
On the Effects of Artificial Data Modification

Antonia Marcu¹ Adam Prügel-Bennett¹

Abstract

Data distortion is commonly applied in vision models during both training (e.g. methods like MixUp and CutMix) and evaluation (e.g. shape-texture bias and robustness). This data modification can introduce artificial information. It is often assumed that the resulting artefacts are detrimental to training, whilst being negligible when analysing models. We investigate these assumptions and conclude that in some cases they are unfounded and lead to incorrect results. Specifically, we show current shape bias identification methods and occlusion robustness measures are biased and propose a fairer alternative for the latter. Subsequently, through a series of experiments we seek to correct and strengthen the community’s perception of how augmenting affects learning of vision models. Based on our empirical results we argue that the impact of the artefacts must be understood and exploited rather than eliminated.

1. Motivation

Augmentation is commonplace when training models. It is a form of data modification where samples are artificially distorted to create larger training sets. Apart from augmentative purposes, data modification is also used for a wide range of model analysis methods. Most recently, distortion-based approaches have been adopted when trying to answer key machine learning questions. To this end, MixUp-like distortions (Zhang et al., 2018) were proposed for empirically predicting generalisation (Schiff et al., 2021; Natekar & Sharma, 2020; Lassance et al., 2020). Thus, data modification is becoming increasingly popular, but little attention is paid to the secondary effects of this practice. As we will demonstrate, *our current understanding of the effects of data modification lies on fundamentally flawed assumptions*. This impacts not only our perception of what features are

important to our models, but also the correctness of the distortion-based approaches we propose as a field.

In this paper we study the assumptions and implications that arise through artificially distorted data. From the model analysis perspective, we take occlusion robustness and shape bias identification methods as examples of where modified data is used. On the training side, we focus on some instances of Mixed Sample Data Augmentation (MSDA), where two images are combined to obtain a new training sample. Visual illustrations of each can be found in Figure 1. We explore some of the *side-effects* of data modification and demonstrate that this practice has resulted in the creation of biased model interpretation tools and poorly informed theories. More specifically, we study a number of assumptions which we show are erroneous and which lie at the heart of the methods we briefly introduce below. Contesting these assumptions has broader implications on the community’s perception of what aspects of the data are important and how distortions can be further used to understand generalisation.

Shape-texture bias: Deep models are known to be sensitive to distribution shifts (Alaiz-Rodríguez & Japkowicz, 2008; Cieslak & Chawla, 2009; Engstrom et al., 2019) and interventions that are imperceptible to humans (Szegedy et al., 2013; Goodfellow et al., 2014). It has been argued that this is intimately linked to networks tending to use texture rather than shape information (Brendel & Bethge, 2019; Geirhos et al., 2019). Recently, input distortions have become a popular way of assessing a model’s texture bias. To this end, images are divided into a grid and the resulting patches are randomly shuffled such that information is preserved locally, while the global shape is altered (Shi et al., 2020; Mummadi et al., 2021; Luo et al., 2019; Zhang & Zhu, 2019). *It is implicitly assumed that patch-shuffling does not introduce misleading shape or texture that could affect model evaluation*. As such, if a model’s accuracy drops when evaluated on patch-shuffled images, this degradation in performance is entirely attributed to the model’s bias for shape information. Thus, any side-effects of the data manipulation process are considered negligible.

Occlusion robustness: Commonly, occlusion robustness is concerned with the amount of information that can be hidden from a model without affecting its ability to classify (Tang et al., 2018; Rajaei et al., 2019). A widely adopted proxy for measuring occlusion robustness is through the raw

¹Vision, Learning and Control research group, University of Southampton. Correspondence to: Antonia Marcu <am1g15@soton.ac.uk>.

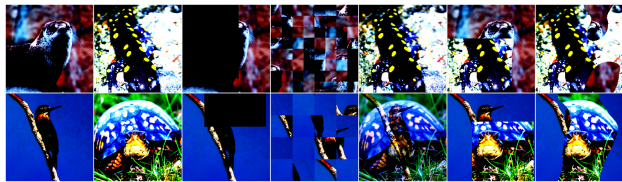


Figure 1. Examples of image distortions. Left to right: source image 1, source image 2, patch occlusion, shuffled grid, MixUp, CutMix, FMix. For mixing augmentations, the first row was generated with a mixing factor of 0.2, while the second one with 0.5. For patch-shuffling, a grid size of 8×8 was used for the top row and 4×4 for the bottom one.

accuracy obtained after superimposing a rectangular patch on an image (Chun et al., 2020; Fawzi & Frossard, 2016; Yun et al., 2019; Zhong et al., 2020; Kokhlikyan et al., 2020). We refer to this approach as CutOcclusion. Just as with shape bias this method relies on the information introduced not to interfere with a model’s learnt representations, such that a decrease in performance can be directly attributed to a lack of robustness. Thus, using CutOcclusion, one implicitly assumes that artefacts do not interfere with the results of robustness evaluation.

Data augmentation studies: In statistical learning, training with augmented data is termed Vicinal Risk Minimisation (VRM) (Vapnik, 1999; Chapelle et al., 2001) and is seen as injecting prior knowledge about the neighbourhood of the data samples. The intuition behind augmentation caused researchers to interpret its effect through the similarity between original and augmented data distributions. This perspective is often challenged by methods which despite generating samples that do not appear to fall under the distribution of natural images, lead to strong learners. Gontijo-Lopes et al. (2020) argue it is the *perceived* distribution shift that needs to be minimised, while maximising the sample vicinity. Formalising these concepts, they introduce augmentation “diversity” and “affinity”. Diversity is defined as the training loss when learning with artificial samples, while affinity quantifies the difference between the accuracy on original test data and augmented test data, for a reference model. The latter penalises augmentations that introduce artificial information to which the model is not invariant, *implicitly assuming that training with that information is detrimental to generalisation*. Thus, in contrast to the evaluation methods mentioned above, the artefacts are considered non-negligible when training with distorted data.

In summary, it is currently assumed that the artefacts introduced by changes in the data are negligible when evaluating models, while those introduced when training are important and undesirable. These assumptions implicitly shape the community’s perception of how machine learning works. Does the artificial information added by analysis methods

not have major side-effects or does it lead to biased results? Conversely, are the artefacts important when training with modified data? Do they cause models to learn better or worse representations?

We set out to answer these questions. We find that results can be misleading when not accounting for the secondary effects of data manipulation, especially in comparative studies. Taking an oversimplified example to illustrate this for robustness, we can imagine a binary cat–truck image classification problem and two models: model *A*, which identifies cats solely by the presence of pointy ears and model *B*, which has a more holistic approach. Generally, masking out the ears will cause model *A* to misclassify cats and we would consider this model not robust to occlusion, while model *B* will continue to correctly classify them. However, if the ears are covered with a large rectangle that introduces horizontal and vertical edges strongly associated with the truck class, this will cause model *B* to also misclassify. In this case, because the misclassification is not caused by the absence of a feature but rather by the presence of a distractor, we would still consider model *B* robust to occlusion although its performance degrades. In such a case CutOcclusion would be unable to distinguish between a model that incorrectly classifies because of lack of information, or because of the presence of confounding artefacts, making it an incorrect proxy for measuring occlusion robustness. Thus, the side-effects of data distortion must be taken into account to create fair evaluation methods.

