that spread density across larger regions of the space than clean images.

The ingredients (training data) of the method are pairs: $(x_1, x_2, y = 1)$ when x_1, x_2 are from the same category or $(x_1, x_2, y = 0)$ when they are not. Note again that Z is not observed – it is a latent variable.

Estimating $P(Z \mid X = x)$ is, thus, even more challenging than estimating $P(Y \mid X = x)$ with complete observations (the typical aleatoric uncertainty case, where (x, y) pairs are available in a large dataset). Here we do not have any observations of Z. We have to learn good representations from noisy observations and a lack of observations of the true latents.

How to handle representation learning with aleatoric uncertainty? Just as in deterministic representation learning, we will compute a quantity that expresses how likely x_1 and x_2 are to be similar. The catch is to introduce a latent distribution $Q(Z \mid X = x)$ in the middle for each input x. This estimates the noisy inverse generative process $P(Z \mid X = x)$. We then marginalize over this distribution, i.e., compute the similarity for all possible latents the image may represent:

$$
\begin{aligned}
Q(Y = 1 \mid x_1, x_2) &= \int Q(Y = 1, z_1, z_2 \mid x_1, x_2) \, dz_1 dz_2 \\
&= \int Q(z_1, z_2 \mid x_1, x_2) Q(Y = 1 \mid z_1, z_2, x_1, x_2) \, dz_1 dz_2 \\
&= \int Q(z_1 \mid x_1, x_2) Q(z_2 \mid x_1, x_2) Q(Y = 1 \mid z_1, z_2) \, dz_1 dz_2 \\
&= \int Q(z_1 \mid x_1) Q(z_2 \mid x_2) Q(Y = 1 \mid z_1, z_2) \, dz_1 dz_2
\end{aligned}
$$

where we used the following assumptions in order:

$$
\begin{aligned}
Z_1 &\perp\!\!\!\perp Z_2 \mid X_1, X_2 \\
Y &\perp\!\!\!\perp X_1 \mid Z_1, Z_2 \\
Y &\perp\!\!\!\perp X_2 \mid Z_1, Z_2 \\
Z_1 &\perp\!\!\!\perp X_2 \mid X_1 \\
Z_2 &\perp\!\!\!\perp X_1 \mid X_2.
\end{aligned}
$$

We marginalize over the predicted posteriors. $Q(Y = 1 \mid z_1, z_2)$ is a distance-based probability. The distance of z_1 and z_2 should tell us whether they correspond to the same category.

To make this DNN-friendly (computational-graph-wise), we model the probabilistic embeddings by parameterizing heteroscedastic, isotropic Gaussians:

$$
Q(z \mid x) = \mathcal{N}(z \mid \mu(x), \sigma^2(x)I)
$$

where $\mu(\cdot)$ and $\sigma^2(\cdot)$ are DNNs. (We predict the parameters of the latent Gaussian from the input X.) We parameterize the match probability

$$Q(Y = 1 \mid z_1, z_2) = \mathrm{sigmoid}\left(-a\|z_1 - z_2\|_2 + b\right) \in (0, 1)$$

with some smooth calibration function (sigmoid of affine transformation) with learnable parameters $a > 0$ and b. We want to keep a notion of Euclidean distance in the latent space. Therefore, we let the match probability depend on the Euclidean distance in some way. For example, $b = 0$ is not a good choice: In this case, distance 0 means 0.5 probability in the match. Thus, the simple sigmoid is not well-calibrated to the 0-1 range. $a > 0$ is required to ensure that a is not flipping the sign of distance. We want a closer distance to encode more similarity; therefore, we need $-a$ to be negative. The exact linear scaling and translation have to be calibrated.

To make the formulation even more DNN-friendly, we perform a Monte Carlo sampling to approximate the integral

$$Q(Y = 1 \mid x_1, x_2) = \int Q(z_1 \mid x_1) Q(z_2 \mid x_2) Q(Y = 1 \mid z_1, z_2) \, dz_1 dz_2$$

$$= \mathbb{E}_{Z_1 \sim Q(Z_1 \mid X_1 = x_1), Z_2 \sim Q(Z_2 \mid X_2 = x_2)} \left[Q(Y = 1 \mid z_1, z_2)\right]$$

$$\stackrel{\mathrm{MC}}{\approx} \frac{1}{K^2} \sum_{k_1, k_2} Q(Y = 1 \mid z_1^{(k_1)}, z_2^{(k_2)})$$

where $z_i^{(k_i)} \sim Q(z_i \mid x_i) = \mathcal{N}\left(z_i \mid \mu(x_i), \sigma^2(x_i)I\right), i \in \{1, 2\}$. We sum over the match probability values for every possible pair of samples from each distribution. With that, we obtain the final match probability estimate.

We differentiate through the sampling back to the network parameters via the reparameterization trick:

1. We sample $\epsilon_i^{(k_i)} \sim \mathcal{N}(0, I)$.
2. We set $z_i^{(k_i)} = \sigma(x_i) \cdot \epsilon_i^{(k_i)} + \mu(x_i)$.

This way, sampling is separated from backpropagation.

Finally, the NLL loss is minimized w.r.t. the parameters of our neural network (MLE):

$$\mathcal{L}(x, y) := -y \log\left(\frac{1}{K^2} \sum_{k_1, k_2} Q(Y = 1 \mid z_1^{(k_1)}, z_2^{(k_2)})\right)$$

$$- (1 - y) \log\left(1 - \frac{1}{K^2} \sum_{k_1, k_2} Q(Y = 1 \mid z_1^{(k_1)}, z_2^{(k_2)})\right)$$

$$= -\log\left(\frac{1}{K^2} \sum_{k_1, k_2} Q(Y = y \mid z_1^{(k_1)}, z_2^{(k_2)})\right)$$

where $y \in \{0, 1\}$ is the GT label for the (x_1, x_2) pair and $z_i^{(k_i)} \sim Q(z_i \mid x_i) = \mathcal{N}\left(z_i \mid \mu(x_i), \sigma^2(x_i)I\right)$.

Although the NLL was a proper scoring rule for many of the previous problems, we now

have no guarantee of propriety. It is not guaranteed that we recover the true latent distribution, not even if the true distribution is in our parametric family $\mathcal{Q}$. This is because, unlike previously, where we said that if we use log-prob that guarantees the recovery of the true distribution in expectation, now we are talking about latent variables, for which we do not have many guarantees (so far). But we will come back to this in Section 4.14.4.

Overview of the method

An overview of the method is shown in Figure 4.58. Here, a CNN outputs the posterior parameters. Sampling is differentiable via the reparameterization trick. The learned sigmoid gives us the match probabilities: we average over K^2 of them. Finally, we take the negative log of the final match probability w.r.t. the true label $y \in \{0, 1\}$.
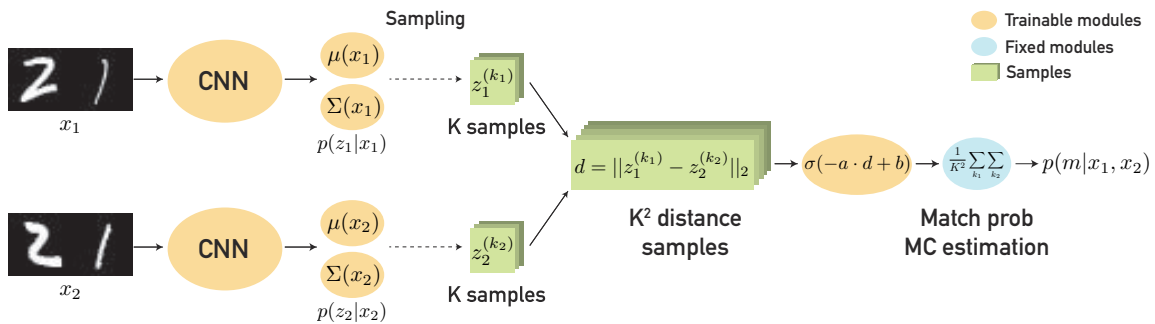
Figure 4.58: Overview of [150], a method for representing aleatoric uncertainty in learned embeddings. CNNs predict the mean and covariance corresponding to the latent posterior conditioned on the input. These are used to calculate semantic distances between the two inputs that are further transformed into a match probability. Figure taken from [150].

Results on [N ∈ {2, 3}]-digit MNIST

Samples from [N ∈ {2, 3}]-digit MNIST are shown in Figure 4.59. Statistics of the used datasets are shown in Table 4.7. We perform data corruption which introduces more aleatoric uncertainty. The clean inputs are original MNIST samples; corrupted inputs contain randomly occluded parts. This leads to ambiguities and variance in $P(y \mid x)$.

Figure 4.59: Samples from [N ∈ {2, 3}]-digit MNIST. Digits from the original MNIST dataset are concatenated together.

Table 4.7: Statistics of [N ∈ {2, 3}]-digit MNIST. Table taken from [150].

Number Digits	Total Classes	Training Classes	Unseen Test Classes	Seen Test Classes	Training Images	Test Images
2	100	70	30	70	100 000	10 000
3	1000	700	100	100	100 000	10 000

In Figure 4.60, we compare the embedding distributions of clean and corrupt images from 2-digit MNIST. 2D Hedged Instance Embeddings (MoGs with a single component) are used.

Although there was no guarantee regarding the correct recovery of the true latent posterior, we do see a high correlation between the image's cleanliness and the embeddings' variance. For clean images, we have barely any variance in the embeddings. For corrupted images, however, we see larger variances in the latent posteriors. Here, we do not have an isotropic Gaussian posterior family, which makes the method more expressive.
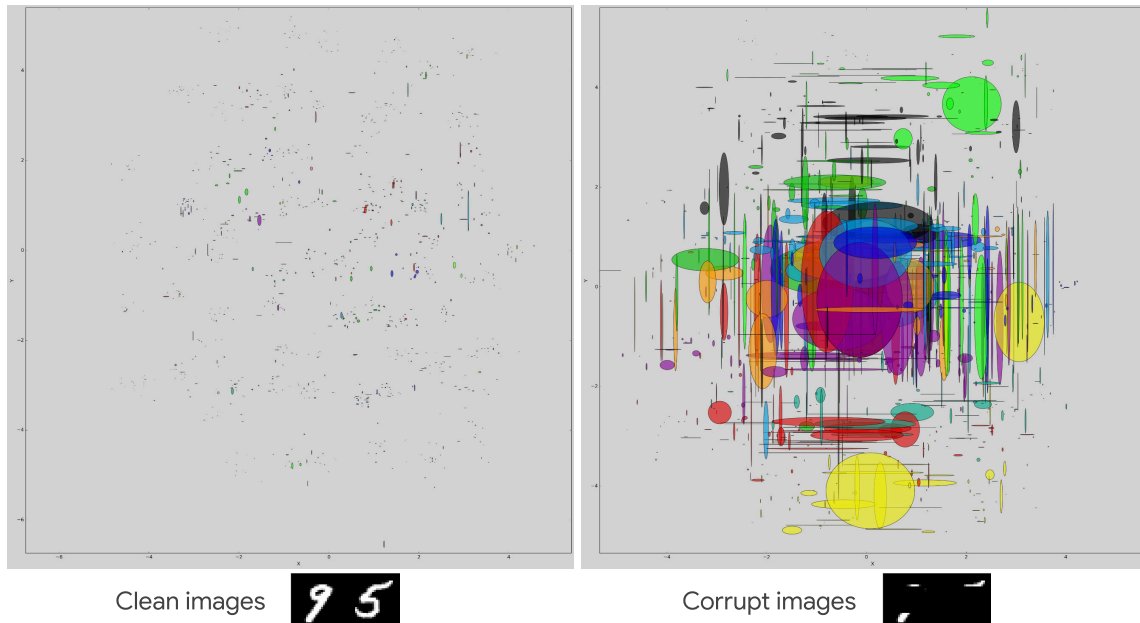
Figure 4.60: Comparison of 2D embedding distributions of clean (left) and corrupted (right) images. As expected, corrupted images also have a latent posterior with significantly more spread. Base figure taken from [150].

In Figure 4.61, we turn to the 3-digit MNIST dataset and use 3D Hedged Instance Embeddings. We observe the same trend as in Figure 4.61, supporting that these observations are quite general when training with this method.

Finally, let us discuss Figure 4.62, a visualization of the embedding distributions of particular clean and corrupt images. The squares correspond to the centroids of each class in the 2D embedding space. (Average mean of the posterior for the same class using clean samples.) For the clean 03 image, the latent posterior is a very narrow Gaussian close to the centroid of class 03. For the corrupted '68' image, the Gaussian has a lot more variance. For the second digit, the network is uncertain whether it is 8, 1, or 9. It hedges its bets across the three class centroids corresponding to 61, 68, and 69. For the corrupted '52' image, the Gaussian is, again, broad. This is because nearly the entire 5 is blocked out. The model finds 52, 32, 22, 42, 12, and 72 all plausible. Humans would also find these plausible. We also see that the neighboring centroids are well structured, allowing the model to hedge its bets across the plausible labels well (as the model was also *trained* with corrupted samples). In general, similar images are embedded in neighboring 2D regions.

The variance of the Gaussians not only encodes aleatoric uncertainty but also shows us a good set of plausible classes the sample could belong to (according to human inspection). It kind of recovers the true $P(z \mid x)$ in that sense, as it empirically recovers all latent variations. Still, we have no guarantees.

Clean images   Corrupt images
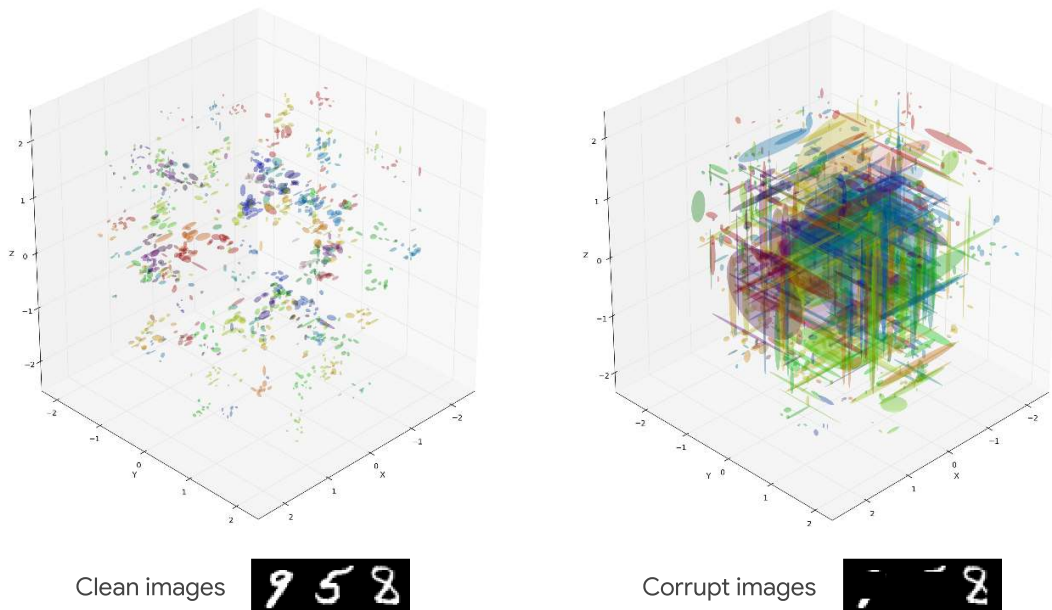
Figure 4.61: Comparison of 3D embedding distributions of clean and corrupted images. Corrupted images have latent posteriors with more spread. Base figures taken from [150].

Note: The latent dimensionality does not have to correspond to the number of digits. However, this way, we get interpretable embeddings (but nothing forces the embeddings to be axis-aligned w.r.t. the digits).
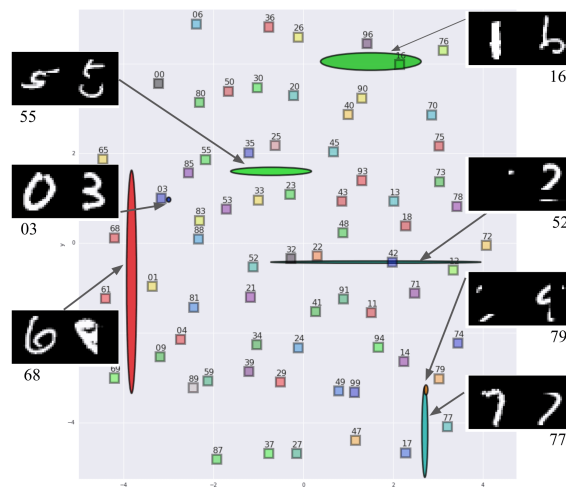
Figure 4.62: Visualization of 2D embedding distributions of particular clean and corrupt images. The distributions are aligned across plausible numbers (reconstructions). Figure taken from [150].

