
Characterizing Structural Regularities of Labeled Data in Overparameterized Models

Ziheng Jiang^{*123} Chiyuan Zhang^{*4} Kunal Talwar⁴⁵ Michael C. Mozer⁴⁶

Abstract

Humans are accustomed to environments that contain both regularities and exceptions. For example, at most gas stations, one pays prior to pumping, but the occasional rural station does not accept payment in advance. Likewise, deep neural networks can generalize across instances that share common patterns or structures, yet have the capacity to memorize rare or irregular forms. We analyze how individual instances are treated by a model via a *consistency score*. The score characterizes the expected accuracy for a held-out instance given training sets of varying size sampled from the data distribution. We obtain empirical estimates of this score for individual instances in multiple data sets, and we show that the score identifies out-of-distribution and mislabeled examples at one end of the continuum and strongly regular examples at the other end. We identify computationally inexpensive proxies to the consistency score using statistics collected during training. We show examples of potential applications to the analysis of deep-learning systems.

1. Introduction

Human learning requires both inferring regular patterns that generalize across many distinct examples and memorizing irregular examples. The boundary between regular and irregular examples can be fuzzy. For example, in learning the past tense form of English verbs, there are some verbs whose past tenses must simply be memorized ($GO \rightarrow WENT$, $EAT \rightarrow ATE$, $HIT \rightarrow HIT$) and there are many *regular* verbs that obey the rule of appending “ed”

($KISS \rightarrow KISSED$, $KICK \rightarrow KICKED$, $BREW \rightarrow BREWED$, etc.). Generalization to a novel word typically follows the “ed” rule, for example, $BINK \rightarrow BINKED$. Intermediate between the exception verbs and regular verbs are subregularities—a set of exception verbs that have consistent structure (e.g., the mapping of $SING \rightarrow SANG$, $RING \rightarrow RANG$). Note that rule-governed and exception cases can have very similar forms, which increases the difficulty of learning each. Consider one-syllable verbs containing ‘ee’, which include the regular cases $NEED \rightarrow NEEDED$ as well as exception cases like $SEEK \rightarrow SOUGHT$. Generalization from the rule-governed cases can hamper the learning of the exception cases and vice-versa. For instance, children in an environment where English is spoken over-regularize by mapping $GO \rightarrow GOED$ early in the course of language learning. Neural nets show the same interesting pattern for verbs over the course of training (Rumelhart & McClelland, 1986).

Intuitively, memorizing irregular examples is tantamount to building a look-up table with the individual facts accessible for retrieval. Generalization requires the inference of statistical regularities in the training environment, and the application of procedures or rules for exploiting the regularities. In deep learning, memorization is often considered a failure of a network because memorization implies no generalization. However, mastering a domain involves knowing when to generalize and when not to generalize, because the data manifolds are rarely unimodal.

Consider the two-class problem of chair vs non-chair with training examples illustrated in Figure 1a. The iron throne (lower left) forms a sparsely populated mode (*sparse mode* for short) as there may not exist many similar cases in the data environment. Generic chairs (lower right) lie in a region with a consistent labeling (a densely populated mode, or *dense mode*) and thus seems to follow a strong regularity. But there are many other cases in the continuum of the two extreme. For example, the rocking chair (upper right) has a few supporting neighbors but it lies in a distinct neighborhood from the majority of same-label instances (the generic chairs).

In this article, we formalize this continuum of the structural regularities of data sets in the context of training overparameterized deep networks. Let $D \stackrel{\mathcal{N}}{\sim} \mathcal{P}$ be an i.i.d. sample of

^{*}Equal contribution ¹Paul G. Allen School of Computer Science, University of Washington, Seattle, WA, USA. ²OctoML.ai, Seattle, WA, USA. ³Work done while interning at Google. ⁴Google Research, Brain Team, Mountain View, CA, USA. ⁵Presently at Apple Inc., Cupertino, CA, USA. ⁶Department of Computer Science, University of Colorado Boulder, Boulder, CO, USA.. Correspondence to: Chiyuan Zhang <chiyuan@google.com>.

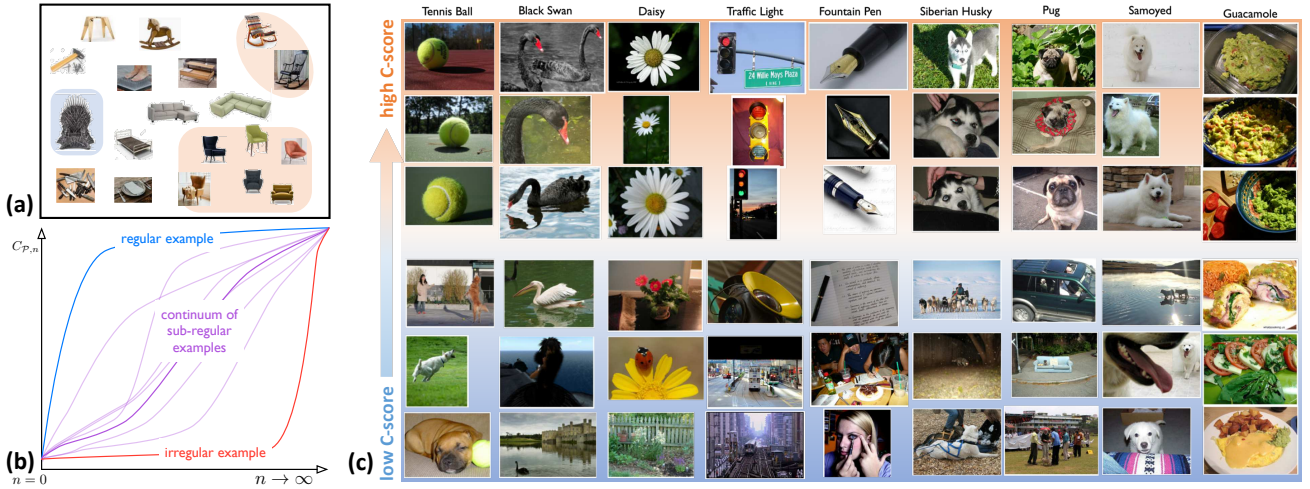


Figure 1. Regularities and exceptions in a binary chairs vs non-chairs problem. (b) illustration of consistency profiles. (c) Regularities (high C-scores) and exceptions (low C-scores) in ImageNet.