We start by introducing a measure that indicates the existence of such side-effects which we then use to disprove the negligibility of data distortion artefacts when evaluating models. We then build on the importance of artefacts to construct empirical counter-examples which disprove common beliefs in the augmentation literature, and highlight the importance of understanding the changes data manipulation introduces. Our contributions are as follows:

- We introduce a quantity for highlighting the interference of artificially introduced information with a model’s learnt representations (Section 2);
- We show that increasingly popular model interpretation and analysis methods are biased (Section 2);
- We propose a fairer alternative for measuring occlusion robustness (Section 3);
- We show that in contrast to what is widely assumed, not preserving the data distribution can lead to learning better representations (Section 4).

An immediate application of our work is in robustness evaluation of real-world vision models. Correctly evaluating the ability of a model to perform in the presence of visual

occluders is of paramount importance for safety-critical applications such as self-driving cars and medical diagnosis and intervention. While the impact of our practical contributions is relevant to the community, we believe combating erroneous research directions is more important for the future development of the field. Correctly understanding the increasingly popular mixed-sample augmentation is essential for trusting its usage in sensitive applications where the data can be out of distribution. Moreover, studying the effects of image distortion can shed further light on what *models* perceive as being within distribution. Most importantly, we believe this could set a new direction in capturing the relationship between data and learned representations, which could ultimately play a role in understanding generalisation.

2. Are Artefacts Negligible when Analysing Classifiers?

In this section we show that artificially introduced artefacts may not be negligible, and distorting data at evaluation time could have side-effects not previously considered. Specifically, the artefacts can interfere with the representations learnt by the model, which in turn leads to incorrect evaluation. We highlight this interference by showing that *the distortion can be consistently associated with a particular class in an image classification task*. We do so by looking at the increase in misclassifications per predicted category; from the number of incorrect predictions of a model evaluated on modified data, we subtract the incorrect predictions when testing on original data. If, across multiple runs, there is a significant increase for a specific class, this indicates that the distortion introduces features the model associates with that class. We refer to this phenomenon as “data interference”. By a model “run” we refer to a model instance trained with a different seed for the initialisation of weights and for the randomised data augmentation. For computing the *DI* index we train 5 different instances of the same model and only consider that data interference occurs when runs consistently display a bias for the same class.

Considering only positive differences, we denote the increase in the percentage of misclassifications for predicted class c for a run r by c^r . Note that the class is taken to be that predicted by the classifier. To keep the score within a consistent range across data sets, we scale c^r by the number of classes. We define the Data Interference (*DI*) index as

$$\mathbb{E}_r \left[\frac{c_{max}^r}{\sum_c c^r} \mathbb{E}_{r'} \left[c_{max}^{r'} \right] \right],$$

where c_{max} is that of the class with the highest mean increase across all runs. The *DI* index measures the proportion represented by the dominant class weighted by its

average increase across runs. A high index value indicates a sharp increase for a particular class which is consistent across runs. We associate this with an overlap between introduced artefacts and learnt representations, thus highlighting the *side effects of distorting data for model evaluation*.

To obtain models with different behaviours in a controlled manner, we make use of Mixed Sample Data Augmentation (MSDA). Since it is sufficient to identify some common cases in which models are disfavoured, we choose to reduce our environmental impact by restricting the analysis to simple MSDAs that combine images without incurring additional computation time or external models. As will be argued in Section 3, we expect the unfairness to be present in most settings, thus the exact choice of augmentation is irrelevant. We focus on two popular MSDAs, MixUp and CutMix (Yun et al., 2019). MixUp linearly interpolates between two images to obtain a new training example, while CutMix masks out a rectangular region of an image with the corresponding region of another image. Besides the aforementioned methods, we also employ FMix (Harris et al., 2020) due to its irregularly shaped masks sampled from Fourier space, which will play an important role in our analysis. Note that although the masking methods sample the size of the occluding patch from the same distribution, in CutMix part of the rectangle can be outside the image, which leads to less occluded samples overall. We refer to models by the augmentations they were trained with and use “basic” to label the models trained without MSDA.

Throughout the paper, we do 5 runs of each experiment with PreAct-ResNet18 (He et al., 2016) as the default architecture. We include a number of results for BagNet (Brendel & Bethge, 2019), VGG (Simonyan & Zisserman, 2015), WideResNet-28-10 (Zagoruyko & Komodakis, 2016) and DenseNet-BC-190 (Huang et al., 2017) in the Appendices (e.g. B.7, B.8) where we also include a cross-architecture analysis.

The main data sets we report results on are CIFAR-10/100 (Krizhevsky, 2009), Tiny ImageNet (Stanford, 2015), FashionMNIST (Xiao et al., 2017), and ImageNet (Rusakovsky et al., 2015). For ImageNet we use pretrained ResNet-101 models made publicly available by Harris et al. (2020). Note that for ImageNet we do not report results with repeated runs since only one model per augmentation was available. For full experimental details, see Appendix A. We provide the code at <https://github.com/AntoniaMarcu/Data-Modification>.

2.1. Shape Bias

Using the *DI* measure, we want to show the existence of side-effects that occur when measuring shape bias by the accuracy after patch-shuffling images. This bias makes the patch-shuffling evaluation method unreliable. For assessing

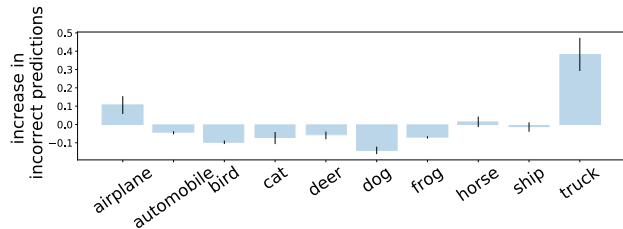


Figure 2. Change in incorrect predictions for five basic CIFAR-10 models when evaluated on patch-shuffled versus original samples. The models consistently associate patch-shuffled images with the “truck” class as indicated by the high increase in incorrect predictions for this class.

shape bias through sample manipulation, the standard procedure is to choose between dividing the image in 4, 16 or 64 patches to be shuffled. Since FashionMNIST images are smaller, we choose a 2×2 grid, while for CIFAR-10/100, Tiny ImageNet and ImageNet we use a 4×4 grid. The results for this experiment are given in Table 1. To provide a sense of scale note that when distorting with Gaussian noise, the basic (trained without MSDA) model has a DI index of 0.09 ± 0.02 (see Appendix B.2). There is no fixed threshold that indicates data interference as this would require an objective definition of bias. Instead, the results are meant to be interpreted in a comparative manner. As such, we are interested in determining if there exists a clear gap between different models across the same task. We observe that for all four data sets such a gap exists, with the basic and MixUp models generally having high index values.

Reiterating, a large index indicates that a models tends to associate the features artificially introduced by patch-shuffling with a certain class. Taking a closer look at the distribution of misclassifications for CIFAR-10 we can see that in the case of the basic model this class is “Truck” (Figure 2). The result is not at all surprising, given that the strong horizontal and vertical edges are highly indicative of this class. Similar observations can be made for other data sets (Appendix B.4). Thus, we believe the patch-shuffling approach could cause models which are not invariant to strong horizontal and vertical edges to falsely appear to rely more on shape information. To verify this claim, we first need to check if sensitivity to patch shuffling is necessarily a symptom of shape bias. Thus, we want to verify if models with similar shape bias can have different DI index values; if so, patch-shuffling would be an unfair basis for evaluation of the shape bias of such models.