4.14.3 Multimodal Joint Representations and Aleatoric Uncertainty

The work on aleatoric uncertainty in multimodal joint representations [39] extends the previous work to more scenarios. The idea was that aleatoric uncertainty could also naturally arise in multimodal joint embeddings. An overview of text-to-image and image-to-text retrieval is given in Figure 4.63. The common embedding space accommodates both image and text
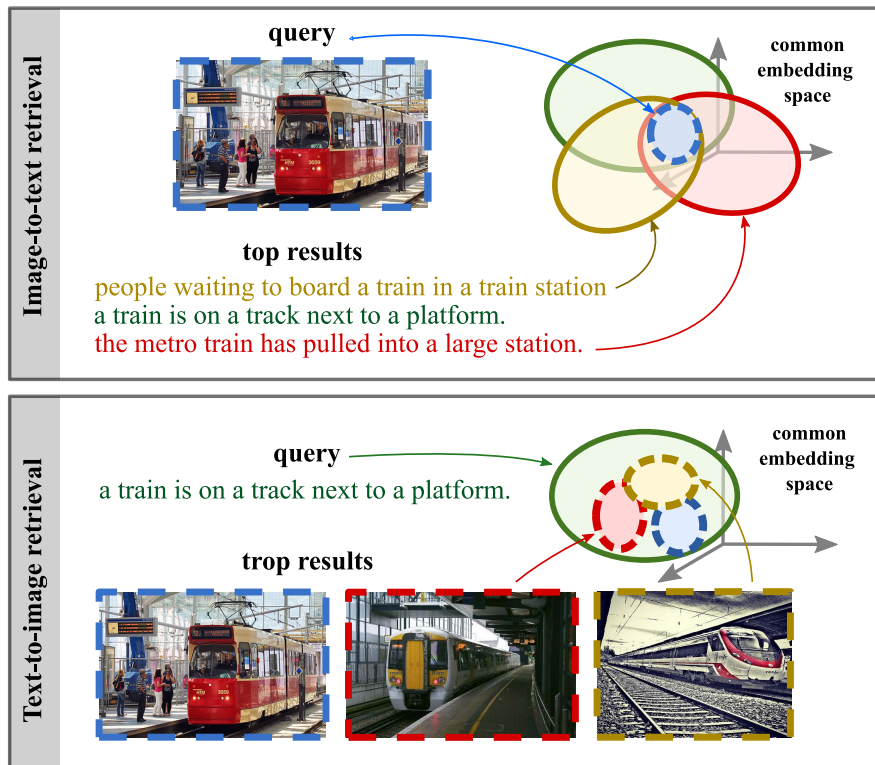
Figure 4.63: Visualization of the joint embedding space of images and texts, and their inherent uncertainty. Many images correspond to the same query, but also many queries correspond to the same image. Figure taken from [39].

embeddings. Given an image, we might be interested in what captions it (likely) corresponds to. However, given the image, we have multiple possibilities for the captions that correctly describe the image. Probabilistic embeddings, therefore, make more sense than point embeddings for multimodality. This also holds the other way around: Multiple different images make sense for the same query. To model this multiplicity, one can embed every image and caption into a Gaussian in the embedding space.
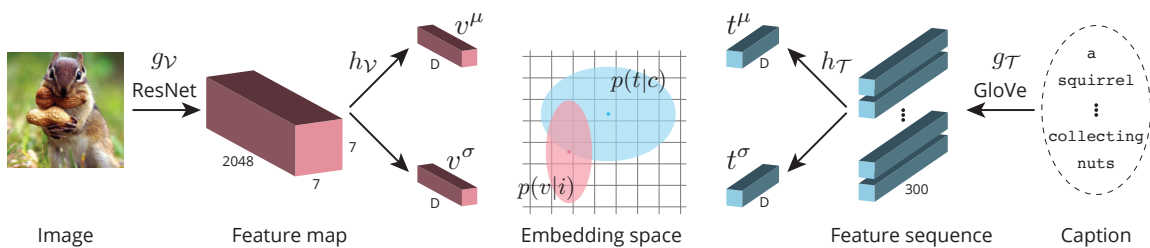
Figure 4.64: High-level overview of representing aleatoric uncertainty in multimodal joint embeddings. The method is quite similar to Hedged Instance Embeddings [150]. Figure taken from [39].

We can train probabilistic representations using the same loss as before, as illustrated in Figure 4.64. We sample K times from each Gaussian, measure Euclidean distances, perform a calibration, use MC integration, and use NLL on top. This trains reasonably well; the authors also ended up with nice embeddings for these joint embeddings.

4.14.4 Theoretical Guarantees for Recovering $P(Z \mid X)$
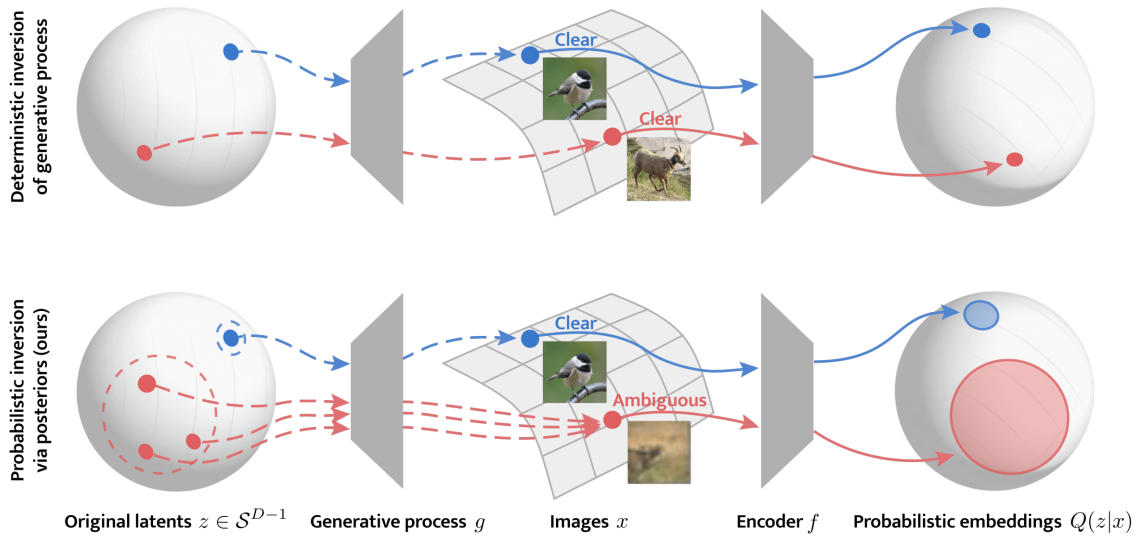
Figure 4.65: Overview of [102]. "Deterministic encoders embed images to points in the latent space. This recovers the latent vectors that generated them, up to a rotation (top). However, if an image is ambiguous, multiple possible latents could have generated it (bottom). An encoder trained with MCInfoNCE correctly recovers this posterior of the generative process, up to a rotation, from contrastive supervision." [102] Figure taken from [102].

Finally, let us discuss some theoretical guarantees for recovering the true latent posterior (up to a rotation and flip) that we hinted at previously. This is a work of Kirchhof et al. on theoretical guarantees for recovering $P(Z \mid X = x)$ [102]. An overview of the method is given in Figure 4.65. Here, we only observe images. Our job is to recover the true latent posteriors $P(Z \mid X = x)$ by only observing which pairs of images are positive and which are not. We assume that images are coming from the true latent space ($P(X \mid Z = z)$, generation).

The used loss function is a probabilistic version of a representation learning objective (InfoNCE), called MCInfoNCE (where MC stands for Monte Carlo):

$$\mathcal{L} := \underset{\substack{x \sim P(x) \\ x^+ \sim P(x^+|x) \\ x_m^- \sim P(x^-), m=1,...,M}}{\mathbb{E}} \left(L_f\left(x, x^+, \{x_m^-\}_{m=1,...,M}\right)\right), \text{ with}$$

$$L_f := -\log \underset{\substack{z \sim Q(z|x) \\ z^+ \sim Q(z^+|x^+) \\ z_m^- \sim Q(z_m^-|x_m^-), m=1,...,M}}{\mathbb{E}} \left(\frac{e^{\kappa_{\text{pos}} z^\top z^+}}{\frac{1}{M} e^{\kappa_{\text{pos}} z^\top z^+} + \frac{1}{M} \sum_{m=1}^{M} e^{\kappa_{\text{pos}} z^\top z_m^-}} \right).$$

This is the loss function we should use to recover the true latent posterior. Of course, there are some assumptions to guarantee the recovery, e.g., the true posterior has to be a von Mises-Fisher distribution, we are sampling from some uniformity, and the positives are determined by the distance from the point.

The theorem that gives the guarantee for the recovery of the true posterior (under some assumptions) is as follows.

Theorem 1 ($\mathcal{L}$ identifies the correct posteriors)**.** *Let $\mathcal{Z} = \mathcal{S}^{D-1}$ and $P(z) = \int P(z|x)dP(x)$ and*

$\int Q(z|x)dP(x)$ *be the uniform distribution over* $\mathcal{Z}$*. Let* g *be a probabilistic generative process [...].* *Let* g *have vMF posteriors* $P(z|x) = \mathrm{vMF}(z; \mu(x), \kappa(x))$ *with* $\mu : \mathcal{X} \to \mathcal{S}^{D-1}$ *and* $\kappa : \mathcal{X} \to \mathbb{R}_{>0}$*. Let* *an encoder* $f(x)$ *parametrize vMF distributions* $\mathrm{vMF}(z; \hat{\mu}(x), \hat{\kappa}(x))$*. Then* $f^* = \arg\min_f \lim_{M \to \infty} \mathcal{L}$ *has the correct posteriors up to a rotation of* $\mathcal{Z}$*, i.e.,* $\hat{\mu}(x) = R\mu(x)$ *and* $\hat{\kappa}(x) = \kappa(x)$*, where* R *is an* *orthogonal rotation matrix,* $\forall x \in \{x \in \mathcal{X} | P(x) > 0\}$*.*

Pay special attention to the $\arg\min$ formulation: This statement is, once more, similar to a strictly proper scoring rule. In summary, it says that the minimizer $Q(Z \mid X)$ of the loss (think maximizer of the score) has true latent posteriors $P(Z \mid X)$ (up to rotation and flip), as visualized in Figure 4.66.
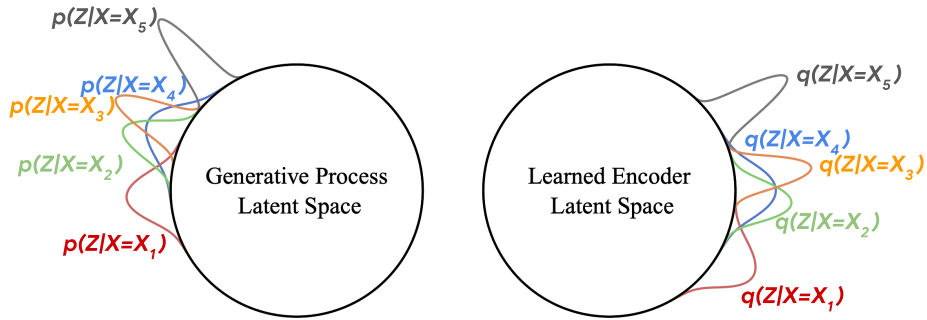
Figure 4.66: Visualization of the degree to which [102] can recover the true embedding posteriors: up to a rotation and a flip. Base figure taken from [102].

Evaluation and Scalability

5 Evaluation and Scalability

5.1 Benchmarks and Evaluation

In this section, we will see common pitfalls of evaluations in trustworthy machine learning.

5.1.1 Why do we do evaluation?

Evaluation enables the ranking of methods. We have a 1D line to put different methods at different positions. We can design new methods that are better than previous ones (w.r.t. the metric) and advance the field. We often compare to prior state-of-the-art (SotA) methods, but comparing a method's performance against human performance often also makes sense. Sometimes there is also a derived theoretical upper bound for performance, either from previous works or our current work. When a model goes over a theoretical upper bound, one has to explain how that is possible. Either there is a bug in the evaluation, the upper bound is flawed, or the model assumes a different set of ingredients than the upper bound.

It is essential to talk about evaluation because it is hard to do it right. There have been many cases in the literature where the evaluation was wrong, and the field had to pay a huge price for that.

5.1.2 What are the costs of wrong evaluation?

(a) The trend according to papers

(b) The trend according to reality
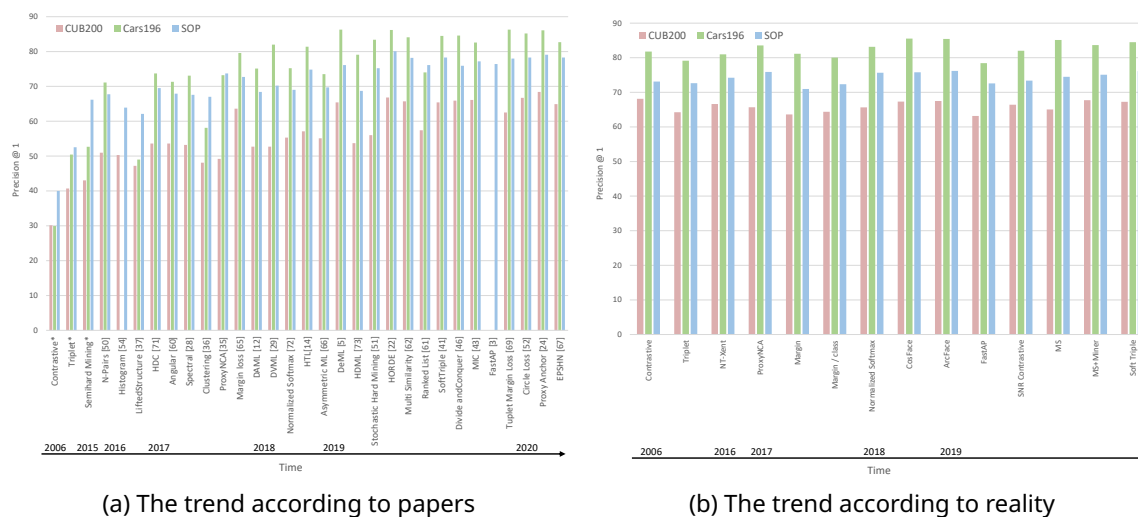
Figure 5.1: "Papers versus Reality: the trend of Precision@1 of various methods over the years. In a), the baseline methods have * next to them, which indicates that their numbers are the average reported accuracy from all papers that included those baselines." [144] In reality (b)), the baseline contrastive loss has comparable performance to all other methods, and no real progress can be seen.

For example, we consider a metric learning benchmark. The "expectation vs. reality" check is

shown in Figure 5.1. We first discuss what the papers claim over time. The colors correspond to different datasets for measuring metric learning performance. The contrastive loss is the starting point of metric learning methods that the later works built upon. This is the standard method we would use to learn a deep metric representation space. Over time, people have developed complicated tricks to improve upon the baselines and the previous year's SotA method. Importantly, *there is a clear upward trend*.

However, the actual reality is *much worse*. The paper unveils many details of the unfair comparisons that lead to the distorted results seen above. In particular, if we tune the hyperparameters super well for contrastive learning and the recent (so-called) SotA methods and consider a fair comparison of them, we get almost the same performance among the methods.

The costs of the wrong evaluation above are severe.

For researchers, 4+ years of effort was put into pursuing the wrong evaluation protocol where we do not unify the set of ingredients among the methods (e.g., the effort put into fine-tuning previous works). We have a false sense of improvement over time. This also translates to opportunity cost: What if they worked on other "real" challenges and were satisfied with contrastive loss instead of working on all these complicated methods?

Practitioners need to select the loss function for their business problems. They waste time looking into all these recent methods, although the most straightforward solution (contrastive loss) probably gives them a good result and requires much less human effort to get it working. This leads to a misinformed selection of methods based on the wrong ranking. They suffer the cost of neglecting a simple solution that works equally well.

Additional Information: Similar "evaluation scandals" in many CV and ML tasks

We consider a list of similar cases in ML where poor evaluation wasted human effort and money. Typically the papers unveiling the problems with the evaluations tend to be very entertaining to read and interesting; thus, we recommend reading them. They can also be very valuable for practitioners who want an unbiased and correct evaluation of methods they can choose from.