size n from the underlying data distribution \mathcal{P} , and $f(\cdot; D)$ be a model trained on D . For an instance x with label y , we trace out the following *consistency profile* by increasing n :

$$C_{\mathcal{P},n}(x, y) = \mathbb{E}_{D \sim \mathcal{P}}[\mathbb{P}(f(x; D) \setminus \{(x, y)\}) = y], \quad (1)$$

Note by taking expectation over (x, y) , this measures the *generalization performance* with respect to the underlying distribution \mathcal{P} . In contrast to the average behavior, we focus on the *per-instance generalization* here, as it helps to reveal the internal regularity structures of the data distribution. This article focuses on multi-class classification problems, but the definition can be easily extended to other problems by replacing the 0-1 classification loss with another suitable loss function.

$C_{\mathcal{P},n}(x, y)$ also encodes our high-level intuition about the structural regularities of the training data during (human or machine) learning. In particular, we can characterize the multimodal structure of an underlying data distribution by grouping examples in terms of a model’s generalization profile for those examples. An (x, y) with high per-instance generalization lies in a region on the data manifold that is well supported by other regular instances.

For $n = 0$, the model makes predictions entirely based on its prior belief. As n increases, the model collects more information about \mathcal{P} and makes better predictions. For an (x, y) instance belonging to a dense mode (e.g., the generic chairs in Figure 1a), the model prediction is accurate even for small n because even small samples have many class-consistent neighbors. The blue curve in the cartoon sketch of Figure 1b illustrates this profile. For instances belonging to sparse modes (e.g., the iron throne in Figure 1a), the prediction will be inaccurate for even large n , as the red curve illustrates. Most instances fill the continuum between these two extreme cases, as illustrated by the purple curves in Figure 1b. To obtain a *total ordering* for all examples, we pool the consistency profile into a scalar *consistency score*,

or *C-score* by taking expectation over n . Figure 1c shows examples from the ImageNet data set ranked by estimated C-scores, using a methodology we shortly describe. The images show that on many ImageNet classes, there exist dense modes of center-cropped, close-up shot of the representative examples; and at the other end of the C-score ranking, there exist sparse modes of highly ambiguous examples (in many cases, the object is barely seen or can only be inferred from the context in the picture).

With strong ties to both theoretical notions of generalization and human intuition, the consistency profile is an important tool for understanding the regularity and subregularity structures of training data sets and the learning dynamics of models trained on those data. The C-score based ranking also has many potential uses, such as detecting out-of-distribution and mislabeled instances; balancing learning between dense and sparse modes to ensure fairness when learning with data from underrepresented groups; or even as a diagnostic used to determine training priority in a curriculum learning setting (Bengio et al., 2009; Saxena et al., 2019). In this article, we focus on formulating and analyzing consistency profiles, and apply the C-score to analyzing the structure of real world image data sets and the learning dynamics of different optimizers. We also study efficient proxies and further applications to outlier detection.

Our key contributions are as follows:

- We formulate and analyze a consistency score that takes inspiration from generalization theory and show that it matches our intuitions about statistical regularities in natural-image data sets.
- We estimate the C-scores with a series of approximations and apply the measure to analyze the structural regularities of the MNIST, CIFAR-10, CIFAR-100, and ImageNet training sets.
- We evaluate computationally efficient proxies for the C-score. We demonstrate that proxies based on distances

between instances of the same class in latent space, while intuitively sensible, are in practice quite sensitive to the underlying distance metric. In contrast, learning-speed based proxies correlate very well with the C-score. This observation is non-trivial because learning speed is measured on training examples and the C-score is defined for hold-out generalization.

- We demonstrate potential application of the C-score as a tool for quantitative analysis of data sets, learning dynamics, and diagnosing and improving deep learning.
- To facilitate future research, we have released the pre-computed C-scores at (URL anonymized). Model checkpoints, code, and extra visualizations are available too.

2. Related Work

Analyzing the structure of data sets has been a central topic for many fields like Statistics, Data Mining and Unsupervised Learning. In this paper, we focus on supervised learning and the interplay between the regularity structure of data and overparameterized neural network learners. This differentiates our work from classical analyses based on input or (unsupervised) latent representations. The distinction is especially prominent in deep learning where a supervised learner jointly learns the classifier and the representation that captures the semantic information in the labels.

In the context of deep supervised learning, Carlini et al. (2018) proposed measures for identifying *prototypical* examples which could serve as a proxy for the complete data set and still achieve good performance. These examples are not necessarily the center of a dense neighborhood, which is what our high C-score measures. Two prototype measures explored in Carlini et al. (2018), *model confidence* and the *learning speed*, are also measures we examine. Their *holdout retraining* and *ensemble agreement* metrics are conceptually similar to our C-score estimation algorithm. However, their retraining is a two-stage procedure involving pre-training and fine-tuning; their ensemble agreement mixes architectures with heterogeneous capacities and ignores labels. Feldman (2020) and Feldman & Zhang (2020) studied the positive effects of memorization on generalization by measuring the influence of a training example on a test example, and identifying pairs with strong influences. To quantify memorization, they defined a memorization score for each (x, y) in a training set as the drop in prediction accuracy on x when (x, y) is removed. A point evaluation of our consistency profile on a fixed data size n resembles the second term of their score. Our empirical C-score estimation is based on the estimator proposed in Feldman & Zhang (2020). A key difference is that we are interested in the profile with increasing n , i.e. the sample complexity required to correctly predict (x, y) .