Is a model necessarily more affected by patch-shuffling if it has a higher shape bias? To explore this we use an alternative, well-known, method of measuring shape and texture, which we call the GST approach. This implies evaluating ImageNet models on the Geirhos Style-Transfer (GST) data

Table 1. DI index for PreAct-ResNet18 on grid-shuffled images for four different types of models. The higher the index, the more a model would be disadvantaged by side-effects caused by the distortion. For each data set, we show the results with the lowest average in blue and the ones with the highest average in green. Information introduced by shuffling tends to interfere less with FMix and CutMix models.

	basic	MixUp	FMix	CutMix
CIFAR-10	2.82 \pm 0.44	2.40 \pm 0.59	0.59 \pm 0.12	0.31 \pm 0.10
CIFAR-100	0.99 \pm 0.27	0.88 \pm 0.24	0.18 \pm 0.10	0.09 \pm 0.04
Fashion	1.23 \pm 0.15	2.42 \pm 0.92	1.06 \pm 0.23	0.68 \pm 0.11
Tiny	1.28 \pm 1.13	0.57 \pm 0.11	0.67 \pm 0.10	0.25 \pm 0.11
ImageNet	0.82	1.49	0.58	—

Table 2. Accuracy of augmentation-trained models on the GST data set when the label is taken to be either the shape or texture. Higher accuracy indicates a stronger bias. There is no clear link between masking and low texture bias.

	ImageNet		Tiny ImageNet	
	Shape	Texture	Shape	Texture
basic	20.31	53.28	10.56 \pm 0.65	26.04 \pm 1.77
MixUp	24.14	60.31	12.02 \pm 0.33	27.77 \pm 1.56
FMix	21.25	53.43	10.40 \pm 0.39	19.90 \pm 2.12
CutMix	—	—	10.54 \pm 0.38	23.72 \pm 2.42

set (Geirhos et al., 2019). GST contains artificially generated images where the shape belongs to one class and the texture to another. The data set was specifically designed for shape and texture bias identification and it represents a gold standard in the field. The limitation associated with it is that it can only assess the bias of models trained on data sets with which it is compatible. For this reason, alternative universal methods such as patch-shuffling were proposed. The data set contains 16 coarse categories that encompass a number of ImageNet classes to which they are mapped. The bias of the models is given by the accuracy obtained when the label is set to either the shape or texture. Using this well-known method we want to find models which have similar biases but different DI indices when patch-shuffling. This would indicate that sensitivity to shuffling is not necessarily linked to increased shape bias, which in turn would mean that models evaluated using patch-shuffling can artificially appear more shape biased.

The results in Table 2 show that the basic model does not have a higher shape bias than masking methods although it has a significantly higher DI index compared to these, as we have seen in Table 1. We repeat the same experiment on the Tiny ImageNet data set. Geirhos et al. (2019) used WordNet (Miller, 1995) to map the 1000 categories to the 16 classes of the GST data set. We used the same method

Table 3. DI index when occluding with black patches. We show the results with the lowest average in blue and the ones with the highest average in green. For each data set, there exists a non-negligible gap between masking and non-masking training regimes, with models like MixUp and basic being disadvantaged due to higher data interference.

	basic	MixUp	FMix	CutMix
CIFAR-10	1.25 \pm 0.17	0.47 \pm 0.11	0.11 \pm 0.04	2.20 \pm 0.81
CIFAR-100	1.24 \pm 0.35	0.34 \pm 0.09	0.12 \pm 0.10	1.06 \pm 0.32
FashionMNIST	0.21 \pm 0.08	0.38 \pm 0.06	0.16 \pm 0.05	0.12 \pm 0.05
Tiny ImageNet	0.52 \pm 0.17	0.39 \pm 0.03	0.14 \pm 0.04	3.46 \pm 2.45
ImageNet	0.50	1.50	0.50	—

to create a mapping between Tiny ImageNet and GST. A number of ImageNet categories that belong to the 16 higher-level classes of GST are missing. For this reason, a poorer overall performance is expected and the results could differ slightly given a better fit between the sets. We find again no significant correlation between masking augmentation and texture bias. Table 13 of Appendix B.7 also includes the results for BagNet models, which have smaller receptive fields and so are forced to use more local information. Even in this case, we find a high DI for the basic model and no difference in texture bias compared to MSDA. Thus, a model which is more affected by the side-effects of patch-shuffling is not necessarily more shaped-bias. In other words, the artefacts introduced by patch-shuffling can cause models to have different accuracies on these distorted images, albeit not differing in their shape and texture bias.

We further confirm this conclusion using the GST approach on the two data sets with which it is compatible. Table 4 gives the accuracy when patch-shuffling on Tiny ImageNet and ImageNet. In both cases, for comparable levels of shape bias, different accuracies are obtained. This confirms that *models can appear to have vastly different shape bias when evaluated on randomly rearranged patches, when in reality their actual shape bias is similar*. The sensitivity of the patch shuffling approach to artefacts makes it unfair and unreliable measure of shape bias.

2.2. Occlusion Measurement

We next want to determine whether the same issue applies to occlusion robustness measures. We focus on CutOcclusion, where a rectangular black patch is superimposed on test images and the robustness is given by the resulting accuracy. There is no standardised distortion when measuring CutOcclusion, with the size and positioning of the obstructing patch varying between studies. Most often in prior art a lack of robustness is noted for large occluders (Chun et al., 2020; Zhong et al., 2020). For this reason, we uniformly sample the size of the patch from [0.7, 1], allowing the occluding patch to lie outside the image (as it is done when augment-

Table 4. Accuracy on patch-shuffled images. It is considered in the literature that the higher the accuracy, the more texture-oriented the model is. Unlike evidenced through the GST method (Table 2), the basic model appears to have a higher shape bias.

	ImageNet	Tiny ImageNet
basic	49.49	13.53 \pm 2.02
MixUp	52.16	17.81 \pm 0.16
FMix	56.20	34.53 \pm 0.62
CutMix	—	42.30 \pm 0.29

ing with CutMix and CutOut (DeVries & Taylor, 2017)). This allows us to capture both the cases in which either the centre or the border area is masked out but requires a non-uniform distribution to counter for the patches existing outside the image.

We perform the same experiment as before, where the DI index is now measured when testing on rectangle-occluded images and present results in Table 3. Once again, we can observe a significant gap in the DI index for each data set. This indicates that some models will again be disadvantaged. We find that data interference is present for different architectures, when overlapping patches from external images or using differently shaped masks (Appendix B.6). Thus, the result of CutOcclusion and its variants is highly dependent on the problem at hand. Just as for randomly shuffling tiles, by occluding images using a particularly shaped patch, one implicitly measures a model’s affinity to certain features. In other words, models for which a distinctive feature is artificially introduced when performing CutOcclusion are going to be penalised by this method even though the feature is perfectly valid in a normal classification setting. This deems such methods inappropriate for fairly assessing robustness and texture bias.