- **Face detection**: Mathias et al. "Face Detection without Bells and Whistles" [137]. ECCV'14.
- **Zero-shot learning**: Xian et al. "Zero-Shot Learning – The Good, the Bad and the Ugly" [214]. CVPR'17.
- **Semi-supervised learning**: Oliver et al. "Realistic Evaluation of Deep Semi-Supervised Learning Algorithms" [151]. NeurIPS'18.
- **Unsupervised disentanglement**: Locatello et al. "Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations" [125]. ICML'19.
- **Image classification**: Recht et al. "Do ImageNet Classifiers Generalize to ImageNet?" [163] ICML'19.
- **Scene text recognition**: Baek et al. "What is Wrong with Scene Text Recognition Model Comparisons? Dataset and Model Analysis" [17]. ICCV'19.

- **Weakly-supervised object localization**: Choe et al. "Evaluating Weakly-Supervised Object Localization Methods Right" [36]. CVPR'20.
- **Deep metric learning**: Musgrave et al. "A Metric Learning Reality Check" [144]. ECCV'20.
- **Natural language QA**: Lewis et al. "Question and Answer Test-Train Overlap in Open-Domain Question Answering Datasets" [120]. ArXiv'20.

These papers cover many domains in CV and NLP in general.

5.1.3 "Recipes" for Wrong Benchmark Evaluation

What are the typical patterns in wrong benchmarking/evaluation? We provide an incomplete list of possible failure modes.

Everyone writes their own evaluation metric code.

Even if things are mathematically the same, when it comes to coding, everyone has different ways of handling corner cases. There are non-trivial code-level details in some evaluation metrics. For example, for computing average precision (AP = AUPR), how should we handle precision values for high-confidence bins where the threshold is very high, and thus there are no positive predictions at all? In such cases

$$\text{Precision}(p) = \frac{|\,\text{TP}(p)|}{|\,\text{TP}(p)| + |\,\text{FP}(p)|} = \frac{0}{0}.$$

This is undefined. Some argue that it should be considered 0, some decide to use 1, and others say it should be excluded from the integral computation (for calculating the AP). There must be some agreement on handling such cases in practice. What probably works best in these cases is to have an evaluation server or a library for computing the metrics. That way, we can ensure that all methods use the same implementations of metrics.

Confounding multiple factors when comparing methods.

An example is shown in Figure 5.2. Consider the paper "Sampling Matters in Deep Embedding Learning" [213] They argue in a benchmark that their novel loss function is bringing them gains. But do the improvements really come from the loss function? They do not disclose that the architectures used for training with the respective losses were different. In particular, for training with their loss, they used a more modern architecture (ResNet-50) than for the others (GoogleNet, Inception-BN – archaic). Then it is naturally expected that a Resnet-50 performs better than a GoogleNet. By confounding multiple factors, it becomes hard to rank the losses alone.

Hiding extra resources needed to make improvements.

We mention the work "What Is Wrong With Scene Text Recognition Model Comparisons? Dataset and Model Analysis" [16] If we only care about accuracy, we might be missing the other important axis: computational cost and efficiency. When we only look at an accuracy

Architecture	k	1	10	100	1000	NMI
GoogleNet	Histogram [34]	63.9	81.7	92.2	97.7	-
GoogleNet	Binomial Deviance [34]	65.5	82.3	92.3	97.6	-
Inception-BN	Triplet Semi-hard [25, 29]	66.7	82.4	91.9	-	89.5
Inception-BN	LiftedStruct [22, 29]	62.5	80.8	91.9	-	88.7
Inception-BN	StructClustering [29]	67.0	83.7	93.2	-	89.5
Inception-BN	N-pairs [28]	67.7	83.8	93.0	97.8	88.1
GoogleNet	HDC [41]	69.5	84.4	92.8	97.7	-
ResNet50	**Margin**	**72.7**	**86.2**	**93.8**	**98.0**	**90.7**

Figure 5.2: Example of a scenario where some part of the setup stays hidden (the used architecture) that would clarify an unfair comparison of methods. The first column is not shown in the original paper, although it is very influential of the results. Base figure taken from [213].

plot, we are more inclined to select the method with the highest accuracy. However, if we also considered the inference time (latency) or other computational costs, maybe we would want a different method than the one with the highest accuracy (that is a bit less accurate but much faster).

Training and test samples overlap.

This problem is illustrated in Table 5.1. We consider the paper "Question and Answer Test-Train Overlap in Open-Domain Question Answering Datasets" [120]. In this work, a general problem is highlighted where a fraction of the test sets overlap with the training set for the natural language Q&A task.

Table 5.1: "Fractions of open-domain [question answering] test sets that overlap with their training sets." [120].

Dataset	% Answer overlap	% Question overlap
Natural Questions	63.6	32.5
TriviaQA	71.7	33.6
WebQuestions	57.9	27.5

Sometimes, our evaluation set is contaminated: We see many test samples during training. The Natural Questions, TriviaQA, and WebQuestions datasets are popular benchmarks for the Q&A task in NLP. It turns out that for all three datasets, there is a $> 50\%$ answer overlap (up to 70%!) with the test answers. By memorizing the training answers, it becomes much easier for the model to produce a good answer at test time. Questions are also overlapping quite a bit, as shown in Table 5.2. The authors show that the models solve the task by memorizing rather than generalizing. Many models achieve 0% accuracy for no overlap samples.

Lack of validation set.

This problem is shown in Figure 5.3. Sometimes, there is no published validation set. The CIFAR and ImageNet classification benchmarks lack validation sets. To be precise, ImageNet has a validation but not a test set. Therefore, people are using the validation set as a test set. This brings us many problems. When there is an improvement on the ImageNet validation

Table 5.2: Model performances (exact match scores) in different partitions of the test set. The question/answer overlaps are significant and severe. Table taken from [120].

Model		Open Natural Questions			
		Total	Question Overlap	Answer Overlap Only	No Overlap
Closed book	T5-11B+SSM	36.6	77.2	22.2	9.4
	BART	26.5	67.6	10.2	0.8
Nearest Neighbor	Dense	26.7	69.4	7.0	0.0
	TF-IDF	22.2	56.8	4.1	0.0

set benchmark, it is usually the pointwise samples in the validation set that are addressed rather than the general image classification task. There have been questions like "Are we solving ImageNet or image classification?".

Another question for the same narrative is "Do ImageNet Classifiers Generalize to ImageNet?" [163]. In this case, the design choices and hyperparameter tuning are performed over the test set, spoiling its measure of generalization. There is some evidence that ImageNet classifiers do not generalize to ImageNet, and they are overfitted to the test set. The same holds for CIFAR. The authors of the referenced paper collected another ImageNet validation set, following the same collection procedure. They found that compared to the original ImageNet validation set, the version 2 validation set accuracy is notably lower. If we plot the performances of individual models on the original validation set against the corresponding performances on the version 2 validation set as a scatter plot, it tends to follow a line below the $x = y$ line, indicating that the models do not seem to generalize well. The models' dropping performances on new samples from the same distribution is evidence of overfitting the design choices to the test set over time.
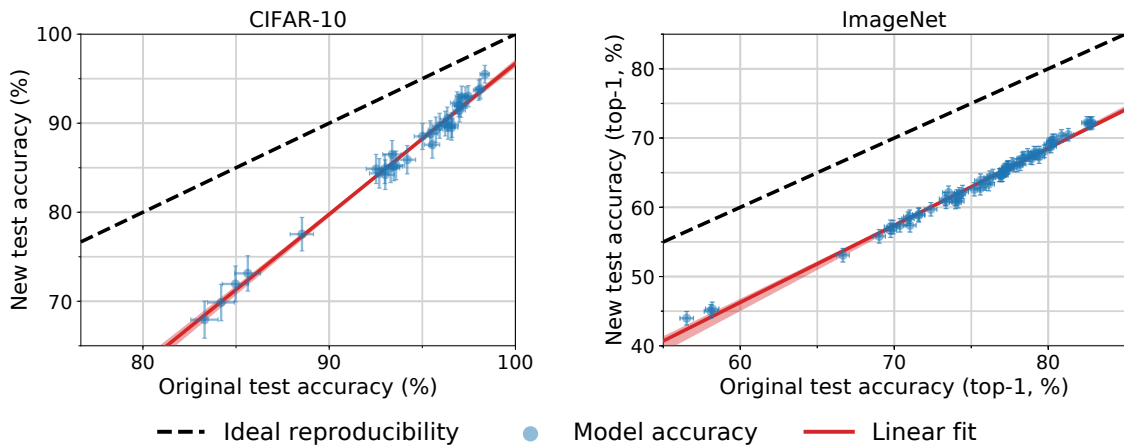
Figure 5.3: Comparison of model accuracy on the original test sets vs. the new test sets collected by the authors. Ideal reproducibility is the line of identity – the performances should not differ at all if the models are not overfit. This is not the case: All models are overfit to both CIFAR-10 and ImageNet. Figure taken from [163].

> **Additional Information: Practical Pointers for Failure Modes**
>
> When we start a new research project in a particular field, how should we find the common failure modes, and how can we avoid them in evaluation?
>
> A good first thing to check is whether there is any paper about the fair comparison of all the methods we are interested in / we want to improve upon.
>
> If not, we have three choices.
> 1. Write such a paper ourselves to "unify all the numbers". This takes the most work, but it can be very rewarding.
> 2. Say that we trust the benchmark because we think there are not so many complicated ingredients involved in the setup; therefore, there is not so much room for the researchers to confound multiple factors during evaluation, e.g., by introducing architectural changes. When the task and the ingredients are both simple, we might want to trust the benchmark.
> 3. Choose to stay skeptical and leave the field until someone performs a trustworthy unified evaluation.

It is crucial to do the evaluation right; otherwise, we are losing much money, time, and research effort. There are currently many domains where this is going sideways, as seen from the list above.

5.2 Scalability

If we look at TML papers, TML is often studied with "toy" datasets. These have the following properties:

- Low-dimensional data ($\leq$ order of 1000 dims per sample).
- Small number of training samples ($\leq$ order of 100k samples).
- Benefit: More extensive and precise labels are available per sample. For example, we have all kinds of attributes labeled for the sample (not just the task label but also other attribute labels like the domain or bias label).
- They make quick evaluation possible.
- Controlled experiments are also possible.
- This kind of dataset accommodates complicated methods with many hyperparameters. Typically we can bring in many of them and tune them the right way to generate the best results here.

Example datasets used in OOD generalization can be seen in Figure 2.19. The real impact, however, comes from results on large-scale datasets. These have the following properties:

- High-dimensional data ($\geq$ order of 10k dims per sample).
- Large number of training samples ($\geq$ order of 1M samples). These days this is not that large-scale either; we can go up to 1B samples if we have the resources.
- High-quality labels are dearer. Often a large portion of our data is even unlabelled or very noisily labelled.

- The validation of an idea on large-scale datasets may take days-weeks. We cannot validate hyperparameter settings super frequently.
- It is hard to analyze the contributing factors. We do not have all the labels we had for the toy dataset, and we also do not have the time and resources to determine which kind of factor contributes to the performance. It is hard to gain knowledge and insights from this kind of data. Therefore, we need some simple methods without many design choices (few hyperparameters).

An example large-scale dataset is the Open Images Dataset (V7), illustrated in Figure 5.4.
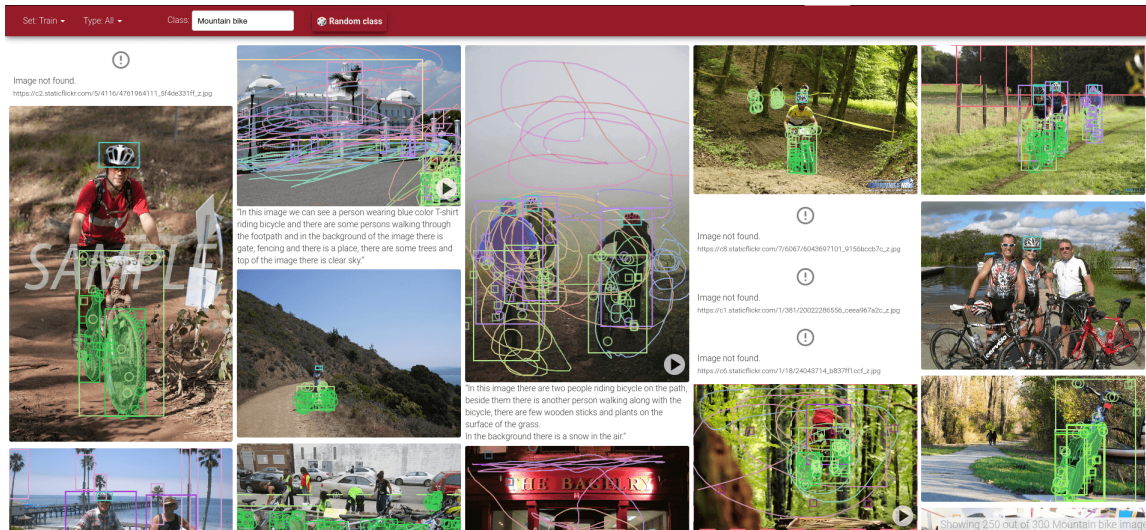
Figure 5.4: Sample of the Open Images Dataset (V7) [108].

5.2.1 Possible Roadmap to Scaling Up TML

There is a possibility to combine toy-ish data and real data to make an impact. This is the method of scaling up from toy data to real data.

First, we work with toy data (e.g., MNIST). Here we have to be creative and propose new ideas (potentially complicated methodology) through quick experiments and tuning hyperparameters. Our goal is to understand why things work.

Based on the insights from toy data and a set of candidate tools, we can go to real data (scaling up). Here we have to identify and remove unnecessary complexities based on knowledge from toy data and aim for something simple. Based on the understanding obtained on toy datasets, we make a good guess about what will work on real data.

The point here is that eventually, to scale up, we need something simple. Of course, going simple is not always easy. We will see some examples where simple wins. There are tons more on arXiv and Twitter.

5.2.2 Simple Wins

Table 5.3: "Average out-of-distribution test accuracies for all algorithms, datasets, and model selection criteria included in the initial release of DomainBed. [...]" [67] Tuned fairly, ERM – the simplest method – is not worse at all than other complicated methods.

	Model selection method: Test-domain validation set *(oracle)*							
Algorithm	**CMNIST**	**RMNIST**	**VLCS**	**PACS**	**Office-Home**	**TerraInc**	**DomainNet**	**Avg**
ERM	58.5 ± 0.3	98.1 ± 0.1	77.8 ± 0.3	87.1 ± 0.3	67.1 ± 0.5	52.7 ± 0.2	41.6 ± 0.1	68.9
IRM	70.2 ± 0.2	97.9 ± 0.0	77.1 ± 0.2	84.6 ± 0.5	67.2 ± 0.8	50.9 ± 0.4	36.0 ± 1.6	69.2
DRO	61.2 ± 0.6	98.1 ± 0.0	77.4 ± 0.6	87.2 ± 0.4	67.7 ± 0.4	53.1 ± 0.5	34.0 ± 0.1	68.4
Mixup	58.4 ± 0.2	98.0 ± 0.0	78.7 ± 0.4	86.4 ± 0.2	68.5 ± 0.5	52.9 ± 0.3	40.3 ± 0.3	69.0
MLDG	58.4 ± 0.2	98.0 ± 0.1	77.8 ± 0.4	86.8 ± 0.2	67.4 ± 0.2	52.4 ± 0.3	42.5 ± 0.1	69.1
CORAL	57.6 ± 0.5	98.2 ± 0.0	77.8 ± 0.1	86.9 ± 0.2	68.6 ± 0.4	52.6 ± 0.6	42.1 ± 0.1	69.1
MMD	63.4 ± 0.7	97.9 ± 0.1	78.0 ± 0.4	87.1 ± 0.5	67.0 ± 0.2	52.7 ± 0.2	39.8 ± 0.7	69.4
DANN	58.3 ± 0.2	97.9 ± 0.0	80.1 ± 0.6	85.4 ± 0.7	65.6 ± 0.3	51.6 ± 0.6	38.3 ± 0.1	68.2
C-DANN	62.0 ± 1.1	97.8 ± 0.1	80.2 ± 0.1	85.7 ± 0.3	65.6 ± 0.3	51.0 ± 1.0	38.9 ± 0.1	68.7

OOD Generalization

Consider addressing the OOD generalization [67] problem and the fair evaluation shown in Figure 5.3. Tuned fairly, ERM – the simplest method – is not worse at all than other complicated methods. (ERM: Training on the whole combined dataset without domain label.) We discussed DRO and DANN in this book.
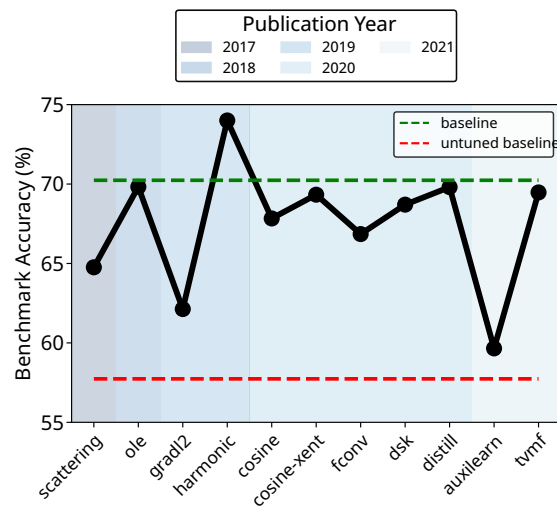
Figure 5.5: "Accuracy of state-of-the-art methods and baselines on the proposed benchmark. The untuned baseline (red dashed line) is trained with default hyper-parameters, i.e., a learning rate of 0.1, and weight decay of 10^{-4}. Conversely, for all other methods, including the baseline (green dashed line), we performed hyper-parameter optimization. Methods are ordered on the x-axis according to their publication year." [25] Seemingly, a lot of human effort went to waste: One could have tuned the baseline instead of coming up with new methods that give almost identical or worse results.