We evaluate various cheap-to-compute *proxies* for the C-




Figure 2. Consistency profiles of training examples. Each curve in the figure corresponds to the average profile of a set of examples, partitioned according to the area under the profile curve of each example.

score and found that the learning speed has a strong correlation with the C-score. Learning speed has been previously studied in contexts quite different from our focus on generalization of individual examples. Mangalam & Prabhu (2019) show that examples learned first are those that could be learned by shallower nets. Hardt et al. (2016) present theoretical results showing that the generalization gap is small if SGD training completes in relatively few steps. Toneva et al. (2019) study forgetting (the complement of learning speed) and informally relate forgetting to examples being outliers or mislabeled. There is a large literature of criteria with no explicit ties to generalization as the C-score has, but provides a means of stratifying instances. For example, Wu et al. (2018) measure the difficulty of an example by the number of residual blocks in a ResNet needed for prediction.

3. The Consistency Profile and the C-score

The consistency profile (Equation 1) encodes the structural consistency of an example with the underlying data distribution \mathcal{P} via expected performance of models trained with increasingly large data sets sampled from \mathcal{P} . However, it is not possible to directly compute this profile because \mathcal{P} is generally unknown for typical learning problems. In practice, we usually have a fixed data set $\hat{\mathcal{D}}$ consisting of N i.i.d. samples from \mathcal{P} . So we can estimate the consistency profile with the following *empirical consistency profile*:

$$\hat{C}_{\hat{\mathcal{D}},n}(x, y) = \hat{\mathbb{E}}_{D \sim \hat{\mathcal{D}}}^r [\mathbb{P}(f(x; D \setminus \{(x, y)\}) = y)], \quad (2)$$

where $n = 0, 1, \dots, N - 1$, D is a subset of size n uniformly sampled from $\hat{\mathcal{D}}$ excluding (x, y) , and $\hat{\mathbb{E}}^r$ denotes empirical averaging with r i.i.d. samples of such subsets. To obtain a reasonably accurate estimate (say, $r = 1000$), calculating the empirical consistency profile is still computationally prohibitive. For example, with each of the 50,000 training example in the CIFAR-10 training set, we need to train more than 2 trillion models. To obtain an estimate within the capability of current computation resources, we make two observations. First, model performance is generally stable when the training set size varies within a small range. Therefore, we can sample across the range of n that

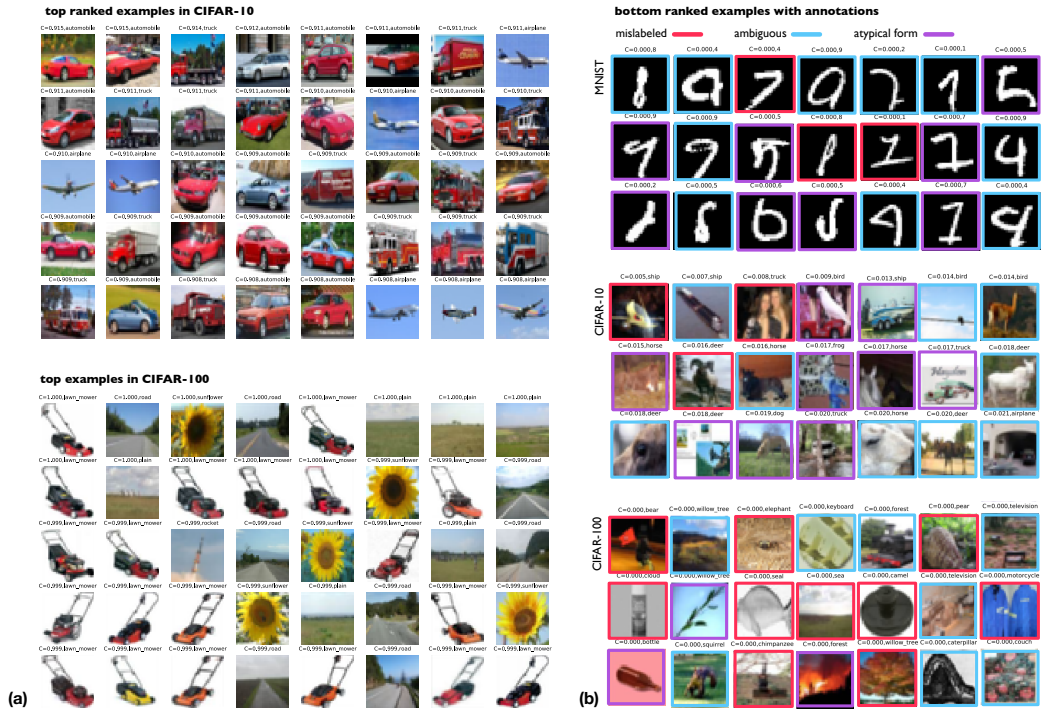


Figure 3. (a) Top ranked examples in CIFAR-10 and CIFAR-100. (b) Bottom ranked examples with annotations.

we’re concerned with and obtain the full profile via smooth interpolation. Second, let D be a random subset of training data, then the single model $f(\cdot; D)$ can be reused in the estimation of all of the held-out examples $(x, y) \in \mathcal{D} \setminus D$. As a result, with clever grouping and reuse, the number of models we need to train can be greatly reduced (See Algorithm 1 in the Appendix).

In particular, we sample n dynamically according to the *subset ratio* $s \in \{10\%, \dots, 90\%\}$ of the full available training set. We sample 2,000 subsets for the empirical expectation of each n and visualize the estimated consistency profiles for clusters of similar examples in Figure 2. One interesting observation is that while CIFAR-100 is generally more difficult than CIFAR-10, the top ranked examples (magenta lines) in CIFAR-100 are more likely to be classified correctly when the subset ratio is low. Figure 3a visualizes the top ranked examples from the two data sets. Note that in CIFAR-10, the dense modes from the truck and automobile classes are quite similar.