To demonstrate that CutOcclusion is sensitive to the occluder specifics, we superimpose patches taken from a different data set and compare the results to those obtained when occluding with black patches only. For clarity, Figure 3 presents the results for the least and most robust models. The experiment indicates that it is sufficient to simply change the colour pattern of the patch for CutOcclusion to give different results outside the margin of error.

Moreover, this is not the only limitation of CutOcclusion. Another problem that occurs when looking purely at post-masking accuracy is that weaker models would erroneously appear less robust. We show this by reversing the problem: we evaluate the same model on two different subsets of the CIFAR-100 data set: typical and tail images as categorised by Feldman & Zhang (2020). They consider a train-test

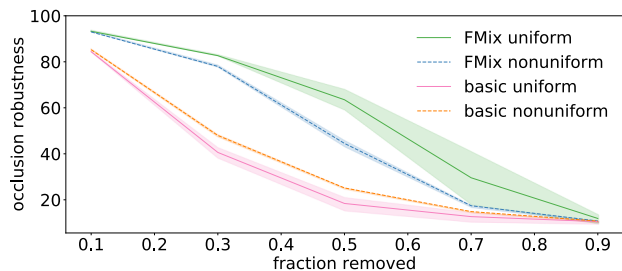


Figure 3. CutOclusion when occluding with black patches (uniform) and patches taken from other images (nonuniform). CutOclusion gives different robustness results for the two types of colour patterns, which means that the measure is sensitive to the characteristics of the occluder.

sample pair to belong to the tail of the data distribution if the test sample is correctly classified when a model is trained with the train sample, and incorrectly without it. CutOclusion would indicate that models are significantly more robust to occluding typical examples. However, a closer analysis makes us doubt this conclusion. The raw accuracy on both train and test data for tail examples is lower than for the typical ones. In fact, the performance when occluding images decreases at the same *rate* for the two subsets, indicating similar robustness.

We have seen that the results obtained using CutOclusion are sensitive to a number of factors. To better reflect how much information can be hidden from a model without affecting its performance, a fair alternative should be invariant to the model’s goodness-of-fit and to the characteristics of the occluder. Concerning the latter, a related observation was made by Hooker et al. (2019) who note the pitfalls of manipulating data to determine feature importance. They point out that when simply superimposing uniform patches over image features, it is difficult to assess how much of the reduction in accuracy is caused by the absence of those features and how much is due to images becoming out of distribution. To address this, the most important features identified by an estimator are masked out both on train and test data, closing the gap between the two sets. Hooker et al. then train and evaluate models on the newly generated images. Unlike with interpretability, the subject of occlusion robustness studies is the model itself, which makes training with a modified version of the data inviable. In the following section we explore ways of overcoming this bias when measuring occlusion robustness.

3. What are Fairer Alternatives?

We propose a simple, carefully defined measure that aims to *decouple the model’s edge bias and generalisation ability from the occlusion robustness*. We refer to our measure

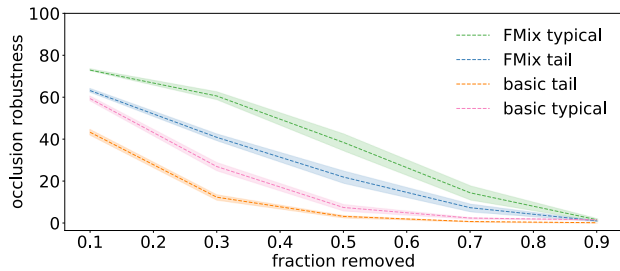


Figure 4. CutOclusion for the basic and FMix models on two subsets of the same data set: tail and typical. CutOclusion reports different robustness results on two subsets although the relative drop in accuracy is the same. This is because CutOclusion does not abstract away the model’s overall performance.

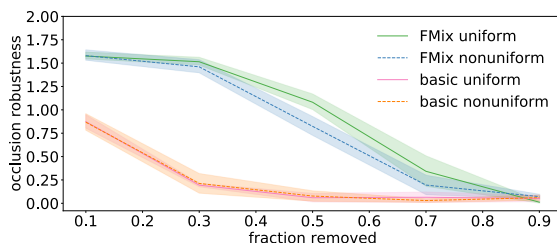


Figure 5. iOclusion when occluding with black patches (uniform) and patches taken from other images (nonuniform). Unlike CutOclusion, our measure gives consistent results when varying the colour pattern of the occluder.

as “interplay occlusion” (iOclusion). Interplay occlusion reflects the change in the interplay between performance on seen and unseen data. Formally,

$$iOclusion_p = \left| \frac{\mathcal{A}(\mathcal{D}_{train}^p) - \mathcal{A}(\mathcal{D}_{test}^p)}{\mathcal{A}(\mathcal{D}_{train}) - \mathcal{A}(\mathcal{D}_{test})} \right|, \quad (1)$$

where $\mathcal{A}(\mathcal{D})$ denotes the accuracy on a given data set \mathcal{D} , and \mathcal{D}^p is the data set resulting from removing $p\%$ pixels of each image. The intuition is that on training data, robust models are less sensitive to the artefacts of the occlusion policy for small levels of occlusion, resulting in a large difference in accuracy from that on unseen data. The performance of both train and test gets close to random as the proportion of occluded data approaches 90%. We expect the gap to fall off quicker for less robust models. This change in interplay is taken with respect to the generalisation gap of the model, such that the quality of the model fit in itself does not interfere with the robustness measure.

A number of factors have to be considered when choosing a masking method for computing \mathcal{D}^p , such as the number of contiguous components or the amount of salient information masked out. In this paper we choose to generate masks using Grad-CAM (Selvaraju et al., 2017). Since

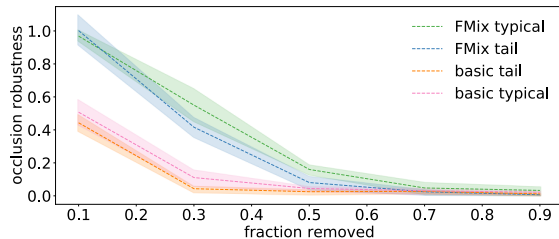


Figure 6. iOcclusion for basic and FMix models on tail and typical data. As opposed to CutOcclusion, measuring robustness with iOcclusion reflects the overlapping robustness level of the models on the two subsets.

Table 5. DI index and occlusion robustness for models trained on CIFAR-10 when obstructing 30% of the image pixels with non-uniform patches. Given in bold is the closest result to that of RM. The higher DI index of RM compared to the other masking methods is causing CutOcclusion to evaluate the robustness of RM as closer to MixUp, rather than FMix or CutMix.

	DI index	CutOcclusion	iOcclusion
basic	1.67 ± 0.17	47.97 ± 0.52	0.21 ± 0.10
MixUp	0.98 ± 0.21	58.65 ± 1.01	0.57 ± 0.18
CutMix	0.14 ± 0.08	76.56 ± 6.36	1.09 ± 0.17
FMix	0.15 ± 0.01	78.00 ± 0.45	1.46 ± 0.07
RM	0.39 ± 0.05	60.79 ± 5.03	1.20 ± 0.23

methods could be sensitive to either occluding contextual or core information, we get an average performance by choosing with equal probability between occluding the *most* or *least* salient $p\%$ pixels for each batch. It must be noted that this method implicitly assumes there could be multiple occluders, which has the downside of incurring a higher computational cost. We also experiment with using rectangular or Fourier-sampled masks and conclude that although random masking makes the process noisier, the exact choice of masking method is of secondary importance as long as the occluder’s granularity is accounted for. Appendix C.5 provides discussion and results on these alternative instances of iOcclusion. Throughout this section, for a fair comparison, when measuring CutOcclusion we do not allow the obstructing patch to lie outside the image. As such, the fraction removed is exact.