Loss Functions

We have a lot of different variants of loss functions [25] we can use for training a classifier. An interesting contrast between a tuned and an untuned baseline is shown in Figure 5.5. Here, the baseline is vanilla CE. The red line is what papers report as the performance of vanilla CE. They say their method works better than vanilla CE. However, it does not. They just did not tune the baseline properly. We should always do it as well as we possibly can and be

completely clear about our methodology.

Note: It is true that nowadays, reviewers are much more careful with checking evaluation setups than in 2017. However, one can always be not 100% clear about how they performed the evaluation. They can say that they did the tuning, but they can hide the fact that they did so, e.g., using a tiny search window. They can also, e.g., leave out weight decay from the baseline ("Who cares about weight decay..."). Weight decay actually turns out to be quite important. Papers that properly tune hyperparameters for fair comparison tend to recognize the importance of weight decay. We can see a mismatch between what people generally know and what is actually true. By not being 100% descriptive, people can still sneak in papers to conferences by saying they did the hyperparameter search, leaving out subtle and important details (What they did not do right.)

Weakly-Supervised Object Localization

Weakly-Supervised Object Localization [37] is another field that had such a scandal. In Figure 5.6, we see a table adapted from the authors' work. The coverage of the paper's re-evaluation is extensive. CAM has been the simplest method for WSOL for a while now. All the other papers reported better results than CAM in general. However, when we do everything correctly with the same set of ingredients, CAM is the best method on average.

Dataset →	ImageNet (MaxBoxAccV2)				CUB (MaxBoxAccV2)				OpenImages (PxAP)				Total
Architecture →	V	I	R	Mean	V	I	R	Mean	V	I	R	Mean	Mean
CAM CVPR'16	60.0	63.4	63.7	62.4	63.7	56.7	63.0	61.1	58.3	63.2	58.5	60.0	61.2
HaS ICCV'17	60.6	63.7	63.5	62.6	63.7	53.4	64.7	60.6	58.1	58.1	55.9	57.4	60.2
ACoL CVPR'17	57.4	63.7	62.3	61.2	57.4	56.2	66.5	60.0	54.3	57.2	57.3	56.3	59.2
SPG ECCV'18	59.9	63.3	63.3	62.2	56.3	55.9	60.4	57.5	58.3	62.3	56.7	59.1	59.6
ADL CVPR'19	59.8	61.4	63.7	61.7	66.3	58.8	58.4	61.1	58.7	56.8	55.2	56.9	59.9
CutMix ICCV'19	59.4	63.9	63.3	62.2	62.3	57.5	62.8	60.8	58.1	62.5	57.7	59.4	60.9

Figure 5.6: Re-evaluation of various methods on WSOL. CAM, a method from 2016, is still the best when tuned appropriately. Table adapted from [37].

5.2.3 One Right Way to Tune Hyperparameters

The rule of thumb we propose is random search.

1. Set a sensible range of hyperparameters by searching the exponential space. From this heuristic, we could already see that, e.g., 10^{-20} for the learning rate is not sensible at all; no learning happens there. We cut off parts of the parameter space that are not sensible.

2. Once we have this search space defined by sensible ranges of 5-10 hyperparameters, we perform a random search with a fixed number of iterations/samples (between different methods). Random search in practice is excellent. The intuition for that: In practice, not all parameters contribute equally to the performance. In particular, some might not contribute to it at all, only a selection of them. In that case, randomly searching the exponential grid is already good enough because all irrelevant dimensions will not

contribute anyway, and all the search samples in this exponential grid are effectively just searching in the relevant dimensions of the hyperparameter space.

Of course, if many hyperparameters contribute to the final result, we might want to consider more principled techniques, e.g., Bayesian Optimization. However, the authors have never used it for hyperparameter search, only random search.

Suppose the viable hyperparameter regions are super non-convex and very wiggly. In that case, this range-based approach might not work, as we will probably lose a lot of reasonable solutions by cutting the space. We hope this is not the case and the loss is close to unimodal in the space of hyperparameters we are optimizing over. We also assume the independence of the involved hyperparameters. (And that many of these do not matter, actually.)

5.3   Transition from "What" to "How"

Now let us consider future research ideas from the authors. We walk about these from the perspective of going from the "What" to the "How" question we have seen before. In ML 2.0 we learn $P(X, Z, Y)$ from (X, Y) ("What") data. If we look at the ingredients, there is a historical artifact. For ML 1.0, we trained on (X, Y) data, but for ML 2.0, we are still using the same ingredients for solving the derivative problems. Is this right? Is this going to work? We argue that the answer is likely no.

5.3.1   Our Vision
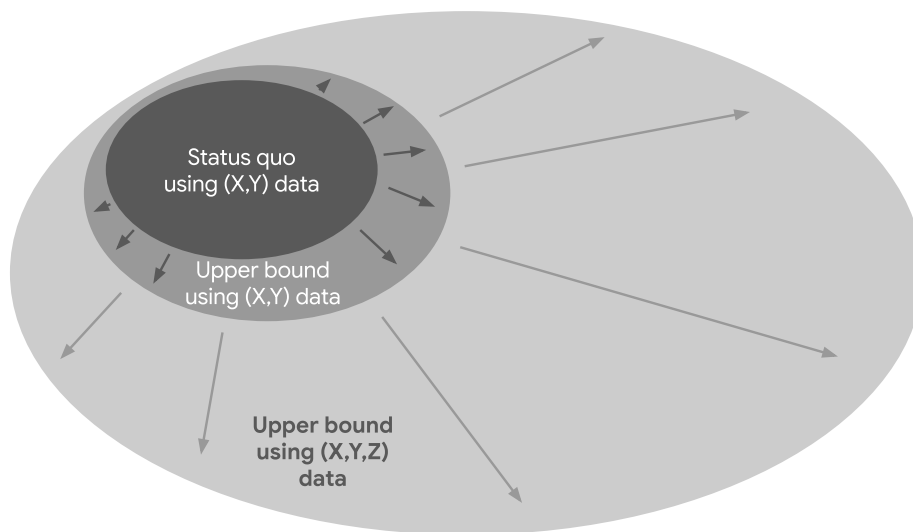
Figure 5.7: Different perspectives of different settings. (X, Y) data correspond to limited ingredients. The upper bound of performance using such data is also very limited. Using (X, Y, Z) opens much broader perspectives.

We can go simple in terms of the method, but what could be really interesting in the future is to collect new types of datasets for scalability and trustworthiness. We discuss Figure 5.7.

The inner two ovals correspond to the *benchmarking* approach. This is the typical approach we have been discussing so far. In a fixed benchmark, everyone uses the same ingredients:

an (X, Y) dataset. If we allow them to use more ingredients, it is no longer a benchmark. It is unfair. (Although that is what people do sometimes still.) The *goal* is to compete for the highest accuracy by using the ingredients most efficiently and smartly. We want to generate the maximal performance from a limited set of ingredients. The *key contribution* is usually the learning algorithm. This approach used to work well for "What" problems. However, we think we should probably use new types of data (new ingredients) that also involve Z. If we discuss with reviewers, we learn that this "learning algorithm contribution + using the same ingredients" is the default mode of thinking for many people.

The outer oval corresponds to the *data hunting* approach. We are still doing some competition, but we are not confined to using the same ingredients. Searching for the ingredient itself is part of the game. When we allow people to use new ingredients, we invite creative new ways to find cheap sources of information that could give us hints about the Z data from all kinds of places. Competitors are allowed to use other ingredients: $(X, Y) + Z$. The *source of value* is the discovery of new, efficient data sources. This is the future of addressing the "How" task. This is also the general research direction the authors of this book want to pursue in the future.

5.3.2   Data as Compressed Human Knowledge

Data are a *compression of human knowledge*. When training an ML model, there are typically two sources of human knowledge.

The first comes from the data and is embedded in the ML model through training. There is a transition of the abstract concept of knowledge in the real world (from annotators) into a dataset in the computational domain. Usually, the dataset contains labels crowdsourced by some annotators.
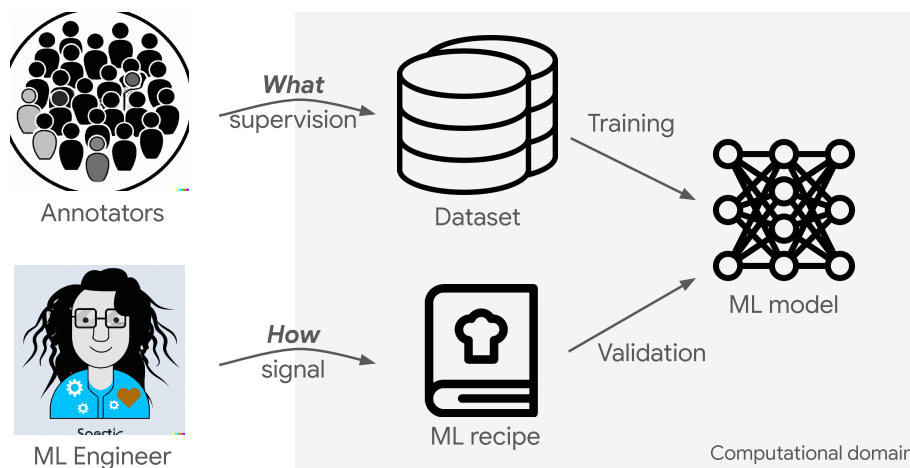
Figure 5.8: The currently dominating types/sources of supervision. Annotators only give "What" supervision through the labeling process. The "How" signal is only supplied by ML engineers through the recipe of creating and training the model.

The second source of human knowledge comes from the validation loop. There is a transition of knowledge of the ML engineer into the recipes for training an ML model that they develop over time. A general overview of this setup is shown in Figure 5.8. This is typically how people are addressing "How" problems now: through the ML engineer's knowledge. Through many

validations, we can find the right setup and design to achieve the "How" tasks. They use the same kind of (X, Y) dataset and rely on the ML engineers to encode business intentions, such as:

- "We need more transparency in the model."
- "We need more robustness."
- "We need better OOD generalization."
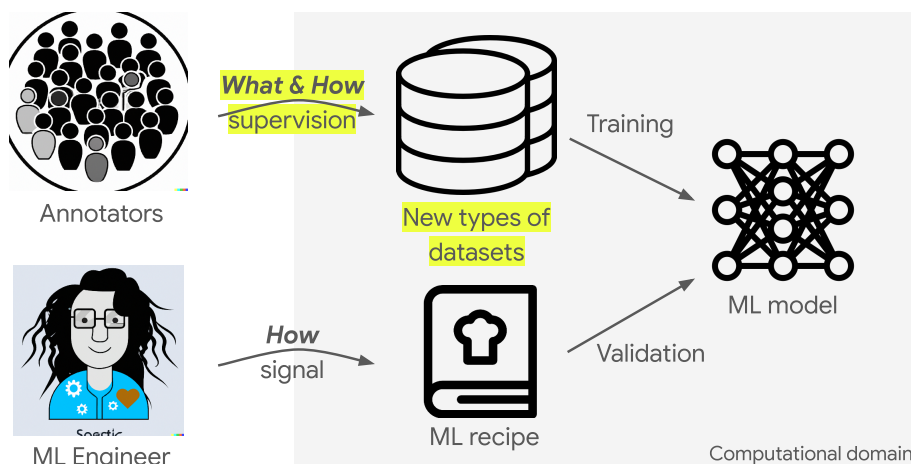
Figure 5.9: A possible future paradigm for types/sources of supervision. Here, the annotators also provide "How" supervision, which can lead to much more robust models.

We argue that in the future, we should also look for methods or datasets for addressing the "How" problem *from the dataset side*. In the future, "How" will probably also be handled through data collection. This is illustrated in Figure 5.9. We wish to not only collect "What" supervision from the annotators but also information related to the "How" task. This way, we obtain a new type of dataset that could be very interesting to the community.

We will now specify two examples of "How" data: We will consider *interventional data* and *additional supervision* on top of our standard annotations. An illustration is given in Figure 5.10.

Figure 5.10: Different types of "How" data. Interventional data specify the "How" aspect by breaking spurious correlations that lead to the incorrect selection of cues. Additional supervision provides explicit new information to specify our needs more thoroughly.

5.3.3 Interventional Data

We will discuss the paper "Weakly Supervised Semantic Segmentation Using Out-of-Distribution Data" [118]. First, let us consider an example of the spurious correlation between trains and rails. If we visualize where our model is looking for the class 'train', we are probably going to get something like in Figure 5.11. The models often look a lot on the rail pixels.

Input image X   Which part is Y?
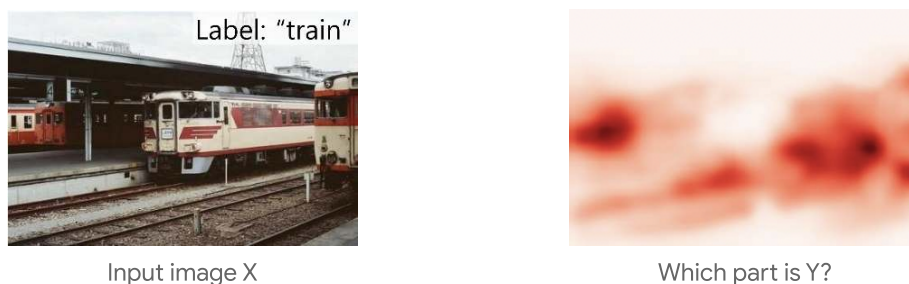
Figure 5.11: Example (input, attribution map) pair that highlights spurious correlations between the label 'train' and the rails. Considering most natural images, the model can get away with looking at the rails because it is quite uncommon to see a train without rails. However, this choice of cue is misspecified and does not lead to robust generalization. Base figure taken from [118].

This is a well-known problem. The reason is that if we collect data naturally arising from the way people take pictures, then we will probably see many images where the trains are on the rails. Models can recognize trains based on rails, leveraging spurious correlation. The learned "How" by the model is wrong. The existence of spurious correlations already indicates that interventional data does not arise naturally in natural data. (If they arose naturally, they would have been part of the training data already, and there would not be any spurious correlation at all.) However, we did not encode "How" requirements in our dataset. *Therefore, we cannot expect our model to get it right.*

Hard-negative images as new p(X)

Figure 5.12: Example hard-negative samples w.r.t. train samples, containing no trains but still including rails. Base Figures are taken from [118].

One way to combat the problem of spurious correlation between rail and train is to introduce interventional data. If we are more cautious when collecting data, we can also collect hard-negative images (rail with no train). This is illustrated in Figure 5.12. Hard-negative images are (1) hard because they target spurious correlations, so the models employing such spurious correlations will get them wrong; and (2) negative because there is no train in the image. Here, we explicitly target the possible bias – we eliminate "rail" from a plausible set of cues for detecting trains.

More examples of interventional data are shown in Figure 5.13. Of course, as discussed, this kind of data does not arise very often naturally; thus, there should be a way to go and find them. We need a data crawling mechanism that supplies such examples.