In contrast, Figure 2 indicates that the bottom-ranked examples (cyan lines) have persistently low probability of correct classification—sometimes below chance—even with a 90% subset ratio. We visualize some bottom-ranked examples and annotate them as (possibly) mislabeled, ambiguous (easily confused with another class or hard to identify the contents), and atypical form (e.g., burning “forest”, fallen “bottle”). As the subset ratio grows, regularities in the data distribution systematically pull the ambiguous instances in

the wrong direction. This behavior is analogous to the phenomenon we mentioned earlier that children over-regularize verbs (GO→GOED) as they gain more linguistic exposure.

To get a total ordering of the examples in a data set, we distill the consistency profiles into a scalar *consistency score*, or C-score, by taking the expectation over n :

$$\hat{C}_{\mathcal{D}}(x, y) = \mathbb{E}_n[\hat{C}_{\mathcal{D},n}(x, y)] \quad (3)$$

For the case where n is sampled according to the subset ratio s , the expectation is taken over a uniform distribution over sampled subset sizes.

4. The Structural Regularities of Common Image Data Sets

We apply the C-score estimate to analyze several common image data sets: MNIST (LeCun et al., 1998), CIFAR-10 / CIFAR-100 (Krizhevsky, 2009), and ImageNet (Russakovsky et al., 2015). See the supplementary materials for details on architectures and hyperparameters.

Figure 4a shows the distribution of $\hat{C}_{\mathcal{D},n}$ on CIFAR-10 for the values of n corresponding to each subset ratio $s \in \{10, \dots, 90\}$. For each s , 2000 models are trained and held-out examples are evaluated. The Figure suggests that depending on s , instances may be concentrated near floor or ceiling, making them difficult to distinguish (as we elaborate further shortly). By taking an expectation over s , the

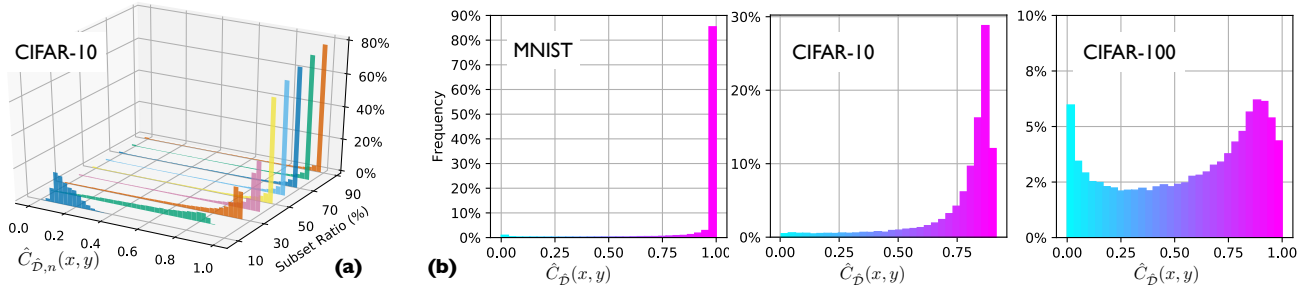


Figure 4. (a) Histogram of $\hat{C}_{\mathcal{D},n}$ for each subset ratio on CIFAR-10. (b) Histogram of the C-score $\hat{C}_{\mathcal{D}}$ averaged over all subset ratios on 3 different data sets.

C-score is less susceptible to floor and ceiling effects. Figure 4b shows the histogram of this integrated C-score on MNIST, CIFAR-10, and CIFAR-100. The histogram of CIFAR-10 in Figure 4b is distributed toward the high end, but is more uniformly spread than the histograms for specific subset ratios in Figure 4a.

Visualization of examples ranked by the estimated score can be found in Figure 3. Detailed per-class rankings can be found in the supplementary material.

Next we apply the C-score analysis to the ImageNet data set. Training a standard model on ImageNet costs one to two orders of magnitude more computing resources than training on CIFAR, preventing us from running the C-score estimation procedure described early. Instead, we investigated the feasibility of approximating the C-score with a *point estimate*, i.e., selection of the s that best represents the integral score. This is equivalent to taking expectation of s with respect to a point-mass distribution, as opposed to the uniform distribution over subset ratios. By ‘best represents,’ we mean that the ranking of instances by the score matches the ranking by the score for a particular s .

Figure 5a shows the rank correlation between the integral score and the score for a given s , as a function of s for our three smaller data sets, MNIST, CIFAR-10, and CIFAR-100. Examining the green CIFAR-10 curve, there is a peak at $s = 30$, indicating that $s = 30$ yields the best point-estimate approximation for the integral C-score. That the peak is at an intermediate s is consistent with the observation from Figure 2 that the C-score bunches together instances for low and high s . For MNIST (blue curve), a less challenging data set than CIFAR-10, the peak is lower, at $s = 10$; for CIFAR-100 (orange curve), a more challenging data set than CIFAR-10, the peak is higher, at $s = 40$ or $s = 50$. Thus, the peak appears to shift to larger s for more challenging data sets. This finding is not surprising: more challenging data sets require a greater diversity of training instances in order to observe generalization.

Based on these observations, we picked $s = 70$ for a point estimate on ImageNet. In particular, we train 2,000 ResNet-

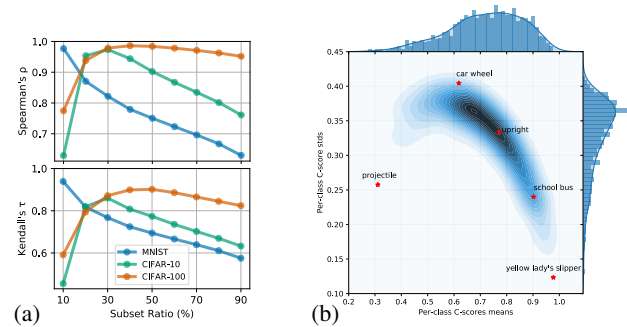


Figure 5. (a) Rank correlation between integral C-score and the C-score for a particular subset ratio, s . The peak of each curve indicates the training set size that best reveals generalization of the model. (b) Joint distribution of C-score per-class means and standard deviations on ImageNet. Samples from representative classes (\star 's) are shown in Figure 6.