Assessing the correctness of such a measure is difficult in the absence of a baseline. We construct varied experiments to attest the validity of our method, highlighting the two important limitations of CutOcclusion that our approach addresses: sensitivity to the specifics of the occluding patch and to model’s generalisation performance. For the curiosity of the reader, we include various additional experiments and analyses in Appendix C, and focus here on the key results. We start by performing the same experiment as

for CutOcclusion measurement where we occlude using black patches and patches taken from a different data set. As shown in Figure 5, iOcclusion using uniform occluders gives similar results to its non-uniform version and thus provides a consistent model evaluation. We then verify that iOcclusion is invariant to the model’s overall performance by evaluating on the same two subsets of the CIFAR-100 data set on which we evaluated CutOcclusion. The results obtained using iOcclusion on tail and typical data are within each other’s margin of error (Figure 6). Thus, iOcclusion abstracts away from the colour of the occluder as well as the model’s generalisation performance.

In addition to not being sensitive to the colour pattern of the patch, a fair measure must also be invariant to the shape of the patch. To empirically confirm iOcclusion reduces the importance of edge information, we aim to obtain a model that is robust to occlusion, but at the same time has a high DI index (it is sensitive to edge information). To this end, we create a variation of FMix, Random Masks (RM), where at the beginning of the training process three masks are randomly sampled from Fourier space. For each batch, one of the three is chosen uniformly at random. When masking with patterned patches the RM model has a higher DI index than FMix, as desired. Table 5 gives results for a fraction of 0.3 pixels covered by a non-uniform occluder. Our measure reflects the robustness of training with RM, situating it close to other masking methods. On the other hand, because CutOcclusion implicitly penalises models with high DI index, according to this measure RM appears almost as sensitive to occlusion as MixUp. Figure 10 shows results for a wider range of fractions.

The unbiased nature of iOcclusion could lead to better understanding and development of training procedures. Equally important, iOcclusion has applicability to real-world vision system deployments where no prior knowledge exists about the possible shapes of the obstructions. While this aspect of generality is the strength of our approach, it must be stressed that when there exists a limited set of known possible occluders, evaluating robustness specifically to them could be safer. For example in an industrial setting or a clinical environment there could be a certain set of objects that could interfere with the subject of interest. We do not propose a universal solution, but rather suggest an alternative to the biased approach for the common scenario in which the environment is not controlled and little is known about all the potential occluders.

The strength of the bias will depend on the data in question and some applications will be more affected than others. We have seen that for natural images this bias does exist. To confirm that we have not just identified isolated cases, we remove the class that has the highest increase in mispredictions and retrain the models on the remaining classes. We

Table 6. Augmentation comparison on CIFAR-10. We consider two variants when calculating diversity. Diversity computes the cross-entropy loss using the label of the majority class, whilst MixDiversity is a linear combination of the two cross-entropy losses. In the literature it is argued that higher affinity and higher diversity are desirable, but we find that this is not necessarily the case.

	Affinity	Diversity	MixDiversity
MixUp	$-12.58_{\pm 0.14}$	$0.41_{\pm 0.01}$	$0.84_{\pm 0.00}$
FMix	$-25.55_{\pm 0.26}$	$0.34_{\pm 0.01}$	$0.65_{\pm 0.00}$

Table 7. DI index measured for non-uniform occlusion when training without the class with highest increase in incorrect predictions. A gap can be noted in this case as well, MixUp being more affected by the distortion’s artefacts as evidenced by the higher *DI* index. This result supports the idea that data interference is not specific to peculiar cases, affecting model evaluation more broadly.

	CIFAR-10	CIFAR-100
MixUp	$0.39_{\pm 0.15}$	$1.22_{\pm 0.19}$
FMix	$0.08_{\pm 0.06}$	$0.50_{\pm 0.21}$

find that the bias is again present (Table 7), but with respect to another class. For example, in the case of CIFAR-10, after removing the “Truck” class, models mispredict rectangle-occluded images as “Boat” (Appendix C.5.1). Thus, the edge artefacts are likely to interfere with learnt representations since they are such fundamental features. From an evaluation perspective, as we have seen, this impacts assessment methods and must be accounted for. From the training perspective, such a widespread data interference would indicate a large perceptual shift when performing MSDA. However, large perceptual shifts are believed to have a negative effect on generalisation, while MSDAs are known to improve performance. Is then the perceptual shift caused by artificially introduced information really detrimental to learning? In the following section we investigate the impact of artefacts on training.

4. Is the Magnitude of the Distribution Shift Important?

Note once again that in this paper we are interested in data augmentation in the classical sense, where distortions are simply used to enlarge the data set. Augmentation has also been used in contrastive representation learning (e.g. Bai et al., 2021; Chen et al., 2020; He et al., 2020), but we are concerned with scenario in which the model learns to perceive distorted and original data as belonging to the same distribution. In this case, it was traditionally believed that a good augmentation should have minimal distribution shift.

However, increasingly many approaches propose heavy distortions (e.g. Yun et al., 2019; Summers & Dinneen, 2019), questioning the traditional view.

Most recently, it has been argued that it is the *degree of shift as perceived by the model* that determines augmentation quality. Gontijo-Lopes et al. (2020) propose to measure this shift by the difference between the performance of the model when presented with original test data and augmented test data. They refer to this gap as “affinity”. We show that *the magnitude of the distribution shift does not determine augmentation quality* of MSDA. We start with the perceptual gap of training with MSDA, as proposed in Gontijo-Lopes et al. (2020). Subsequently, we address the gap in the wider sense, as is often sought in prior art. We first argue that high affinity and high diversity are not necessarily desirable. Indeed, we find FMix, a better performing train-time augmentation in terms of impact on generalisation performance, to have both lower affinity and lower diversity than MixUp (Table 6). For diversity, we compute the cross-entropy loss where the label is taken to be that of the majority class. Similar results are obtained with the MixUp loss, where a weighted average of the true labels is taken.

While intuitively for a high level of affinity, high diversity could correspond to better methods, the converse does not hold. We argue this is because affinity is rather an analysis of the learnt representations of the reference model and cannot give an insight into the quality of the augmentation or its effect on learning. The limitations of affinity are intimately linked to those of CutOcclusion. We have seen in Section 3 that the bias of the basic model is present not only when obstructing an image with a uniform patch, but also when mask-mixing. As such, an augmentation will have a lower affinity if it introduces artefacts that could otherwise lead to learning better representations when used in the training process. We believe this issue extends to other approaches that aim to motivate the success of MSDA through reduced distribution shift. Henceforth, we focus on bringing further supporting evidence that *the importance lies in the invariance introduced by the shift and its interaction with the given problem rather than its magnitude*.