Efficiently Collecting an Interventional Dataset

Our task is to find hard-negative samples for the 'train' class in our running example. We can do this as illustrated in Figure 5.14. The candidate dataset does not fully have to be OOD; it just has to be a large dataset of images. Of course, the more purely OOD the original dataset

Training data — Label: Car

Additional **hard-negative** training data — Label: **Not** a car.

Label: Frog

Label: **Not** a frog.
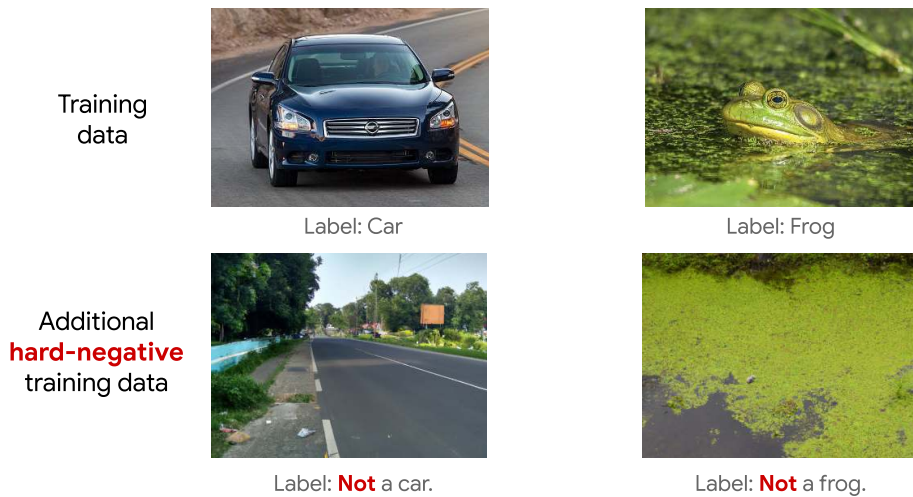
Figure 5.13: Addition of hard-negative samples w.r.t. 'car' and 'frog', containing their corresponding biases but not the objects themselves.

Candidate OoD set

Compute $p(\text{train})$ and sort.

0.98   0.97   0.89   0.85   0.63

0.47   0.38   0.18   0.08   0.03

Remove images with $p(\text{train})<0.5$.

Pruned set

0.98   0.97   0.89   0.85   0.63

Manually filter out images containing train.
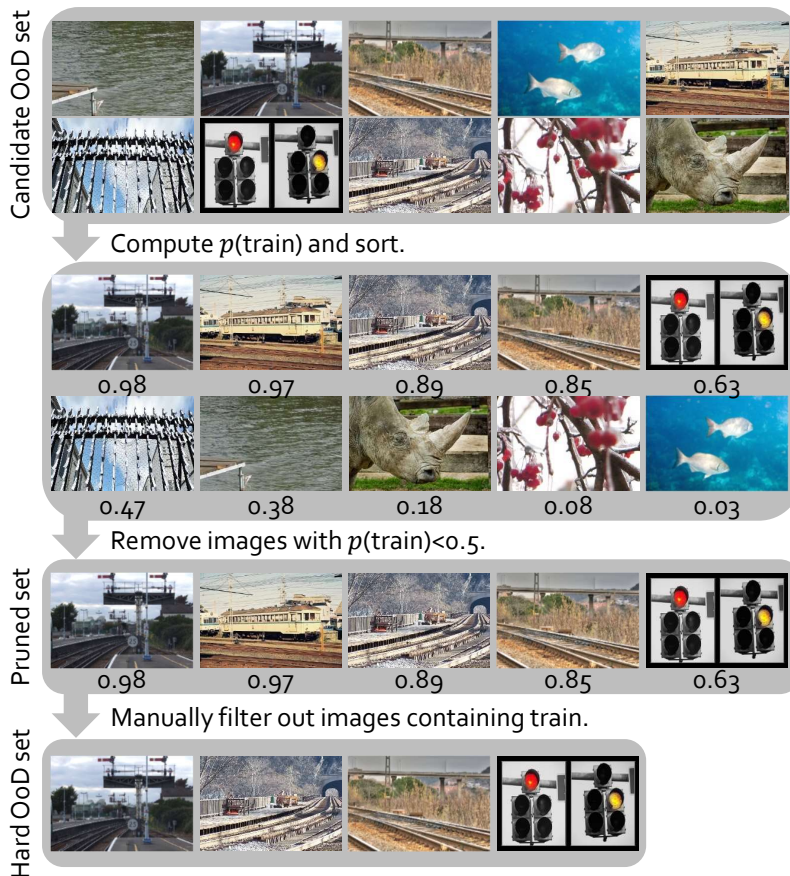
Hard OoD set

Figure 5.14: Possible procedure of collecting hard OOD samples. Figure taken from [118].

is, the more efficiently we can collect relevant data from it. We compute $p(\text{train})$ by running all images in the dataset through our classifier. We only keep images with a 'train' score above a certain threshold. This set will not look as nice as in the figure in practice. There will be a lot of true positives as well. For manual filtering, we need human labor (HITL). Humans are the sources of hard-negative knowledge. There is no solution to finding a clarification of such

spurious correlations without requiring human knowledge. The name "hard OOD dataset" is equivalent to "hard false positive dataset" and also to "hard negative dataset."

We need to minimize our costs whenever we use human labor because it is expensive. The cost depends on two dimensions:

1. **How long does it take a human to remove all the true positive images?** This is very cheap, as the annotators do not have to draw a bounding box/segmentation map or classify the image into 1k classes. It is an easy binary decision (Y/N).

2. **How many hard-negative images are needed per class?** Not many at all. Considering only one image per class, the mIoU (which is a measure for telling how much spurious correlation we have, higher is better) with foreground increases by 2%. Using 100 images per class, the mIoU with foreground increases by 3%. We have poor mIoU without any hard negatives. However, one hard-negative image per class already helps a lot, apparently. If we take more hard negatives per class, we get diminishing returns as the mIoU performance saturates.

Interventional data are, therefore, a cheap source of "How" information. For the Pascal VOC dataset with 20 classes, we only need 20 new hard-negative samples to improve mIoU quite a bit. This is a low-hanging fruit for new types of data.
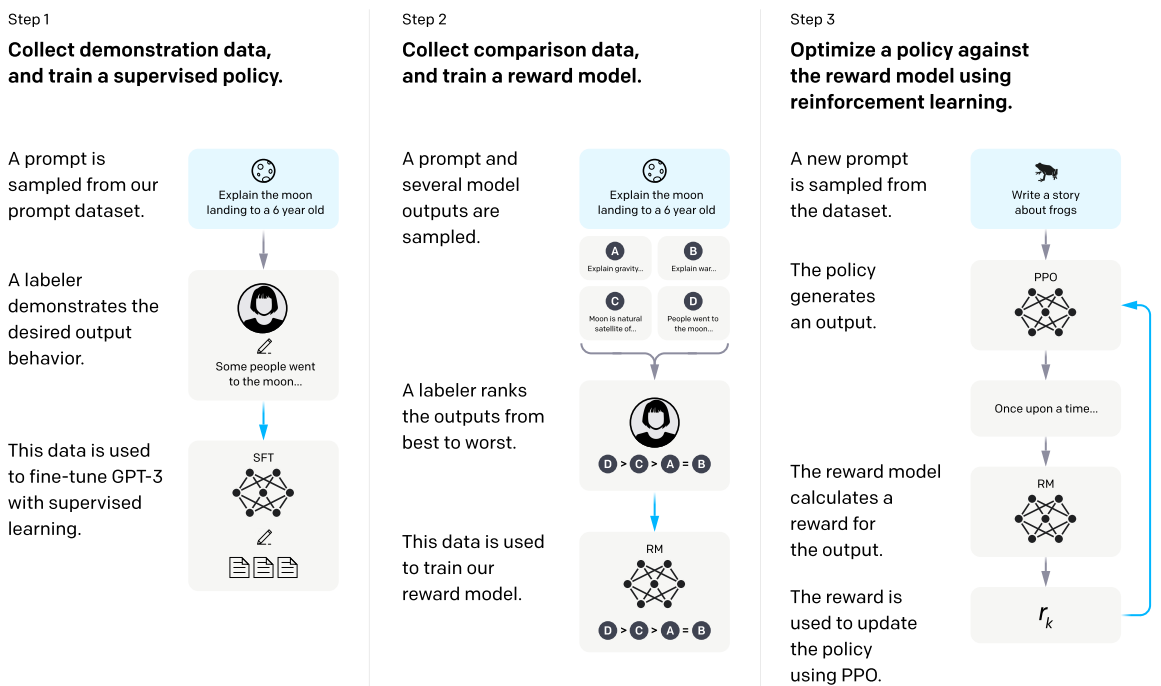
Figure 5.15: Three-step fine-tuning procedure of ChatGPT. HITL is crucial for aligned ML models. Figure taken from [153].

Interventional data collection is gaining momentum now. One example is ChatGPT's fine-tuning, illustrated in Figure 5.15. The research field seems to return to the HITL paradigm. This is good because it is the only way to solve this problem. HITL is used for both InstructGPT and ChatGPT. These improve upon the original GPT-3 in terms of the safety features precisely because they also use HITL to fine-tune the models further. Researchers developing these systems know that humans are the ultimate source of the "How" information. We are also shifting the distribution (data or output) a little bit to what humans would consider more

appropriate/relevant as answers during a chat. This introduces an intervention in the data generation process; we use a novel data source for further training. And this is what matters: There are all kinds of issues around LLMs, like inappropriate outputs and jailbreak (making LLMs output inappropriate things). Humans can teach LLMs "how to behave." Instead of web crawling, one can use humans to generate samples. This improves trustworthiness considerably. On the left of Figure 5.15, humans are used to generate possible answers to questions. This is quite labor-heavy. On the right, humans are only used to rank the outputs of models based on their preferences. This ranking can be used for further fine-tuning with RLHF. This is less labor-heavy and is quite scalable.

5.3.4 Introducing Additional Supervision

ImageNet annotation is performed as follows. First, annotators receive an object category or concept at the top of a webpage. Then they have to click on images containing the concept. Some images from the candidate image set are selected, and some are not. (This is already a pre-filtered set of images that might correspond to the concept.) When we do this, we obtain a set of images for every selected concept. These are then used for training the model.


```
"imageID": "n01440764/n01440764_105",
"originalImageHeight": 375,
"originalImageWidth": 500,
"imageHeight": 243,
"imageWidth": 243,
"imagePosition": {"x": 857, "y": 1976},
"hoveredRecord": [
  {"action": "enter", "time": 1641425051975},
  {"action": "leave", "time": 1641425052740}
],
"mouseTracking": [
  {"x": 0.0030864197530864196, "y": 0.629629629629629, "time": 1641425051975},
  {"x": 0.44135802469135804, "y": 0.6008230452674898, "time": 1641425052027}
],
"selected": true,
"selectedRecord": [
  {"x": 0.5401234567901234, "y": 0.4732510288065844, "time": 1641425052319}
],
"worker_id": "47DBDD543E"
```

Original annotation   Annotation byproducts (Z)

Figure 5.16: Additional, potentially useful meta-data from the annotation procedure. *Blue:* Original annotation ImageNet data collectors have considered so far. This is wasting a ton of auxiliary supervision. *Red:* The annotation byproducts may be irrelevant and noisy, but we should not throw them away, as they can also be informative. We want to use them to improve our model (e.g., by obtaining new ingredients for uncertainty estimation).

However, the action of annotation also contains valuable information in terms of the mouse track, click location, time annotators took between clicks, the full time needed to go through the set of images, and many other factors. We can efficiently collect additional supervision. This is shown in Figure 5.16. Annotation byproducts can be leveraged in several ways. The work of Han et al. [72] gives a thorough demonstration of how these can be used along with task supervision.

There are also a lot of parallel efforts from other groups to obtain additional supervision. One is OpenImages Localized Narratives, illustrated in Figure 5.17. This is way more information
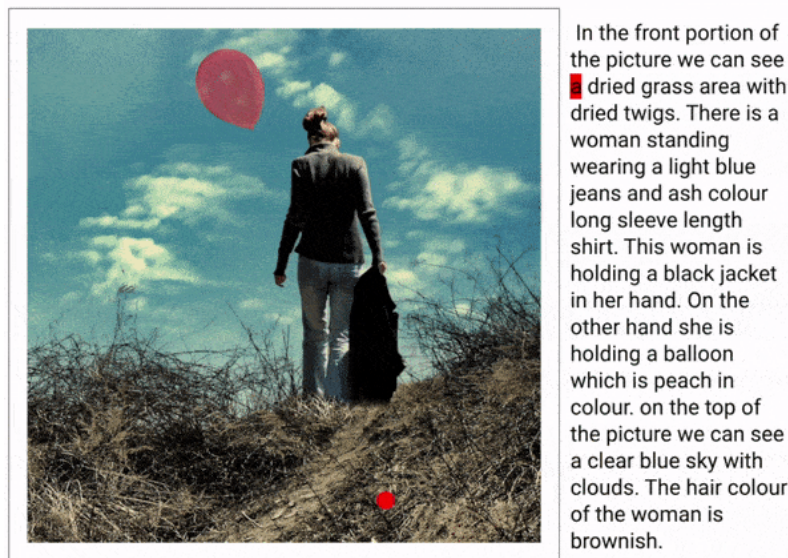
Figure 5.17: Labeling process of OpenImages Localized Narratives. They try to collect as much information as possible for every image. The human annotator speaks out what they see in the image. As they describe every object, they need to hover over the image part they are talking about. For every word, we have a corresponding location in the image. They record the mouse trajectory and voice (1-to-1 correspondence between what they say and what they point to). Then, the voice recording is transcribed into text. This results in huge captions compared to COCO. Figure taken from [157].

and supervision compared to traditional image captioning or object localization datasets. Multimodal annotations are rich in the "How" information content in general. The annotation contains much new information. We should consider how to best exploit this information for the "How" problems. There is not much research on this yet; we are fortunate to work on this now and make an impact.

5.3.5 Method-centric vs. Dataset-centric Solutions

There are two general ways to solve problems, both with pros and cons, detailed below.

Method-centric solutions. These have cheap initial costs. One can use existing benchmarks and training sets. One just needs to devise a clever new method (e.g., loss, architecture, optimizer, regularizer). Typically we end up with highly complex methods because all simple methods have been tried out already. For these complicated methods, we need a lot of computational resources and human brain time. This potentially has enormous costs. We usually do not consider brain time cost as much compared to, e.g., annotation cost. Development is expensive and requires many runs to validate hyperparameters. Furthermore, such solutions are upper-limited by the information cap defined by the benchmark. (What supervision do we have?) As such, scaling up often fails. The complex tricks do not work anymore. We need a lot of effort and experiments to prune down the method into something simpler that scales well.

Dataset-centric solutions. These are relatively new and exciting approaches. It has large initial costs: 10k - 10M EUR for a large-scale dataset. A few thousand might be enough for a small-scale dataset, but it is meaningful when our budget goes up to 100k EUR. (This way,

we can obtain a larger dataset and/or better supervision.) Once built, it brings huge utility to the public. (Everyone can use it to create new methods; it has a huge impact.) One could also expect good transferability of pre-trained models to other tasks. (We can pre-train a model on the dataset and open source it: this is also a huge contribution to the field. They can just download it without needing to train it from scratch.) Notably, there is no information cap (only creativity cap and budget cap). If we have more information available, the method itself can be quite simple. We can just use a vanilla loss/architecture, which will often work best (we have seen that simple methods often work best). Such methods also scale better and are easier to use, as they come with fewer hyperparameters usually.

We think that simple methods with new kinds of data will bring us the biggest gain in the future.

ChatGPT proposed the following closing statement for the book: "Let us harness the power of machine learning to make a difference. Let us make an impact through machine learning." We could not agree more.

Calculus Refresher

A Calculus Refresher

We consider a couple of exercises for calculating partial derivates w.r.t. vectors and matrices. One particularly useful object for calculating partial derivatives is the Kronecker delta, a function of two variables.

> **Definition A.1: Kronecker Delta**
>
> $$\delta_{ij} = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j. \end{cases}$$

Matrix multiplication, which is also important to be able to solve the exercises later, is defined as follows.

> **Definition A.2: Matrix Multiplication**
>
> Let $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{n \times p}, C \in \mathbb{R}^{m \times p}$.
>
> $$C = AB \iff C_{ij} = \sum_{k=1}^{n} A_{ik} B_{kj} \; \forall i \in \{1, \ldots, m\}, j \in \{1, \ldots, p\}.$$

Let us consider the gradient operator.