50 models each with a random 70% subset of the ImageNet training set, and estimate the C-score based on those models.

The examples shown in Figure 1c are ranked according to this C-score estimate. Because ImageNet has 1,000 classes, we cannot offer a simple overview over the entire data set as in MNIST and CIFAR. Instead, we focus on analyzing the behaviors of individual classes. Specifically, we compute the mean and standard deviation (SD) of the C-scores of all the examples in a particular class. The mean C-score indicates the relative difficulty of classes, and the SD indicates the diversity of examples within each class. The two-dimensional histogram in Figure 5a depicts the joint distribution of mean and SD across all classes. We selected several classes with various combinations of mean and SD, indicated by the \star 's in Figure 5a. We then selected sample images from the top 99%, 35% and 1% percentile ranked by the C-score within each class, and show them in Figure 6.

Projectile and *yellow lady's slipper* represent two extreme cases of diverse and unified classes, respectively. Most other classes lie in the high density region of the 2D histogram

in Figure 5b, and share a common pattern of a densely populated mode of highly regular examples and a tail of rare, ambiguous examples. The tail becomes smaller from the class *car wheel* to *upright* and *school bus*.

5. C-score Proxies

We are able to reduce the cost of estimating C-scores from infeasible to feasible, but the procedure is still very expensive. Ideally, we would like to have more efficient *proxies* that do not require training multiple models. We use the term *proxy* to refer to any quantity that is well correlated with the C-score but does not have a direct mathematical relation to it, as contrasted with *approximations* that are designed to mathematically approximate the C-score (e.g., approximating the expectation with empirical averaging). The possible candidate set for C-score proxies is very large, as any measure that reflects information about difficulty or regularity of examples could be considered. Our Related Work section mentions a few such possibilities. In this paper, we primarily study two variants: *pairwise distance based proxies* and *learning speed based proxies*.

5.1. Pairwise Distance Based Proxies

Pairwise distance matches our intuition about consistency very well. In fact, our motivating example in Figure 1a is illustrated in this way. Intuitively, an example is consistent with the data distribution if it lies near other examples having the same label. However, if the example lies far from instances in the same class or lies near instances of different classes, one might not expect it to generalize. Based on this intuition, we define a relative local-density score:

$$\hat{C}^{\pm L}(x, y) = 1/N \sum_{i=1}^N 2(\mathbf{1}[y = y_i] - \frac{1}{2})K(x_i, x), \quad (4)$$

Table 1. Rank correlation between C-score and pairwise distance based proxies on inputs. Measured with Spearman’s ρ and Kendall’s τ rank correlations, respectively.

		\hat{C}	\hat{C}^L	$\hat{C}^{\pm L}$	\hat{C}^{LOF}
ρ	CIFAR-10	-0.064	-0.009	0.083	0.103
	CIFAR-100	-0.098	0.117	0.105	0.151
τ	CIFAR-10	-0.042	-0.006	0.055	0.070
	CIFAR-100	-0.066	0.078	0.070	0.101

where $K(x, x') = \exp(-\|x - x'\|^2/h^2)$ is an RBF kernel with the bandwidth h , and $\mathbf{1}[\cdot]$ is the indicator function. To evaluate the importance of explicit label information, we study two related scores: \hat{C}^L that uses only same-class examples when estimating the local density, and \hat{C} that uses all the neighbor examples by ignoring the labels.

$$\hat{C}^L(x, y) = 1/N \sum_{i=1}^N \mathbf{1}[y = y_i]K(x_i, x), \quad (5)$$

$$\hat{C}(x) = 1/N \sum_{i=1}^N K(x_i, x). \quad (6)$$

We also study a proxy based on the local outlier factor (LOF) algorithm (Breunig et al., 2000), which measures the local deviation of each point with respect to its neighbours. Since large LOF scores indicate outliers, we use the negative LOF score as a C-score proxy, denoted by $\hat{C}^{\text{LOF}}(x)$.

Table 1 shows the agreement between the proxy scores and the estimated C-score. Agreement is quantified by two rank correlation measures on three data sets. \hat{C}^{LOF} performs slightly better than the other proxies, but none of the proxies has high enough correlation to be useful, because it is very hard to obtain semantically meaningful distance estimations from the raw pixels.

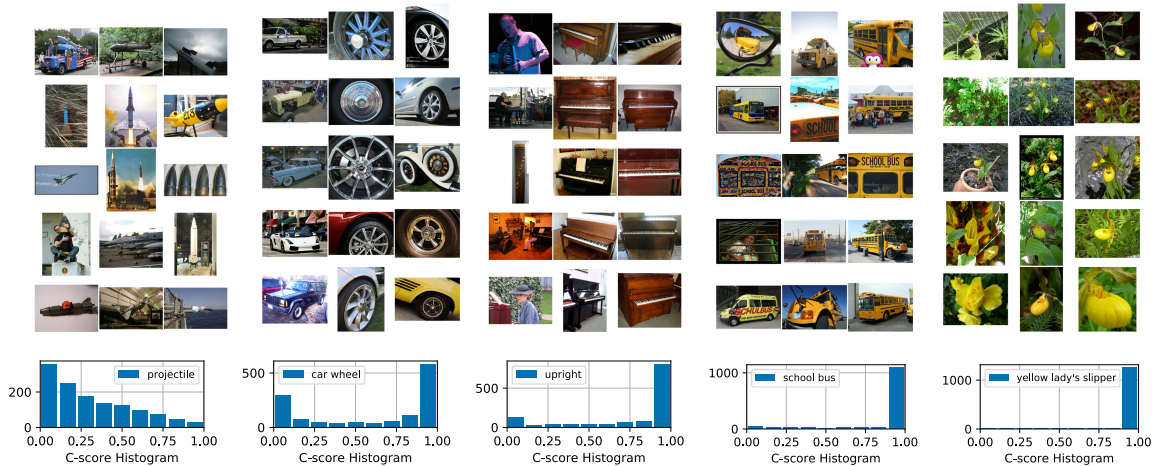


Figure 6. Example images from ImageNet. The 5 classes are chosen to have representative per-class C-score mean–standard-deviation profiles, as shown in Figure 5a. For each class, the three columns show sampled images from the (C-score ranked) top 99%, 35%, and 1% percentiles, respectively. The bottom pane shows the histograms of the C-scores in each of the 5 classes.