4.1. If it is not the Magnitude that Matters, is it the Direction?

We propose the study of introduced bias as a more informative research direction. We use the term “bias” to refer to a drift in the learnt representations introduced by the change in the training procedure. In this section we perform inter-dataset augmentation and show that by distorting the data distribution by a similar *magnitude*, we obtain two opposing results. This suggests that it is the *direction* of the introduced bias that is important for understanding the impact of augmentation.

Table 8. Accuracy on CIFAR-10 (left) and CIFAR-100 (right) upon mixing with samples from a different data set. The baseline is the accuracy when training with a single data set using the reformulated objective. CIFAR-110 refers to mixing with CIFAR-100 when training on the CIFAR-10 problem and vice-versa. Mixing CIFAR-10 images with CIFAR-100 is beneficial for the former (increases the average performance by 0.51%) and detrimental for the latter (decreases the average performance by 2.32%).

	MixUp	FMix	CutMix	MixUp	FMix	CutMix
baseline	94.18 \pm 0.34	94.36 \pm 0.28	94.67 \pm 0.20	74.68 \pm 0.37	75.75 \pm 0.31	74.19 \pm 0.50
CIFAR-110	94.70 \pm 0.27	94.80 \pm 0.32	94.66 \pm 0.12	72.36 \pm 1.04	74.80 \pm 0.55	74.47 \pm 0.39
Fashion	92.28 \pm 0.28	95.03 \pm 0.10	94.61 \pm 0.19	66.40 \pm 1.86	74.46 \pm 0.57	74.06 \pm 0.28

We use the reformulated objective setting (Huszár, 2017; Harris et al., 2020), where targets are not mixed and the mixing coefficient is drawn from an imbalanced Beta distribution. This allows us to apply MSDA between data sets. Thus, for training a model on a data set, we use an additional one whose targets will be ignored. As an example, a model that is learning to predict CIFAR-10 images will be trained on a combination of CIFAR-10 and CIFAR-100 images, with the target of the former. This scenario breaks the added correlation between training examples. Table 8 contains the results of this experiment, showing that an accuracy similar to or better than that of regular MSDA can be obtained by performing inter-dataset MSDA. This invalidates previous arguments that the power of MixUp comes from causing the model to act linearly between samples (Zhang et al., 2018), and calls for a broader explanation of its success beyond that of pushing the examples towards their mean (Carratino et al., 2020). Another observation is that for FMix and MixUp, introducing elements from CIFAR-100 when training models on the CIFAR-10 problem does not harm the learning process. The reciprocal, however, does not hold. Hence, the “distribution shift” is more intimately linked to the problem at hand and aiming to characterise an augmentation based on the distance from the original distribution is a limiting approach, especially when the distance is measured as perceived by a reference model.

We believe an explanation is that the artefacts created when putting together images from CIFAR-10 with those of CIFAR-100 could introduce information that makes the separation of the 10 classes easier. However, if the information happens to interfere with a feature that is important for separating the CIFAR-100 categories, the performance could degrade on this data set. This single experiment is not sufficient to draw any general conclusions. However, it does show that shifting two distributions by a similar amount can have different effects on the model performance. Thus, *the direction of the introduced bias could be more important than its magnitude*. While some level of data similarity has to be preserved when performing MSDA, it is far from being the objective of such data-distorting approaches.

We have highlighted that the shift in learnt representations can lead to better models and simply quantifying the distri-

bution shift can be misleading. An open question remains: How can we better capture the bias that is introduced and measure its quality? We believe understanding how a relatively small change in the data distribution impacts learnt representations could lead the way to characterising the relationship between data and model generalisation.

5. Conclusions

Studying the side-effects of data distortion is crucial for building fair evaluation methods, better training procedures and even understanding image data as a whole. In this paper we show that the artificial information introduced by visual distortion must be accounted for when evaluating models, and it should be studied and exploited during training. A limitation of previous studies that aim to explain the success of MSDA is the focus on trying to argue similarity with original data, rather than explaining the bias that is introduced. Correctly interpreting the bias is important not only for making the models trustable, but also for injecting more informed prior knowledge in future applications. Beyond their practical benefits, we believe MSDAs have the potential to help characterise the interplay between data and learnt representations. We hope this paper encourages better practice when dealing with all forms of data distortions.

Acknowledgements

A.M. is funded by the EPSRC Doctoral Training Partnership (EP/N509747/1). The authors acknowledge the use of the IRIDIS High-Performance Computing Facility, the ECS Alpha Cluster, and associated support services at the University of Southampton in the completion of this work. The authors would like to thank Jonathon Hare for the support and helpful discussions.

References

Alaiz-Rodríguez, R. and Japkowicz, N. Assessing the impact of changing environments on classifier performance. In *Conference of the Canadian Society for Computational Studies of Intelligence*, pp. 13–24. Springer, 2008.

- Bai, Y., Yang, Y., Zhang, W., and Mei, T. Directional self-supervised learning for risky image augmentations. *arXiv preprint arXiv:2110.13555*, 2021.
- Brendel, W. and Bethge, M. Approximating CNNs with bag-of-local-features models works surprisingly well on imagenet. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SkfMWhAqYQ>.
- Carlucci, F. M., D’Innocente, A., Bucci, S., Caputo, B., and Tommasi, T. Domain generalization by solving jigsaw puzzles. In *CVPR*, 2019.
- Carratino, L., Cissé, M., Jenatton, R., and Vert, J.-P. On mixup regularization. *arXiv preprint arXiv:2006.06049*, 2020.
- Chapelle, O., Weston, J., Bottou, L., and Vapnik, V. Vicinal risk minimization. In *Advances in neural information processing systems*, pp. 416–422, 2001.
- 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, 2020.
- Chun, S., Oh, S. J., Yun, S., Han, D., Choe, J., and Yoo, Y. An empirical evaluation on robustness and uncertainty of regularization methods. *arXiv preprint arXiv:2003.03879*, 2020.
- Cieslak, D. A. and Chawla, N. V. A framework for monitoring classifiers’ performance: when and why failure occurs? *Knowledge and Information Systems*, 18(1): 83–108, 2009.
- DeVries, T. and Taylor, G. W. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017.
- Engstrom, L., Tran, B., Tsipras, D., Schmidt, L., and Madry, A. Exploring the landscape of spatial robustness. In *International Conference on Machine Learning*, pp. 1802–1811. PMLR, 2019.
- Fawzi, A. and Frossard, P. Measuring the effect of nuisance variables on classifiers. In Richard C. Wilson, E. R. H. and Smith, W. A. P. (eds.), *Proceedings of the British Machine Vision Conference (BMVC)*, pp. 137.1–137.12. BMVA Press, September 2016. ISBN 1-901725-59-6. doi: 10.5244/C.30.137. URL <https://dx.doi.org/10.5244/C.30.137>.
- Feldman, V. and Zhang, C. What neural networks memorize and why: Discovering the long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33, 2020.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., and Brendel, W. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bygh9j09KX>.
- Gontijo-Lopes, R., Smullin, S. J., Cubuk, E. D., and Dyer, E. Affinity and diversity: Quantifying mechanisms of data augmentation. *arXiv preprint arXiv:2002.08973*, 2020.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Harris, E., Marcu, A., Painter, M., Niranjan, M., Prügell-Bennett, A., and Hare, J. Understanding and enhancing mixed sample data augmentation. *arXiv preprint arXiv:2002.12047*, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.
- Hooker, S., Erhan, D., Kindermans, P.-J., and Kim, B. A benchmark for interpretability methods in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 9737–9748, 2019.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Huszár, F. mixup: Data-dependent data augmentation, 2017. URL <http://www.inference.vc/mixup-data-dependent-data-augmentation/>.
- Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Al-sallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., et al. Captum: A unified and generic model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896*, 2020.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Master’s thesis, University of Toronto, 2009.
- Lassance, C., Béthune, L., Bontonou, M., Hamidouche, M., and Gripon, V. Ranking deep learning generalization using label variation in latent geometry graphs. *arXiv preprint arXiv:2011.12737*, 2020.