> **Definition A.3: Gradient**
>
> Let $f : \mathbb{R}^n \to \mathbb{R}$. Then
> $$\nabla f : \mathbb{R}^n \to \mathbb{R}^n, (\nabla f)_i = \frac{\partial f}{\partial x_i}.$$
>
> Sometimes, the argument is explicit: $\nabla_x f$. In $\nabla_x f(x)$, the subscript indicates "which variable" and the argument indicates "where to evaluate".
>
> **Example**: $f : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^l \to \mathbb{R}$. Then $\nabla_z f : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^l \to \mathbb{R}^l$. We often abuse notation and use the variable name to indicate the position of the argument w.r.t. which we take the gradient. Often this is clear from the context.
>
> The following notations are questionable. They are both abusing the abuse of notation.
> - $\nabla_z f(x + z, z^2, y)$. This notation is unclear. According to general use, the z in the subscript should refer to the position. However, we also have an explicit variable z that can be confusing. One should either use different symbols as the arguments and/or one should write everything down nicely using partial derivatives. Combining the two, one might first declare that f is a function of variables (placeholders)

x', y', and z', and then write

$$\left.\frac{\partial f}{\partial x}\right|_{x'=x+z,\,y'=z^2,\,z'=y}.$$

- $\nabla_{x+y+z} f(x+y+z, x)$. This notation is incorrect. One should use the subscript to refer to the position, and again, either use different symbols as the arguments or write everything down using partial derivatives.

Lastly, we present a simple rule for taking partial derivatives of a tensor element w.r.t. another tensor element.

Definition A.4: Derivative of a Tensor Element w.r.t. Another Element

$$\frac{\partial v_{i_1,\dots,i_n}}{\partial v_{j_1,\dots,j_n}} = \prod_{k=1}^{n} \delta_{i_k j_k}.$$

We are now ready to solve the first exercise.

Task: Gradient of Squared L_2 Norm

Show that for $x \in \mathbb{R}^n$,

$$\nabla_x \|x\|^2 = 2x.$$

$\forall i \in \{1,\dots,n\}$:

$$\left(\nabla_x \|x\|^2\right)_i = \frac{\partial}{\partial x_i} \sum_{j=1}^{n} x_j^2$$

$$= \sum_{j=1}^{n} \frac{\partial}{\partial x_i} x_j^2$$

$$= \sum_{j=1}^{n} \delta_{ij} 2x_j$$

$$= 2x_i.$$

Let us define the trace operator for real matrices.

Definition A.5: Trace

The trace of a square matrix $A \in \mathbb{R}^{n \times n}, n \in \mathbb{N}$ is defined as

$$\mathrm{tr}(A) = \sum_{i=1}^{n} a_{ii}.$$

The second exercise is as follows.

> **Task: Gradient of Trace of Matrix Multiplication**
>
> Show that for $A \in \mathbb{R}^{n \times m}, B \in \mathbb{R}^{m \times n}$,
>
> $$\nabla_A \operatorname{tr}(AB) = B^\top.$$

$\forall i, j \in \{1, \ldots, n\}$:

$$
\begin{aligned}
(\nabla_A \operatorname{tr}(AB))_{ij} &= \frac{\partial}{\partial A_{ij}} \sum_{p=1}^{n} (AB)_{pp} \\
&= \frac{\partial}{\partial A_{ij}} \sum_{p=1}^{n} \sum_{q=1}^{m} A_{pq} B_{qp} \\
&= \sum_{p=1}^{n} \sum_{q=1}^{m} \frac{\partial}{\partial A_{ij}} A_{pq} B_{qp} \\
&= \sum_{p=1}^{n} \sum_{q=1}^{m} \delta_{ip} \delta_{jq} B_{qp} \\
&= B_{ji}.
\end{aligned}
$$

Our last exercise is to compute the gradient of a quadratic form.

> **Task: Gradient of Quadratic Form**
>
> Show that for $x \in \mathbb{R}^n, A \in \mathbb{R}^{n \times n}$,
>
> $$\nabla_x x^\top A x = (A + A^\top) x.$$

$\forall i \in \{1, \ldots, n\}$:

$$
\begin{aligned}
\left(\nabla_x x^\top A x\right)_i &= \frac{\partial}{\partial x_i} \sum_{p,q=1}^{n} x_p A_{pq} x_q \\
&= \sum_{p,q=1}^{n} \frac{\partial}{\partial x_i} (x_p A_{pq} x_q) \\
&= \sum_{p,q=1}^{n} \delta_{ip} A_{pq} x_q + \sum_{p,q=1}^{n} \delta_{iq} x_p A_{pq} \\
&= \sum_{q=1}^{n} A_{iq} x_q + \sum_{p=1}^{n} x_i A_{iq} \\
&= (Ax)_i + (A^\top x)_i \\
&= ((A + A^\top) x)_i.
\end{aligned}
$$

List of Definitions

List of Definitions

Bibliography

Bibliography

[1] Radhakrishna Achanta et al. "SLIC Superpixels Compared to State-of-the-Art Super-pixel Methods". In: *IEEE transactions on pattern analysis and machine intelligence* 34 (May 2012). DOI: 10.1109/TPAMI.2012.120.

[2] Julius Adebayo et al. *Sanity Checks for Saliency Maps*. 2018. DOI: 10.48550/ARXIV.1810.03292. URL: https://arxiv.org/abs/1810.03292.

[3] Naman Agarwal, Brian Bullins, and Elad Hazan. *Second-Order Stochastic Optimization for Machine Learning in Linear Time*. 2017. arXiv: 1602.03943 [stat.ML].

[4] Evidently AI. *Machine Learning Monitoring, Part 5: Why You Should Care About Data and Concept Drift*. URL: https://www.evidentlyai.com/blog/machine-learning-monitoring-data-and-concept-drift (visited on 02/14/2023).

[5] Ekin Akyürek et al. *Towards Tracing Factual Knowledge in Language Models Back to the Training Data*. 2022. DOI: 10.48550/ARXIV.2205.11482. URL: https://arxiv.org/abs/2205.11482.

[6] Michael A. Alcorn et al. *Strike (with) a Pose: Neural Networks Are Easily Fooled by Strange Poses of Familiar Objects*. 2018. DOI: 10.48550/ARXIV.1811.11553. URL: https://arxiv.org/abs/1811.11553.

[7] Ameen Ali et al. *XAI for Transformers: Better Explanations through Conservative Propagation*. 2022. DOI: 10.48550/ARXIV.2202.07304. URL: https://arxiv.org/abs/2202.07304.

[8] David Alvarez-Melis and Tommi S. Jaakkola. *On the Robustness of Interpretability Methods*. 2018. arXiv: 1806.08049 [cs.LG].

[9] Joost van Amersfoort et al. *Uncertainty Estimation Using a Single Deep Deterministic Neural Network*. 2020. DOI: 10.48550/ARXIV.2003.02037. URL: https://arxiv.org/abs/2003.02037.

[10] Jacob Andreas. *Measuring Compositionality in Representation Learning*. 2019. arXiv: 1902.07181 [cs.LG].

[11] Marcin Andrychowicz et al. *Learning to learn by gradient descent by gradient descent*. 2016. DOI: 10.48550/ARXIV.1606.04474. URL: https://arxiv.org/abs/1606.04474.

[12] Martin Arjovsky et al. *Invariant Risk Minimization*. 2019. DOI: 10.48550/ARXIV.1907.02893. URL: https://arxiv.org/abs/1907.02893.

[13] Anish Athalye, Nicholas Carlini, and David Wagner. *Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples*. 2018. DOI: 10.48550/ARXIV.1802.00420. URL: https://arxiv.org/abs/1802.00420.

[14] Peter M. Attia et al. "Closed-loop optimization of fast-charging protocols for batteries with machine learning". In: *Nature* 578 (2020), pp. 397–402.

[15] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. *Layer Normalization*. 2016. DOI: 10.48550/ARXIV.1607.06450. URL: https://arxiv.org/abs/1607.06450.

[16] Jeonghun Baek et al. *What Is Wrong With Scene Text Recognition Model Comparisons? Dataset and Model Analysis*. 2019. DOI: 10.48550/ARXIV.1904.01906. URL: https://arxiv.org/abs/1904.01906.

[17] Jeonghun Baek et al. "What is wrong with scene text recognition model comparisons? dataset and model analysis". In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2019, pp. 4715–4723.

[18] Hyojin Bahng et al. *Learning De-biased Representations with Biased Representations*. 2019. DOI: 10.48550/ARXIV.1910.02806. URL: https://arxiv.org/abs/1910.02806.

[19] European Central Bank. *Transmission mechanism of monetary policy*. URL: https://www.ecb.europa.eu/mopo/intro/transmission/html/index.en.html (visited on 04/01/2023).

[20] Manel Baradad et al. *Learning to See by Looking at Noise*. 2021. DOI: 10.48550/ARXIV.2106.05963. URL: https://arxiv.org/abs/2106.05963.

[21] Vaishak Belle and Ioannis Papantonis. "Principles and practice of explainable machine learning". In: *Frontiers in big Data* (2021), p. 39.

[22] Alexander Binder et al. *Layer-wise Relevance Propagation for Neural Networks with Local Renormalization Layers*. 2016. DOI: 10.48550/ARXIV.1604.00825. URL: https://arxiv.org/abs/1604.00825.

[23] Or Biran and Courtenay Cotton. "Explanation and justification in machine learning: A survey". In: *IJCAI-17 workshop on explainable AI (XAI)*. Vol. 8. 1. 2017, pp. 8–13.

[24] Wieland Brendel and Matthias Bethge. *Approximating CNNs with Bag-of-local-Features models works surprisingly well on ImageNet*. 2019. DOI: 10.48550/ARXIV.1904.00760. URL: https://arxiv.org/abs/1904.00760.

[25] Lorenzo Brigato et al. "Image Classification With Small Datasets: Overview and Benchmark". In: *IEEE Access* 10 (2022), pp. 49233–49250. DOI: 10.1109/access.2022.3172939. URL: https://doi.org/10.1109/access.2022.3172939.

[26] Tom B. Brown et al. *Language Models are Few-Shot Learners*. 2020. DOI: 10.48550/ARXIV.2005.14165. URL: https://arxiv.org/abs/2005.14165.

[27] Jacob Buckman et al. "Thermometer Encoding: One Hot Way To Resist Adversarial Examples". In: *International Conference on Learning Representations*. 2018. URL: https://openreview.net/forum?id=S18Su--CW.

[28] Remi Cadene et al. "RUBi: Reducing Unimodal Biases in Visual Question Answering". In: (2019). DOI: 10.48550/ARXIV.1906.10169. URL: https://arxiv.org/abs/1906.10169.

[29] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. "COCO-Stuff: Thing and stuff classes in context". In: *Computer vision and pattern recognition (CVPR), 2018 IEEE conference on*. IEEE. 2018.

[30] Nicholas Carlini and David Wagner. *Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods*. 2017. DOI: 10.48550/ARXIV.1705.07263. URL: https://arxiv.org/abs/1705.07263.

[31] Nicholas Carlini and David Wagner. *Towards Evaluating the Robustness of Neural Networks*. 2016. DOI: 10.48550/ARXIV.1608.04644. URL: https://arxiv.org/abs/1608.04644.

[32] João Carreira et al. "A Short Note on the Kinetics-700 Human Action Dataset". In: *CoRR* abs/1907.06987 (2019). arXiv: 1907.06987. URL: http://arxiv.org/abs/1907.06987.

[33] AI Addict YouTube Channel. *Self Driving Collision (Analysis)*. URL: https://www.youtube.com/watch?v=Zl9rM8D3k34&list=LL (visited on 02/14/2023).

[34] Liang-Chieh Chen et al. *DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs*. 2016. DOI: 10.48550/ARXIV.1606.00915. URL: https://arxiv.org/abs/1606.00915.

[35] Pin-Yu Chen et al. "ZOO". In: *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*. ACM, Nov. 2017. DOI: 10.1145/3128572.3140448. URL: https://doi.org/10.1145%2F3128572.3140448.

[36] Junsuk Choe et al. "Evaluating weakly supervised object localization methods right". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 3133–3142.

[37] Junsuk Choe et al. *Evaluation for Weakly Supervised Object Localization: Protocol, Metrics, and Datasets*. 2020. DOI: 10.48550/ARXIV.2007.04178. URL: https://arxiv.org/abs/2007.04178.

[38] Jinwoo Choi et al. *Why Can't I Dance in the Mall? Learning to Mitigate Scene Bias in Action Recognition*. 2019. DOI: 10.48550/ARXIV.1912.05534. URL: https://arxiv.org/abs/1912.05534.

[39] Sanghyuk Chun et al. *Probabilistic Embeddings for Cross-Modal Retrieval*. 2021. DOI: 10.48550/ARXIV.2101.05068. URL: https://arxiv.org/abs/2101.05068.

[40] Julien Colin et al. *What I Cannot Predict, I Do Not Understand: A Human-Centered Evaluation Framework for Explainability Methods*. 2021. DOI: 10.48550/ARXIV.2112.04417. URL: https://arxiv.org/abs/2112.04417.

[41] Wikimedia Commons. *File:Roc curve.svg*. URL: https://commons.wikimedia.org/wiki/File:Roc_curve.svg (visited on 03/02/2023).

[42] Alexander D'Amour et al. *Underspecification Presents Challenges for Credibility in Modern Machine Learning*. 2020. DOI: 10.48550/ARXIV.2011.03395. URL: https://arxiv.org/abs/2011.03395.

[43] Jia Deng et al. "ImageNet: A large-scale hierarchical image database". In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 2009, pp. 248–255. DOI: 10.1109/CVPR.2009.5206848.

[44] Jia Deng et al. "Imagenet: A large-scale hierarchical image database". In: *2009 IEEE conference on computer vision and pattern recognition*. Ieee. 2009, pp. 248–255.

[45] Devopedia For Developers By Developers. *Word Embedding*. URL: https://devopedia.org/word-embedding (visited on 03/02/2023).

[46] Guneet S. Dhillon et al. "Stochastic activation pruning for robust adversarial defense". In: *International Conference on Learning Representations*. 2018. URL: https://openreview.net/forum?id=H1uR4GZRZ.

[47] Finale Doshi-Velez and Been Kim. *Towards A Rigorous Science of Interpretable Machine Learning*. 2017. DOI: 10.48550/ARXIV.1702.08608. URL: https://arxiv.org/abs/1702.08608.

[48] Alexey Dosovitskiy et al. *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*. 2020. DOI: 10.48550/ARXIV.2010.11929. URL: https://arxiv.org/abs/2010.11929.

[49] Upol Ehsan et al. "Expanding explainability: Towards social transparency in ai systems". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 2021, pp. 1–19.

[50] M. Everingham et al. "The Pascal Visual Object Classes (VOC) Challenge". In: *International Journal of Computer Vision* 88.2 (June 2010), pp. 303–338.

[51] Michael Firman et al. *DiverseNet: When One Right Answer is not Enough*. 2020. DOI: 10. 48550/ARXIV.2008.10634. URL: https://arxiv.org/abs/2008.10634.

[52] Lukas Fluri. *On Calibration of Modern Neural Networks*. URL: https://fernandoperezc.github.io/Advanced-Topics-in-Machine-Learning-and-Data-Science/Fluri.pdf (visited on 03/02/2023).

[53] Ruth C. Fong and Andrea Vedaldi. "Interpretable Explanations of Black Boxes by Meaningful Perturbation". In: *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE, Oct. 2017. DOI: 10.1109/iccv.2017.371. URL: https://doi.org/10.1109/iccv.2017.371.

[54] Yarin Gal and Zoubin Ghahramani. *Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning*. 2015. DOI: 10.48550/ARXIV.1506.02142. URL: https://arxiv.org/abs/1506.02142.

[55] Ido Galil, Mohammed Dabbah, and Ran El-Yaniv. *What Can We Learn From The Selective Prediction And Uncertainty Estimation Performance Of 523 Imagenet Classifiers*. 2023. arXiv: 2302.11874 [cs.LG].

[56] Yaroslav Ganin et al. "Domain-Adversarial Training of Neural Networks". In: (2015). DOI: 10.48550/ARXIV.1505.07818. URL: https://arxiv.org/abs/1505.07818.

[57] Timur Garipov et al. *Loss Surfaces, Mode Connectivity, and Fast Ensembling of DNNs*. 2018. DOI: 10.48550/ARXIV.1802.10026. URL: https://arxiv.org/abs/1802.10026.

[58] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. *A Neural Algorithm of Artistic Style*. 2015. DOI: 10.48550/ARXIV.1508.06576. URL: https://arxiv.org/abs/1508.06576.