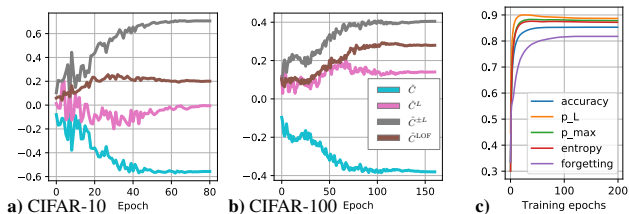


Figure 7. (a-b) Spearman rank correlation between C-score and distance based proxies using learned hidden representations. (c) Spearman rank correlation between C-score and learning speed based proxies on CIFAR-10.

We further evaluate the proxies using the penultimate layer of the network as a representation of an image: $\hat{C}_h^{\pm L}$, \hat{C}_h^L , \hat{C}_h and \hat{C}_h^{LOF} , with the subscript h indicating distance in hidden space. In particular, we train neural network models with the same specification on the full training set. We plot the correlation between the C-score and the proxy based on the learned representation at each epoch as a function of training epoch in Figure 7a,b. For both data sets, the proxy score that correlates best with the C-score is $\hat{C}_h^{\pm L}$ (grey), followed by \hat{C}_h^{LOF} (brown), then \hat{C}_h^L (pink) and \hat{C}_h (blue). Clearly, appropriate use of labels helps with the ranking.

The results reveal interesting properties of the hidden representation. One might be concerned that as training progresses, the representations will optimize toward the classification loss and may discard inter-class relationships that could be potentially useful for other downstream tasks (Scott et al., 2018). However, our results suggest that $\hat{C}_h^{\pm L}$ does not diminish as a predictor of the C-score, even long after training converges. Thus, at least some information concerning the relation between different examples is retained in the representation, even though intra- and inter-class similarity is not very relevant for a classification model. To the extent that the hidden representation—crafted through a discriminative loss—preserves class structure, one might expect that the C-score could be predicted without label reweighting; however, the poor performance of \hat{C}_h suggests otherwise.

Figure 8 visualizes examples ranked by the class weighted local density scores in the input and learned hidden space, respectively, in comparison with examples ranked by the C-score. The ranking calculated in the input space relies heavily on low level features that can be derived directly from the pixels like strong silhouette. The rankings calculated from the learned hidden space correlate better with C-score, though the visualizations show that they could not faithfully detect the dense cluster of highly uniform examples with high C-scores.

In summary, while pairwise distance based proxies are very intuitive to formulate, in practice, the rankings are very sensitive to the underlying distance metrics.




Figure 8. (Left pane) The 3 blocks show examples from CIFAR-10 “automobile” ranked by $\hat{C}_h^{\pm L}$, $\hat{C}_h^{\pm L}$ and the C-score, respectively. The three columns in each block shows the top, middle and middle ranked examples, respectively. (Right pane) Examples from CIFAR-100 “bear” shown in the same layout.

5.2. Learning Speed Based Proxies

Inspired by our observations in the previous section that the speed-of-learning tends to correlate with the C-score rankings, we instead focus on a class of learning-speed based proxies that have the added bonus of being trivial to compute. Intuitively, a training example that is consistent with many others should be learned quickly because the gradient steps for all consistent examples should be well aligned. One might therefore conjecture that strong regularities in a data set are not only better learned at asymptote—leading to better generalization performance—but are also learned *sooner* in the time course of training. This *learning speed* hypothesis is nontrivial, because the C-score is defined for a held-out instance following training, whereas learning speed is defined for a training instance during training. This hypothesis is qualitatively verified from Figure 10. In particular, the cyan examples having the lowest C-scores are learned most slowly and the purple examples having the highest C-scores are learned most quickly. Indeed, learning speed is monotonically related to C-score bin.

Figure 7c shows a quantitative evaluation, where we compute the Spearman’s rank correlation between the C-score of an instance and various proxy scores based on learning speed. In particular, we test *accuracy* (0-1 correctness), p_L (softmax confidence on the correct class), p_{max} (max softmax confidence across all classes) and *entropy* (negative entropy of softmax confidences). We use *cumulative* statistics which average from the beginning of training to the current epoch because the cumulative statistics yield a more stable measure—and higher correlation—than statistics based on a single epoch. We also compare to a *forgetting-event* statistic (Toneva et al., 2019), which is simply a count of the number of transitions from “learned” to “forgotten” during training. All of our proxies show strong correlation with the C-score: p_L reaches $\rho \approx 0.9$ at the peak; p_{max} and *entropy* perform similarly, both slightly worse than p_L . The *forgetting event* statistic slightly underperforms our proxies and takes a larger number of training epochs to reach its peak correlation. We suspect this is because forgetting events hap-

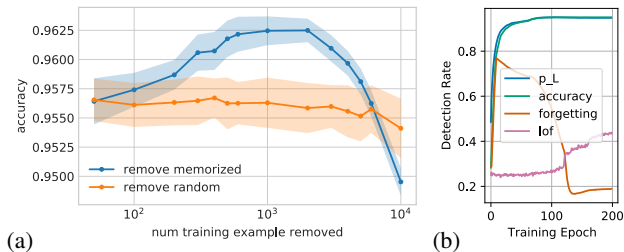


Figure 9. (a) Model performance on SVHN when certain number of examples are removed from the training set. (b) Detection rate of label-flipped outliers on CIFAR-10.

pen only *after* an example is learned, so unlike the proxies studied here, forgetting statistics for hard examples cannot be obtained in the earlier stage of training.