- Luo, T., Cai, T., Zhang, M., Chen, S., He, D., and Wang, L. Defective convolutional layers learn robust CNNs. *arXiv preprint arXiv:1911.08432*, 2019.
- Miller, G. A. Wordnet: a lexical database for english. *Communications of the ACM*, 38:39–41, 1995.
- Mummadi, C. K., Subramaniam, R., Hutmacher, R., Vitay, J., Fischer, V., and Metzen, J. H. Does enhanced shape bias improve neural network robustness to common corruptions? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=yUXUNaj2Sl>.
- Natekar, P. and Sharma, M. Representation based complexity measures for predicting generalization in deep learning. *arXiv preprint arXiv:2012.02775*, 2020.
- Rajaei, K., Mohsenzadeh, Y., Ebrahimpour, R., and Khaligh-Razavi, S.-M. Beyond core object recognition: Recurrent processes account for object recognition under occlusion. *PLoS computational biology*, 15(5):e1007001, 2019.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Schiff, Y., Quanz, B., Das, P., and Chen, P.-Y. Gi and pal scores: Deep neural network generalization statistics. *arXiv preprint arXiv:2104.03469*, 2021.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.
- Shi, B., Zhang, D., Dai, Q., Zhu, Z., Mu, Y., and Wang, J. Informative dropout for robust representation learning: A shape-bias perspective. In *International Conference on Machine Learning*, pp. 8828–8839. PMLR, 2020.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- Stanford. Tiny imagenet visual recognition challenge, 2015. URL <https://tiny-imagenet.herokuapp.com/>.
- Summers, C. and Dinneen, M. J. Improved mixed-example data augmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1262–1270. IEEE, 2019.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Tang, H., Schrimpf, M., Lotter, W., Moerman, C., Paredes, A., Caro, J. O., Hardesty, W., Cox, D., and Kreiman, G. Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences*, 115: 8835–8840, 2018.
- Vapnik, V. *The nature of statistical learning theory*. Springer science & business media, 1999.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Xie, L. Hardhat, 2019. URL <https://doi.org/10.7910/DVN/7CBGOS>.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 6023–6032, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=r1Ddp1-Rb>.
- Zhang, T. and Zhu, Z. Interpreting adversarially trained convolutional neural networks. In *International Conference on Machine Learning*, pp. 7502–7511. PMLR, 2019.
- Zhong, Z., Zheng, L., Kang, G., Li, S., and Yang, Y. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 13001–13008, 2020.

A. Experimental Details

Throughout the paper, we use PreAct-ResNet18 (He et al., 2016) models, trained for 200 epochs with a batch size of 128. For the MSDA parameters we use the same values as Harris et al. (2020). All models are augmented with random crop and horizontal flip and are averaged across 5 runs. We optimise using SGD with 0.9 momentum, learning rate of 0.1 up until epoch 100 and 0.001 for the rest of the training. This is due to an incompatibility with newer versions of the PyTorch library of the official implementation of Harris et al. (2020), which we use as a starting point for model training. However, the difference in learning rate schedule between our work and prior art does not affect our findings since we are not introducing a new method to be applied at training time. In our case, it is sufficient to show that the bias exists in at least one configuration. For the analysis we also used adapted code from (Carlucci et al., 2019) for patch-shuffling. The models were trained on either one of the following: Titan X Pascal, GeForce GTX 1080 Ti or Tesla V100. For the analyses, a GeForce GTX 1050 was also used. The average training time was less than two hours, with the exception of model trained on Tiny-ImageNet, which took around 10 hours to run.

Training Models

The code for model training is largely based on the open-source official implementation of FMix, which also includes those of MixUp, CutOut, and CutMix. For the experiment where we use the reformulated objective to combine data sets, instead of mixing with a permutation of the batch, as it is done in the original implementation of the mixed-augmentations, we now draw a batch from the desired data set. To ensure a fair comparison, for the basic we also perform inter-batch mixing.

Evaluating Robustness

For the CutOcclusion measurement, we modify open-source code to restrict the occluding patch to lie within the margins of the image to be occluded. This is to ensure that the mixing factor λ matches the true proportion of the occlusion. For iOcclusion, the implementation of Grad-CAM is again adapted from publicly available code. With both methods, we evaluate 5 instances of the same model and average over the results obtained.

The added computation time of iOcclusion over the regular CutOcclusion for a fixed occlusion fraction is that of performing Grad-CAM on train and test data, as well as evaluating on the latter. With a batch size of 128, this takes under half an hour.

B. Analysis of Wrong Predictions

B.1. Alternative Index

Table 9 gives the worst-case DI index where we replace c_{max} in Equation 1 by the maximum increase across the runs. As per the original formulation, we note that the masking methods lead to models which are less sensitive to the artefacts resulted after patch-shuffling.

Table 9. Alternative DI index for PreAct-ResNet18 on grid-shuffled images for four different types of models. Again, a bias can be noted for all considered data sets.

	basic	MixUp	FMix	CutMix
CIFAR-10	3.52 \pm 0.56	3.31 \pm 0.82	0.76 \pm 0.16	0.43 \pm 0.13
CIFAR-100	1.40 \pm 0.38	1.09 \pm 0.29	0.38 \pm 0.21	0.16 \pm 0.08
Fashion	1.56 \pm 0.39	3.57 \pm 1.35	1.65 \pm 0.35	0.82 \pm 0.13
Tiny	3.01 \pm 0.48	2.24 \pm 0.30	2.34 \pm 1.86	11.45 \pm 10.54
ImageNet	0.82	1.49	0.58	—

B.2. DI Confidence Evaluation and Gaussian Noise

Given the high standard deviation of our index, as a sanity check we want to determine what the chances are that the DI we measured is a result of chance. For this, we compute the DI index where each misclassified example gets randomly assigned to one of the classes, repeating this experiment 10^5 times. We obtain an average DI of 0.10 with a standard deviation of 4×10^{-5} . Thus, the result we obtain is not by chance and the index captures a real phenomenon.

To make our baseline more challenging, we then randomly assign **all** misclassified examples to one of the classes and compute the index across 5 iterations. This is so as to simulate the case where each specific model instance consistently associates artefacts with a certain class, but the class is not necessarily consistent across runs. We perform this experiment 10^5 times and obtain a DI of $0.10_{\pm 0.30}$.