[59] Robert Geirhos et al. "Don't trust your eyes: on the (un) reliability of feature visualizations". In: *arXiv preprint arXiv:2306.04719* (2023).

[60] Robert Geirhos et al. *ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness*. 2018. DOI: 10.48550/ARXIV.1811.12231. URL: https://arxiv.org/abs/1811.12231.

[61] Ross Girshick. *Fast R-CNN*. 2015. DOI: 10.48550/ARXIV.1504.08083. URL: https://arxiv.org/abs/1504.08083.

[62] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[63] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. *Explaining and Harnessing Adversarial Examples*. 2014. DOI: 10.48550/ARXIV.1412.6572. URL: https://arxiv.org/abs/1412.6572.

[64] Bryce Goodman and Seth Flaxman. "European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation"". In: *AI Magazine* 38.3 (Oct. 2017), pp. 50–57. DOI: 10.1609/aimag.v38i3.2741. URL: https://doi.org/10.1609%5C%2Faimag.v38i3.2741.

[65] Google. *What is Google Lens?* URL: https://lens.google/howlensworks/ (visited on 03/04/2023).

[66] Ishaan Gulrajani and David Lopez-Paz. *In Search of Lost Domain Generalization*. 2020. DOI: 10.48550/ARXIV.2007.01434. URL: https://arxiv.org/abs/2007.01434.

[67] Ishaan Gulrajani and David Lopez-Paz. "In Search of Lost Domain Generalization". In: *CoRR* abs/2007.01434 (2020). arXiv: 2007.01434. URL: https://arxiv.org/abs/2007.01434.

[68] Chuan Guo et al. "Countering Adversarial Images using Input Transformations". In: *CoRR* abs/1711.00117 (2017). arXiv: 1711.00117. URL: http://arxiv.org/abs/1711.00117.

[69] Chuan Guo et al. *On Calibration of Modern Neural Networks*. 2017. DOI: 10.48550/ARXIV.1706.04599. URL: https://arxiv.org/abs/1706.04599.

[70] Han Guo et al. *FastIF: Scalable Influence Functions for Efficient Model Interpretation and Debugging*. 2020. DOI: 10.48550/ARXIV.2012.15781. URL: https://arxiv.org/abs/2012.15781.

[71] Abner Guzman-Rivera, Dhruv Batra, and Pushmeet Kohli. "Multiple Choice Learning: Learning to Produce Multiple Structured Outputs". In: *NIPS'12 Proceedings of the 25th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Dec. 2012, pp. 1799–1807. URL: https://www.microsoft.com/en-us/research/publication/multiple-choice-learning-learning-produce-multiple-structured-outputs/.

[72] Dongyoon Han et al. *Neglected Free Lunch – Learning Image Classifiers Using Annotation Byproducts*. 2023. arXiv: 2303.17595 [cs.CV].

[73] Kaiming He et al. *Deep Residual Learning for Image Recognition*. 2015. DOI: 10.48550/ARXIV.1512.03385. URL: https://arxiv.org/abs/1512.03385.

[74] Kaiming He et al. *Mask R-CNN*. 2017. DOI: 10.48550/ARXIV.1703.06870. URL: https://arxiv.org/abs/1703.06870.

[75] Lisa Anne Hendricks et al. *Generating Counterfactual Explanations with Natural Language*. 2018. DOI: 10.48550/ARXIV.1806.09809. URL: https://arxiv.org/abs/1806.09809.

[76] hendrycks. *Benchmarking Neural Network Robustness to Common Corruptions and Perturbations*. URL: https://github.com/hendrycks/robustness (visited on 03/01/2023).

[77] Dan Hendrycks and Thomas Dietterich. "Benchmarking Neural Network Robustness to Common Corruptions and Perturbations". In: *Proceedings of the International Conference on Learning Representations* (2019).

[78] Dan Hendrycks et al. "Natural Adversarial Examples". In: *arXiv preprint arXiv:1907.07174* (2019).

[79] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. *Distilling the Knowledge in a Neural Network*. 2015. DOI: 10.48550/ARXIV.1503.02531. URL: https://arxiv.org/abs/1503.02531.

[80] Jonathan Ho, Ajay Jain, and Pieter Abbeel. *Denoising Diffusion Probabilistic Models*. 2020. DOI: 10.48550/ARXIV.2006.11239. URL: https://arxiv.org/abs/2006.11239.

[81] Judy Hoffman, Daniel A. Roberts, and Sho Yaida. *Robust Learning with Jacobian Regularization*. 2019. DOI: 10.48550/ARXIV.1908.02729. URL: https://arxiv.org/abs/1908.02729.

[82]   Gao Huang et al. *Deep Networks with Stochastic Depth*. 2016. DOI: 10.48550/ARXIV.1603.09382. URL: https://arxiv.org/abs/1603.09382.

[83]   Zeyi Huang et al. "Self-Challenging Improves Cross-Domain Generalization". In: *CoRR* abs/2007.02454 (2020). arXiv: 2007.02454. URL: https://arxiv.org/abs/2007.02454.

[84]   Eyke Hüllermeier and Willem Waegeman. "Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods". In: *Machine Learning* 110.3 (Mar. 2021), pp. 457–506. DOI: 10.1007/s10994-021-05946-3. URL: https://doi.org/10.1007/s10994-021-05946-3.

[85]   informIT. *What Software Architects Need to Know About DevOps*. URL: https://www.informit.com/articles/article.aspx?p=2424801 (visited on 02/14/2023).

[86]   Saachi Jain et al. *Missingness Bias in Model Debugging*. 2022. DOI: 10.48550/ARXIV.2204.08945. URL: https://arxiv.org/abs/2204.08945.

[87]   Ziwei Ji, Justin Li, and Matus Telgarsky. "Early-stopped neural networks are consistent". In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato et al. Vol. 34. Curran Associates, Inc., 2021, pp. 1805–1817. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/0e1ebad68af7f0ae4830b7ac92bc3c6f-Paper.pdf.

[88]   Heinrich Jiang. "Non-asymptotic uniform rates of consistency for k-nn regression". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 01. 2019, pp. 3999–4006.

[89]   Jeff Johnson, Matthijs Douze, and Hervé Jégou. "Billion-scale similarity search with GPUs". In: *IEEE Transactions on Big Data* 7.3 (2019), pp. 535–547.

[90]   Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 1st. USA: Prentice Hall PTR, 2000. ISBN: 0130950696.

[91]   Dae-Young Kang, Pham Duong, and Jung-Chul Park. "Application of Deep Learning in Dentistry and Implantology". In: *The Korean Academy of Oral and Maxillofacial Implantology* 24 (Sept. 2020), pp. 148–181. DOI: 10.32542/implantology.202015.

[92]   Jared Kaplan et al. "Scaling Laws for Neural Language Models". In: *CoRR* abs/2001.08361 (2020). arXiv: 2001.08361. URL: https://arxiv.org/abs/2001.08361.

[93]   Divyansh Kaushik, Eduard Hovy, and Zachary C. Lipton. *Learning the Difference that Makes a Difference with Counterfactually-Augmented Data*. 2019. DOI: 10.48550/ARXIV.1909.12434. URL: https://arxiv.org/abs/1909.12434.

[94]   Maxime Kayser et al. *e-ViL: A Dataset and Benchmark for Natural Language Explanations in Vision-Language Tasks*. 2021. DOI: 10.48550/ARXIV.2105.03761. URL: https://arxiv.org/abs/2105.03761.

[95]   Frank C Keil. "Explanation and understanding". In: *Annu. Rev. Psychol.* 57 (2006), pp. 227–254.

[96]   Alex Kendall and Yarin Gal. *What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?* 2017. DOI: 10.48550/ARXIV.1703.04977. URL: https://arxiv.org/abs/1703.04977.

[97]   William Ogilvy Kermack and Anderson G McKendrick. "A contribution to the mathematical theory of epidemics". In: *Proceedings of the royal society of london. Series A,*

Containing papers of a mathematical and physical character 115.772 (1927), pp. 700–721.

[98] Been Kim et al. "Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)". In: (2017). DOI: 10.48550/ARXIV.1711.11279. URL: https://arxiv.org/abs/1711.11279.

[99] Jae Myung Kim et al. "Keep calm and improve visual feature attribution". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 8350–8360.

[100] Diederik P Kingma and Max Welling. *Auto-Encoding Variational Bayes*. 2013. DOI: 10.48550/ARXIV.1312.6114. URL: https://arxiv.org/abs/1312.6114.

[101] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG].

[102] Michael Kirchhof, Enkelejda Kasneci, and Seong Joon Oh. *Probabilistic Contrastive Learning Recovers the Correct Aleatoric Uncertainty of Ambiguous Inputs*. 2023. DOI: 10.48550/ARXIV.2302.02865. URL: https://arxiv.org/abs/2302.02865.

[103] Alexander Kirillov et al. *Panoptic Segmentation*. 2019. arXiv: 1801.00868 [cs.CV].

[104] Pang Wei Koh and Percy Liang. *Understanding Black-box Predictions via Influence Functions*. 2017. DOI: 10.48550/ARXIV.1703.04730. URL: https://arxiv.org/abs/1703.04730.

[105] Pang Wei Koh et al. "WILDS: A Benchmark of in-the-Wild Distribution Shifts". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, July 2021, pp. 5637–5664. URL: https://proceedings.mlr.press/v139/koh21a.html.

[106] Takeshi Kojima et al. *Large Language Models are Zero-Shot Reasoners*. 2022. DOI: 10.48550/ARXIV.2205.11916. URL: https://arxiv.org/abs/2205.11916.

[107] Alexander Kolesnikov et al. *Big Transfer (BiT): General Visual Representation Learning*. 2020. arXiv: 1912.11370 [cs.CV].

[108] Ivan Krasin et al. "OpenImages: A public dataset for large-scale multi-label and multi-class image classification." In: *Dataset available from https://storage.googleapis.com/openimages/web/index.html* (2017).

[109] Satyapriya Krishna et al. *The Disagreement Problem in Explainable Machine Learning: A Practitioner's Perspective*. 2022. arXiv: 2202.01602 [cs.LG].

[110] Alex Krizhevsky, Geoffrey Hinton, et al. "Learning multiple layers of features from tiny images". In: (2009).

[111] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*. NIPS'12. Lake Tahoe, Nevada: Curran Associates Inc., 2012, pp. 1097–1105.

[112] Salem Lahlou et al. *DEUP: Direct Epistemic Uncertainty Prediction*. 2023. arXiv: 2102.08501 [cs.LG].

[113] Brenden M Lake. "Towards more human-like concept learning in machines: Compositionality, causality, and learning-to-learn". PhD thesis. Massachusetts Institute of Technology, 2014.

[114] Himabindu Lakkaraju et al. "Rethinking Explainability as a Dialogue: A Practitioner's Perspective". In: *arXiv preprint arXiv:2202.01875* (2022).

[115] Balaji Lakshminarayanan. *Introduction to Uncertainty in Deep Learning*. URL: https://www.gatsby.ucl.ac.uk/~balaji/balaji-uncertainty-talk-cifar-dlrl.pdf (visited on 03/02/2023).

[116] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. *Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles*. 2016. DOI: 10.48550/ARXIV.1612.01474. URL: https://arxiv.org/abs/1612.01474.

[117] Yann LeCun, Corinna Cortes, and CJ Burges. "MNIST handwritten digit database". In: *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist* 2 (2010).

[118] Jungbeom Lee et al. *Weakly Supervised Semantic Segmentation using Out-of-Distribution Data*. 2022. DOI: 10.48550/ARXIV.2203.03860. URL: https://arxiv.org/abs/2203.03860.

[119] Kimin Lee et al. *A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks*. 2018. DOI: 10.48550/ARXIV.1807.03888. URL: https://arxiv.org/abs/1807.03888.

[120] Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. *Question and Answer Test-Train Overlap in Open-Domain Question Answering Datasets*. 2020. DOI: 10.48550/ARXIV.2008.02637. URL: https://arxiv.org/abs/2008.02637.

[121] Da Li et al. *Deeper, Broader and Artier Domain Generalization*. 2017. DOI: 10.48550/ARXIV.1710.03077. URL: https://arxiv.org/abs/1710.03077.

[122] Da Li et al. "Deeper, Broader and Artier Domain Generalization". In: *CoRR* abs/1710.03077 (2017). arXiv: 1710.03077. URL: http://arxiv.org/abs/1710.03077.

[123] Tsung-Yi Lin et al. "Microsoft COCO: Common Objects in Context". In: *CoRR* abs/1405.0312 (2014). arXiv: 1405.0312. URL: http://arxiv.org/abs/1405.0312.

[124] Zachary C Lipton. "The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery." In: *Queue* 16.3 (2018), pp. 31–57.

[125] Francesco Locatello et al. "Challenging common assumptions in the unsupervised learning of disentangled representations". In: *international conference on machine learning*. PMLR. 2019, pp. 4114–4124.

[126] LossLandscape. *Mode Connectivity*. URL: https://losslandscape.com/ (visited on 03/02/2023).

[127] Scott Lundberg and Su-In Lee. *A Unified Approach to Interpreting Model Predictions*. 2017. DOI: 10.48550/ARXIV.1705.07874. URL: https://arxiv.org/abs/1705.07874.

[128] Paula M Luz, Claudio J Struchiner, and Alison P Galvani. "Modeling transmission dynamics and control of vector-borne neglected tropical diseases". In: *PLoS neglected tropical diseases* 4.10 (2010), e761.

[129] Xingjun Ma et al. *Characterizing Adversarial Subspaces Using Local Intrinsic Dimensionality*. 2018. DOI: 10.48550/ARXIV.1801.02613. URL: https://arxiv.org/abs/1801.02613.

[130] Papers With Code: The latest in Machine Learning. *Image Classification on ImageNet*. URL: https://paperswithcode.com/sota/image-classification-on-imagenet (visited on 03/05/2023).

[131] Papers With Code: The latest in Machine Learning. *Language Modelling on Penn Tree-bank (Word Level)*. URL: `https : / / paperswithcode . com / sota / language - modelling-on-penn-treebank-word` (visited on 03/04/2023).

[132] Wesley Maddox et al. *A Simple Baseline for Bayesian Uncertainty in Deep Learning*. 2019. DOI: `10.48550/ARXIV.1902.02476`. URL: `https://arxiv.org/abs/1902.02476`.

[133] Aleksander Madry et al. *Towards Deep Learning Models Resistant to Adversarial Attacks*. 2017. DOI: `10.48550/ARXIV.1706.06083`. URL: `https://arxiv.org/abs/1706.06083`.

[134] Eric Maggard. *Review of "Robust Physical-World Attacks on Machine Learning Models"*. URL: `https://medium.com/@ericmaggard/review-of-robust-physical-world-attacks-on-machine-learning-models-46d19dd5873a` (visited on 02/27/2023).

[135] Osama Makansi et al. "Overcoming Limitations of Mixture Density Networks: A Sampling and Fitting Framework for Multimodal Future Prediction". In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2019. DOI: `10.1109/cvpr.2019.00731`. URL: `https://doi.org/10.1109/cvpr.2019.00731`.

[136] Bertram F Malle. *How the mind explains behavior: Folk explanations, meaning, and social interaction*. MIT press, 2006.

[137] Markus Mathias et al. "Face detection without bells and whistles". In: *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part IV 13*. Springer. 2014, pp. 720–735.

[138] Tom McCoy, Ellie Pavlick, and Tal Linzen. "Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference". In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 3428–3448. DOI: `10.18653/v1/P19-1334`. URL: `https://aclanthology.org/P19-1334`.

[139] Tim Miller. "Explanation in Artificial Intelligence: Insights from the Social Sciences". In: *CoRR* abs/1706.07269 (2017). arXiv: `1706.07269`. URL: `http://arxiv.org/abs/1706.07269`.

[140] Matthias Minderer et al. *Revisiting the Calibration of Modern Neural Networks*. 2021. DOI: `10.48550/ARXIV.2106.07998`. URL: `https://arxiv.org/abs/2106.07998`.

[141] Tom M Mitchell. "The need for biases in learning generalizations". In: (1980).

[142] Christoph Molnar. *Interpretable machine learning*. Lulu. com, 2020.

[143] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. *DeepFool: a simple and accurate method to fool deep neural networks*. 2015. DOI: `10.48550/ARXIV.1511.04599`. URL: `https://arxiv.org/abs/1511.04599`.

[144] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. *A Metric Learning Reality Check*. 2020. DOI: `10.48550/ARXIV.2003.08505`. URL: `https://arxiv.org/abs/2003.08505`.