6. Application

By characterizing the structural regularities in large scale datasets, the C-score provides powerful tools for analyzing data sets, learning dynamics, and to diagnose and potentially improve learning systems. In this section, we provide several illustrative applications along this line.

In the first example, we demonstrate the effects of removing the irregular training examples. In Figure 9a, we show the performance of models trained on the SVHN (Netzer et al., 2011) training set as a function of the number of lowest C-score examples removed. For comparison, we show the performance with the same number of *random* examples removed. We found that the model performance improves as we remove the lowest ranked training examples, but it eventually deteriorates when too many (about 10^4) training examples are removed. This deterioration occurs because the C-score typically ranks mislabeled instances toward the bottom, followed by—at least in this data set—correctly labeled but rare instances. Although the mislabeled instances have no utility, the rare instances do, causing a drop in performance as more rare instances are removed. On data sets with fewer mislabelings, such as CIFAR-10, we did not observe an advantage of removing low-ranked examples versus removing random examples.

To quantitatively evaluate the outlier identification rate, we construct a modified dataset by corrupting a random fraction $\gamma = 25\%$ of the CIFAR-10 training set with random label assignments, so that we have the ground-truth indicators for the outliers. We then identify the fraction γ with the lowest ranking by our two most promising learning-speed based C-score proxies—cumulative accuracy and p_L . Figure 9b shows the detection rate—the fraction of the lowest ranked examples which are indeed outliers; the two C-score proxies successfully identify over 95% of outliers. This is consistent with previous work (Pleiss et al., 2020) showing the loss

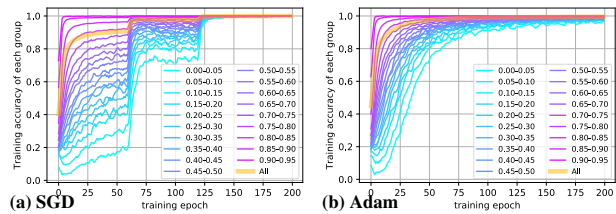


Figure 10. Learning speed of CIFAR-10 examples grouped by C-score. The thick transparent curve shows the average accuracy over the entire training set. SGD achieves test accuracy 95.14%, Adam achieves 92.97%.

curves could be informative at detecting noisy examples. We also evaluated the forgetting-event statistic (Toneva et al., 2019) and the local outlier factor (LOF) (Breunig et al., 2000) algorithm based on distances in the hidden space, but neither is competitive.

In the final example, we demonstrate using the C-score to study the behavior of different optimizers. For this study, we partition the CIFAR-10 training set into subsets by C-score. Then we record the learning curves—model accuracy over training epochs—for each set. Figure 10 plots the learning curves for C-score-binned examples. The left panel shows SGD training with a stagewise constant learning rate, and the right panel shows the Adam optimizer (Kingma & Ba, 2015), which scales the learning rate adaptively. In both cases, the groups with high C-scores (magenta) generally learn faster than the groups with low C-scores (cyan). Intuitively, the high C-score groups consist of mutually consistent examples that support one another during training, whereas the low C-score groups consist of irregular examples forming sparse modes with fewer consistent peers. In the case of true outliers, the model needs to memorize the labels individually as they do not share structure with any other examples.

The learning curves have wider dispersion in SGD than in Adam. Early in SGD training where the learning rate is large, the examples with the lowest C-scores barely learn. In comparison, Adam shows less spread among the groups and as a result, converges sooner. However, the superior convergence speed of adaptive optimizers like Adam does not always lead to better generalization (Wilson et al., 2017; Keskar & Socher, 2017; Luo et al., 2019). We observe this outcome as well: SGD with a stagewise learning rate achieves 95.14% test accuracy, compared to 92.97% for Adam. The visualization generated with the help of the C-score provides an interesting perspective on the difference between the two cases with different generalization performances: SGD with stagewise learning rate effectively enforces a sort of curriculum in which the model focuses on learning the strongest regularities first. This curriculum could help the model building a more solid representation based on domain regularities, when compared to Adam that

learns all examples at similar pace.

7. Discussion

We formulated a *consistency profile* for individual examples in a data set that reflects the probability of correct generalization to the example as a function of training set size. This profile has strong ties to generalization theory as it essentially measures the per-instance generalization. We distilled the profile into a scalar C-score, which provides a total ordering of the instances in a data set by essentially the sample complexity—the amount of training data required—to ensure correct generalization to the instance. By studying the estimated scores on real world datasets, we show that this formulation captures well the basic intuitions about data regularity in both human and machine learning.

To leverage the C-score to analyze structural regularities in complex data sets, we derived a C-score estimation procedure and obtained C-scores for examples in MNIST, CIFAR-10, CIFAR-100, and ImageNet. The C-score estimate helps to characterize the continuum between a densely populated mode consisting of aligned, centrally cropped examples with unified shape and color profiles, and sparsely populated modes of just one or two instances.

We further studied two variants of computationally efficient proxies to the C-score. We found that the pairwise distance based proxies are sensitive to the underlying distance metrics, while the learning speed based proxies generally provide better correlation with the C-score.

We demonstrate examples of potential applications of the C-score as analytical tools to inspect large scale datasets and the learning systems trained on the data, which provides insights to the otherwise complicated and opaque systems. In particular, we show that the C-score could be used to identify outliers and provide detailed analysis of the learning dynamics when comparing different optimizers.

One feature of our formulation is that the C-score depends on the neural network architecture, and more generally on the learning algorithm. Just like how a math major and a music major might have different opinions on the difficulty of courses, different neural networks could have different inductive biases a priori, and the C-score captures this fact. In practice, we found that the C-score estimations are consistent among commonly used convolutional networks, potentially because they are not that different from each other. In particular, we compared the Inception based estimation on CIFAR-10 with ResNet-18, VGG-11 and VGG-16, and found the Spearman’s ρ correlations are above 0.91. Recently some new *convolution-free* architectures based on attention mechanism (Dosovitskiy et al., 2021) or dense connections (Tolstikhin et al., 2021; Melas-Kyriazi, 2021; Touvron et al., 2021) emerged and achieved similar per-

formance as their convolutional counterparts on standard image classification benchmarks. We leave it as future work to conduct extensive comparison on more diverse architectures and emerging algorithms such as finetuning after self-supervised learning (Chen et al., 2020a;b; Grill et al., 2020; Caron et al., 2021), and so on.

In the 1980s, neural nets were touted for learning *rule-governed behavior* without explicit rules (Rumelhart & McClelland, 1986). At the time, AI researchers were focused on constructing expert systems by extracting explicit rules from human domain experts. Expert systems ultimately failed because the diversity and nuance of statistical regularities in a domain was too great for any human to explicate. In the modern deep learning era, researchers have made much progress in automatically extracting regularities from data. Nonetheless, there is still much work to be done to understand these regularities, and how the consistency relationships among instances determine the outcome of learning. By defining and investigating a consistency score, we hope to have made some progress in this direction. We have released the precomputed C-scores on standard deep learning benchmark datasets to foster future research along this direction.

Code and Pre-computed C-scores

We provide code implementing our C-score estimation algorithms, and pre-computed C-scores and associated model checkpoints for CIFAR-10, CIFAR-100 and ImageNet (downloadable from <https://pluskid.github.io/structural-regularity/>). The exported files are in Numpy’s data format saved via `numpy.savez`. For CIFAR-10 and CIFAR-100, the exported file contains two arrays `labels` and `scores`. Both arrays are stored in the order of training examples as defined by the original data sets found at <https://www.cs.toronto.edu/~kriz/cifar.html>. The data loading tools provided in some deep learning library might not be following the original data example orders, so we provided the `labels` array for easy sanity check of the data ordering.

For ImageNet, since there is no well defined example ordering, we order the exported scores arbitrarily, and include a script to reconstruct the data set with index information by using the filename of each example to help identify the example-score mapping.

Acknowledgements

We thank Vitaly Feldman for guidance on simulation design and framing of the research, Samy Bengio for general comments and feedback, and Yoram Singer for making the collaboration possible.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48. ACM, 2009.
- Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 93–104, 2000.
- Carlini, N., Erlingsson, U., and Papernot, N. Prototypical examples in deep learning: Metrics, characteristics, and utility. Technical report, OpenReview, 2018.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., and Joulin, A. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020a.
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., and Hinton, G. Big Self-Supervised models are strong Semi-Supervised learners. In *Advances in Neural Information Processing Systems*, 2020b.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- Feldman, V. Does learning require memorization? A short tale about a long tail. In *ACM Symposium on Theory of Computing (STOC)*, 2020.
- Feldman, V. and Zhang, C. What neural networks memorize and why: Discovering the long tail via influence estimation. In *Advances in neural information processing systems*, 2020.
- Grill, J.-B., Strub, F., Alché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., Piot, B., Kavukcuoglu, K., Munos, R., and Valko, M. Bootstrap your own latent: A new approach to Self-Supervised learning. In *Advances in Neural Information Processing Systems*, 2020.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pp. 1225–1234. PMLR, 2016.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.
- Keskar, N. S. and Socher, R. Improving generalization performance by switching from adam to sgd. *arXiv preprint arXiv:1712.07628*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical Report TR-2009, University of Toronto, 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Luo, L., Xiong, Y., Liu, Y., and Sun, X. Adaptive gradient methods with dynamic bound of learning rate. In *International Conference on Learning Representations*, 2019.
- Mangalam, K. and Prabhu, V. U. Do deep neural networks learn shallow learnable examples first? In *ICML 2019 Workshop on Identifying and Understanding Deep Learning Phenomena*, 2019.
- Melas-Kyriazi, L. Do you even need attention? a stack of feed-forward layers does surprisingly well on imagenet. *arXiv preprint arXiv:2105.02723*, 2021.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- Pleiss, G., Zhang, T., Elenberg, E. R., and Weinberger, K. Q. Detecting noisy training data with loss curves. In *International Conference on Learning Representations*, 2020.

- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- Rumelhart, D. E. and McClelland, J. L. *On Learning the Past Tenses of English Verbs*, pp. 216–271. MIT Press, Cambridge, MA, USA, 1986.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252, 2015.
- Saxena, S., Tuzel, O., and DeCoste, D. Data parameters: A new family of parameters for learning a differentiable curriculum. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 11093–11103. Curran Associates, Inc., 2019.
- Scott, T., Ridgeway, K., and Mozer, M. C. Adapted deep embeddings: A synthesis of methods for k-shot inductive transfer learning. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 76–85. Curran Associates, Inc., 2018.
- Tan, M. and Le, Q. V. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- Tolstikhin, I., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Keysers, D., Uszkoreit, J., Lucic, M., et al. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.
- Toneva, M., Sordoni, A., Combes, R. T. d., Trischler, A., Bengio, Y., and Gordon, G. J. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations*, 2019.
- Touvron, H., Bojanowski, P., Caron, M., Cord, M., El-Nouby, A., Grave, E., Joulin, A., Synnaeve, G., Verbeek, J., and Jégou, H. Resmlp: Feedforward networks for image classification with data-efficient training. *arXiv preprint arXiv:2105.03404*, 2021.
- Wilson, A. C., Roelofs, R., Stern, M., Srebro, N., and Recht, B. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, pp. 4148–4158, 2017.
- Wu, Z., Nagarajan, T., Kumar, A., Rennie, S., Davis, L. S., Grauman, K., and Feris, R. Blockdrop: Dynamic inference paths in residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8817–8826, 2018.