[145] Jun Hyun Nam et al. "Learning from Failure: Training Debiased Classifier from Biased Classifier". In: *CoRR* abs/2007.02561 (2020). arXiv: `2007.02561`. URL: `https://arxiv.org/abs/2007.02561`.

[146] Michael Neely et al. *A Song of (Dis)agreement: Evaluating the Evaluation of Explainable Artificial Intelligence in Natural Language Processing*. 2022. arXiv: `2205.04559 [cs.CL]`.

[147] Yuval Netzer et al. "Reading Digits in Natural Images with Unsupervised Feature Learning". In: *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*. 2011.

URL: http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf.

[148] Anh Nguyen, Jason Yosinski, and Jeff Clune. *Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images*. 2014. DOI: 10.48550/ARXIV.1412.1897. URL: https://arxiv.org/abs/1412.1897.

[149] NVIDIA. *NVIDIA A100 Tensor Core GPU – Unprecedented acceleration at every scale*. URL: https://www.nvidia.com/en-us/data-center/a100/ (visited on 02/14/2023).

[150] Seong Joon Oh et al. *Modeling Uncertainty with Hedged Instance Embedding*. 2018. DOI: 10.48550/ARXIV.1810.00319. URL: https://arxiv.org/abs/1810.00319.

[151] Avital Oliver et al. "Realistic evaluation of deep semi-supervised learning algorithms". In: *Advances in neural information processing systems* 31 (2018).

[152] OpenAI. *GPT-4 Technical Report*. 2023. arXiv: 2303.08774 [cs.CL].

[153] OpenAI. *Introducing ChatGPT*. URL: https://openai.com/blog/chatgpt#OpenAI (visited on 03/03/2023).

[154] Nicolas Papernot et al. *Practical Black-Box Attacks against Machine Learning*. 2016. DOI: 10.48550/ARXIV.1602.02697. URL: https://arxiv.org/abs/1602.02697.

[155] Joshua C. Peterson et al. *Human uncertainty makes classification more robust*. 2019. arXiv: 1908.07086 [cs.CV].

[156] Mohammad Pezeshki et al. *Gradient Starvation: A Learning Proclivity in Neural Networks*. 2020. DOI: 10.48550/ARXIV.2011.09468. URL: https://arxiv.org/abs/2011.09468.

[157] Jordi Pont-Tuset. *Open Images V6 — Now Featuring Localized Narratives*. URL: https://ai.googleblog.com/2020/02/open-images-v6-now-featuring-localized.html (visited on 03/03/2023).

[158] Garima Pruthi et al. *Estimating Training Data Influence by Tracing Gradient Descent*. 2020. DOI: 10.48550/ARXIV.2002.08484. URL: https://arxiv.org/abs/2002.08484.

[159] Alec Radford et al. *Learning Transferable Visual Models From Natural Language Supervision*. 2021. DOI: 10.48550/ARXIV.2103.00020. URL: https://arxiv.org/abs/2103.00020.

[160] Edward Raff et al. "Barrage of Random Transforms for Adversarially Robust Defense". In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019, pp. 6521–6530. DOI: 10.1109/CVPR.2019.00669.

[161] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. *Certified Defenses against Adversarial Examples*. 2018. DOI: 10.48550/ARXIV.1801.09344. URL: https://arxiv.org/abs/1801.09344.

[162] Aditya Ramesh et al. *Zero-Shot Text-to-Image Generation*. 2021. DOI: 10.48550/ARXIV.2102.12092. URL: https://arxiv.org/abs/2102.12092.

[163] Benjamin Recht et al. *Do ImageNet Classifiers Generalize to ImageNet?* 2019. DOI: 10.48550/ARXIV.1902.10811. URL: https://arxiv.org/abs/1902.10811.

[164] Jie Ren and Balaji Lakshminarayanan. *Improving Out-of-Distribution Detection in Machine Learning Models*. URL: https://ai.googleblog.com/2019/12/improving-out-of-distribution-detection.html (visited on 03/02/2023).

[165] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "" Why should i trust you?" Explaining the predictions of any classifier". In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2016, pp. 1135–1144.

[166] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. *"Why Should I Trust You?": Explaining the Predictions of Any Classifier*. 2016. DOI: 10.48550/ARXIV.1602.04938. URL: https://arxiv.org/abs/1602.04938.

[167] Robin Rombach et al. *High-Resolution Image Synthesis with Latent Diffusion Models*. 2021. DOI: 10.48550/ARXIV.2112.10752. URL: https://arxiv.org/abs/2112.10752.

[168] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. *U-Net: Convolutional Networks for Biomedical Image Segmentation*. 2015. DOI: 10.48550/ARXIV.1505.04597. URL: https://arxiv.org/abs/1505.04597.

[169] Andrew Slavin Ross et al. *Ensembles of Locally Independent Prediction Models*. 2019. DOI: 10.48550/ARXIV.1911.01291. URL: https://arxiv.org/abs/1911.01291.

[170] Shiori Sagawa and Pang Wei Koh. *WILDS: A Benchmark of in-the-Wild Distribution Shifts*. URL: https://ai.stanford.edu/blog/wilds/ (visited on 03/01/2023).

[171] Shiori Sagawa et al. *Distributionally Robust Neural Networks for Group Shifts: On the Importance of Regularization for Worst-Case Generalization*. 2019. DOI: 10.48550/ARXIV.1911.08731. URL: https://arxiv.org/abs/1911.08731.

[172] Andrea Schioppa et al. *Scaling Up Influence Functions*. 2021. DOI: 10.48550/ARXIV.2112.03052. URL: https://arxiv.org/abs/2112.03052.

[173] Luca Scimeca et al. *Which Shortcut Cues Will DNNs Choose? A Study from the Parameter-Space Perspective*. 2021. DOI: 10.48550/ARXIV.2110.03095. URL: https://arxiv.org/abs/2110.03095.

[174] Ramprasaath R Selvaraju et al. "Grad-cam: Visual explanations from deep networks via gradient-based localization". In: *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 618–626.

[175] Mohammad Hossein Shaker and Eyke Hüllermeier. *Ensemble-based Uncertainty Quantification: Bayesian versus Credal Inference*. 2021. arXiv: 2107.10384 [cs.LG].

[176] Vaishaal Shankar et al. "Evaluating Machine Accuracy on ImageNet". In: *Proceedings of the 37th International Conference on Machine Learning*. Ed. by Hal Daumé III and Aarti Singh. Vol. 119. Proceedings of Machine Learning Research. PMLR, July 2020, pp. 8634–8644. URL: https://proceedings.mlr.press/v119/shankar20c.html.

[177] Lloyd S Shapley et al. "A value for n-person games". In: (1953).

[178] Xinyue Shen et al. *"Do Anything Now": Characterizing and Evaluating In-The-Wild Jailbreak Prompts on Large Language Models*. 2023. arXiv: 2308.03825 [cs.CR].

[179] Rakshith Shetty, Bernt Schiele, and Mario Fritz. *Not Using the Car to See the Sidewalk: Quantifying and Controlling the Effects of Context in Classification and Segmentation*. 2018. DOI: 10.48550/ARXIV.1812.06707. URL: https://arxiv.org/abs/1812.06707.

[180] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. "Learning Important Features Through Propagating Activation Differences". In: (2017). DOI: 10.48550/ARXIV.1704.02685. URL: https://arxiv.org/abs/1704.02685.

[181] Nathan Silberman et al. *NYU Depth Dataset V2*. URL: https://cs.nyu.edu/~silberman/datasets/nyu_depth_v2.html (visited on 03/02/2023).

[182] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. *Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps*. 2013. DOI: 10.48550/ARXIV.1312.6034. URL: https://arxiv.org/abs/1312.6034.

[183] Sandeep Kumar Singh et al. *Leveraging Graph and Deep Learning Uncertainties to Detect Anomalous Trajectories*. 2021. DOI: 10.48550/ARXIV.2107.01557. URL: https://arxiv.org/abs/2107.01557.

[184] Daniel Smilkov et al. *SmoothGrad: removing noise by adding noise*. 2017. DOI: 10.48550/ARXIV.1706.03825. URL: https://arxiv.org/abs/1706.03825.

[185] Samuel L. Smith et al. *On the Origin of Implicit Regularization in Stochastic Gradient Descent*. 2021. DOI: 10.48550/ARXIV.2101.12176. URL: https://arxiv.org/abs/2101.12176.

[186] Yang Song et al. *PixelDefend: Leveraging Generative Models to Understand and Defend against Adversarial Examples*. 2017. DOI: 10.48550/ARXIV.1710.10766. URL: https://arxiv.org/abs/1710.10766.

[187] Jost Tobias Springenberg et al. *Striving for Simplicity: The All Convolutional Net*. 2014. DOI: 10.48550/ARXIV.1412.6806. URL: https://arxiv.org/abs/1412.6806.

[188] Nitish Srivastava et al. "Dropout: A Simple Way to Prevent Neural Networks from Overfitting". In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. URL: http://jmlr.org/papers/v15/srivastava14a.html.

[189] Doug Steen. *Precision-Recall Curves*. URL: https://medium.com/@douglaspsteen/precision-recall-curves-d32e5b290248 (visited on 03/02/2023).

[190] Yu Sun et al. *Test-Time Training with Self-Supervision for Generalization under Distribution Shifts*. 2020. arXiv: 1909.13231 [cs.LG].

[191] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. "Axiomatic attribution for deep networks". In: *International conference on machine learning*. PMLR. 2017, pp. 3319–3328.

[192] Christian Szegedy et al. *Intriguing properties of neural networks*. 2013. DOI: 10.48550/ARXIV.1312.6199. URL: https://arxiv.org/abs/1312.6199.

[193] Damien Teney, Maxime Peyrard, and Ehsan Abbasnejad. *Predicting is not Understanding: Recognizing and Addressing Underspecification in Machine Learning*. 2022. DOI: 10.48550/ARXIV.2207.02598. URL: https://arxiv.org/abs/2207.02598.

[194] Damien Teney et al. "Evading the Simplicity Bias: Training a Diverse Set of Models Discovers Solutions with Superior OOD Generalization". In: *arXiv preprint arXiv:2105.05612* (2021).

[195] S. Thys, W. Ranst, and T. Goedeme. "Fooling Automated Surveillance Cameras: Adversarial Patches to Attack Person Detection". In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2019, pp. 49–55. DOI: 10.1109/CVPRW.2019.00012. URL: https://doi.ieeecomputersociety.org/10.1109/CVPRW.2019.00012.

[196] Ilya Tolstikhin et al. *MLP-Mixer: An all-MLP Architecture for Vision*. 2021. arXiv: 2105.01601 [cs.CV].

[197] Du Tran et al. "A Closer Look at Spatiotemporal Convolutions for Action Recognition". In: *CoRR* abs/1711.11248 (2017). arXiv: 1711.11248. URL: http://arxiv.org/abs/1711.11248.

[198] Harshit Tyagi. *What is MLOps – Everything You Must Know to Get Started*. URL: https://towardsdatascience.com/what-is-mlops-everything-you-must-know-to-get-started-523f2d0b8bd8 (visited on 02/14/2023).

[199] Ivan Ukhov. *Heteroscedastic Gaussian process regression*. URL: https://blog.ivanukhov.com/2020/06/22/gaussian-process.html (visited on 03/02/2023).

[200] Matias Valdenegro-Toro and Daniel Saromo. *A Deeper Look into Aleatoric and Epistemic Uncertainty Disentanglement*. 2022. arXiv: `2204.09308 [cs.LG]`.

[201] Ashish Vaswani et al. *Attention Is All You Need*. 2017. DOI: `10.48550/ARXIV.1706.03762`. URL: `https://arxiv.org/abs/1706.03762`.

[202] Li Wan et al. "Regularization of Neural Networks using DropConnect". In: *Proceedings of the 30th International Conference on Machine Learning*. Ed. by Sanjoy Dasgupta and David McAllester. Vol. 28. Proceedings of Machine Learning Research 3. Atlanta, Georgia, USA: PMLR, June 2013, pp. 1058–1066. URL: `https://proceedings.mlr.press/v28/wan13.html`.

[203] Jason Wei et al. *Chain-of-Thought Prompting Elicits Reasoning in Large Language Models*. 2023. arXiv: `2201.11903 [cs.CL]`.

[204] Philippe Weinzaepfel and Grégory Rogez. "2019". In: *arXiv* ().

[205] Max Welling and Yee W Teh. "Bayesian learning via stochastic gradient Langevin dynamics". In: *Proceedings of the 28th international conference on machine learning (ICML-11)*. 2011, pp. 681–688.

[206] Wikipedia. *ImageNet*. URL: `https://en.wikipedia.org/wiki/ImageNet` (visited on 08/17/2023).

[207] Wikipedia. *Scarlet tanager*. URL: `https://en.wikipedia.org/wiki/Scarlet_tanager` (visited on 03/18/2023).

[208] Wikipedia. *Thermodynamic system*. URL: `https://en.wikipedia.org/wiki/Thermodynamic_system` (visited on 02/14/2023).

[209] Adina Williams, Nikita Nangia, and Samuel Bowman. "A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, 2018, pp. 1112–1122. URL: `http://aclweb.org/anthology/N18-1101`.

[210] Andrew Gordon Wilson. *Bayesian Deep Learning and a Probabilistic Perspective of Model Construction*. URL: `https://cims.nyu.edu/~andrewgw/bayesdlicml2020.pdf` (visited on 03/02/2023).

[211] Andrew Gordon Wilson and Pavel Izmailov. *Bayesian Deep Learning and a Probabilistic Perspective of Generalization*. 2020. DOI: `10.48550/ARXIV.2002.08791`. URL: `https://arxiv.org/abs/2002.08791`.

[212] David H Wolpert and William G Macready. "No free lunch theorems for optimization". In: *IEEE transactions on evolutionary computation* 1.1 (1997), pp. 67–82.

[213] Chao-Yuan Wu et al. *Sampling Matters in Deep Embedding Learning*. 2017. DOI: `10.48550/ARXIV.1706.07567`. URL: `https://arxiv.org/abs/1706.07567`.

[214] Yongqin Xian, Bernt Schiele, and Zeynep Akata. "Zero-shot learning-the good, the bad and the ugly". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 4582–4591.

[215] Chaowei Xiao et al. "Spatially Transformed Adversarial Examples". In: *International Conference on Learning Representations*. 2018. URL: `https://openreview.net/forum?id=HyydRMZC-`.

[216] Han Xiao, Kashif Rasul, and Roland Vollgraf. *Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms*. 2017. DOI: 10.48550/ARXIV.1708.07747. URL: https://arxiv.org/abs/1708.07747.

[217] Cihang Xie et al. "Mitigating Adversarial Effects Through Randomization". In: *International Conference on Learning Representations*. 2018. URL: https://openreview.net/forum?id=Sk9yuql0Z.

[218] Jason Yosinski et al. *Understanding Neural Networks Through Deep Visualization*. 2015. DOI: 10.48550/ARXIV.1506.06579. URL: https://arxiv.org/abs/1506.06579.

[219] H Peyton Young. "Monotonic solutions of cooperative games". In: *International Journal of Game Theory* 14.2 (1985), pp. 65–72.

[220] Jianming Zhang et al. *Top-down Neural Attention by Excitation Backprop*. 2016. DOI: 10.48550/ARXIV.1608.00507. URL: https://arxiv.org/abs/1608.00507.

[221] Richard Zhang et al. "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric". In: *CoRR* abs/1801.03924 (2018). arXiv: 1801.03924. URL: http://arxiv.org/abs/1801.03924.

[222] Lei Zhao et al. "Uctgan: Diverse image inpainting based on unsupervised cross-space translation". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020, pp. 5741–5750.

[223] Bolei Zhou et al. *Learning Deep Features for Discriminative Localization*. 2015. DOI: 10.48550/ARXIV.1512.04150. URL: https://arxiv.org/abs/1512.04150.

[224] Kaiyang Zhou et al. *Deep Domain-Adversarial Image Generation for Domain Generalisation*. 2020. DOI: 10.48550/ARXIV.2003.06054. URL: https://arxiv.org/abs/2003.06054.

[225] Roland S Zimmermann et al. "How Well do Feature Visualizations Support Causal Understanding of CNN Activations?" In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 11730–11744.

[226] Luisa M Zintgraf et al. *Visualizing Deep Neural Network Decisions: Prediction Difference Analysis*. 2017. DOI: 10.48550/ARXIV.1702.04595. URL: https://arxiv.org/abs/1702.04595.