Lastly, we compute the DI index for distorting with Gaussian noise. We distort the images after the image normalisation step, at which point the pixel values are between 0 and 1. For each image, we uniformly sample the standard deviation of the Gaussian noise from the interval $[0, 0.1]$, while the mean we set to 0. For the basic model, we obtain a DI value of $0.09_{\pm 0.02}$. This is significantly lower than the index we observe in the case of patch-shuffling or rectangular occlusion.

B.3. Varying the Grid Size

Table 10 gives the results obtained when varying the number of image tiles to be randomly rearranged. We observe that data interference appears for different grid sizes.

Table 10. DI index for alternative grid sizes. Gaps can be identified for various grid sizes, with more pronounced differences for finer-grained grids.

		basic	MixUp	FMix	CutMix
CIFAR-10	2×2	$0.61_{\pm 0.24}$	$0.56_{\pm 0.33}$	$0.19_{\pm 0.14}$	$0.12_{\pm 0.06}$
	8×8	$6.41_{\pm 0.55}$	$6.95_{\pm 1.96}$	$2.75_{\pm 1.46}$	$1.41_{\pm 1.15}$
CIFAR-100	2×2	$1.03_{\pm 0.29}$	$0.46_{\pm 0.14}$	$0.21_{\pm 0.14}$	$0.12_{\pm 0.07}$
	8×8	$9.16_{\pm 6.15}$	$3.10_{\pm 4.59}$	$1.62_{\pm 0.89}$	$0.65_{\pm 0.50}$
Tiny ImageNet	8×8	$5.76_{\pm 6.61}$	$5.73_{\pm 3.82}$	$2.49_{\pm 1.38}$	$0.60_{\pm 0.69}$
	16×16	$44.01_{\pm 36.47}$	$14.06_{\pm 14.63}$	$11.94_{\pm 17.79}$	$1.86_{\pm 1.98}$
ImageNet	16×16	0.72	1.38	0.55	–
	64×64	4.89	41.16	12.77	–

B.4. Patch-shuffling

We look at the classes which have the highest increase in incorrect predictions and note that their shapes are characterised by strong horizontal and vertical edges. For example, on CIFAR-100, varying the grid size between 2×2 , 4×4 and 8×8 gives “Lamp”, “Bus” and “Table” as dominant c_{max} classes, while the model trained on Fashion MNIST with the standard procedure tends to predict grid-shuffled images as “Bag”. Figure 7 shows that on ImageNet, the basic model tends to wrongly identify the patch-shuffled images as belonging to class “Envelope”.

B.5. Increase in Wrong Predictions for the Shape-biased Models

We attempted to train a model on the GST data set using the shape labels. Although we have identified a bias, upon a closer inspection we observed that all model instances were heavily overfitting. This is because the GST data set is meant as an evaluation data set, only containing 10 unique shape examples for each class. Due to time constraints, we could not further explore settings that would give sufficiently well-performing models on which the bias evaluation could be trusted. Nonetheless, this remains an interesting future direction to explore.

Can we propose a fairer shape bias evaluation starting from patch-shuffling?

The shape-texture issue is far more complicated and subtle than occlusion. We would argue that the definition of shape and texture on which the patch-shuffling method is based is flawed as it locally maintains both the shape and the texture information. Although it is mostly the texture that is preserved in the patches, edge information is often preserved as well (e.g. a cat’s pointy ears might be fully preserved in a patch). Thus, great care must be taken to make sure that a correct and reliable evaluation of the bias can be constructed with patch-shuffling.

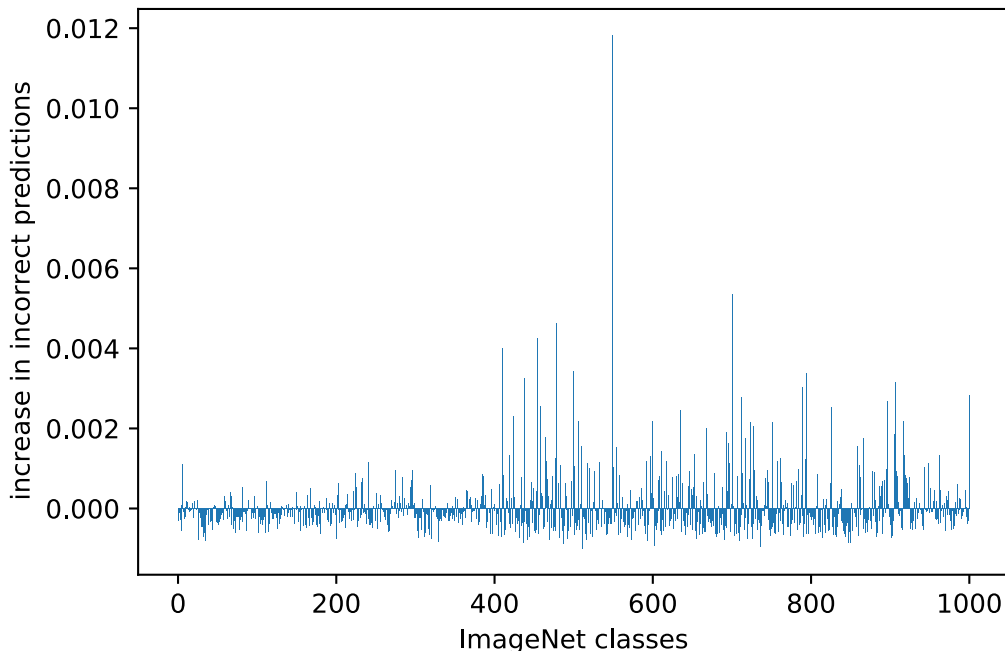


Figure 7. Difference between the number of times a class was wrongly predicted when presented with regular ImageNet samples and patch-shuffled data.

B.6. CutOcclusion

In this section we experiment with alternative masking methods when computing CutOcclusion. We note that the bias exists when occluding with patches taken from images belonging to different data sets (Table 11). Figure 9 gives a visual account of the results obtained for CIFAR-10 when mix-patching, while the results obtained when occluding with a uniform patch are given in Figure 8. Note that when mix-patching, for Fashion MNIST we use MNIST, for Tiny ImageNet we use ImageNet, while for CIFAR-10 we mix with CIFAR-100 and vice versa. Since ImageNet images are significantly larger than those of the other data sets, mixing would imply padding large areas, which would give results very similar to uniform patching. We also experiment with VGG models, where on CIFAR-10 the basic has a DI index of $0.80_{\pm 0.40}$ compared to $0.18_{\pm 0.11}$ of MixUp.

Table 11. DI index for occluding with images from another data set.

	basic	MixUp	FMix	CutMix
CIFAR-10	$0.18_{\pm 0.05}$	$0.39_{\pm 0.15}$	$0.12_{\pm 0.11}$	$0.08_{\pm 0.06}$
CIFAR-100	$0.48_{\pm 0.09}$	$0.61_{\pm 0.27}$	$0.90_{\pm 0.15}$	$1.25_{\pm 0.25}$
Fashion	$3.40_{\pm 0.29}$	$3.06_{\pm 1.07}$	$1.81_{\pm 0.55}$	$2.61_{\pm 0.80}$
Tiny	$0.25_{\pm 0.12}$	$0.17_{\pm 0.04}$	$0.06_{\pm 0.03}$	$0.12_{\pm 0.04}$

On the Effects of Artificial Data Modification

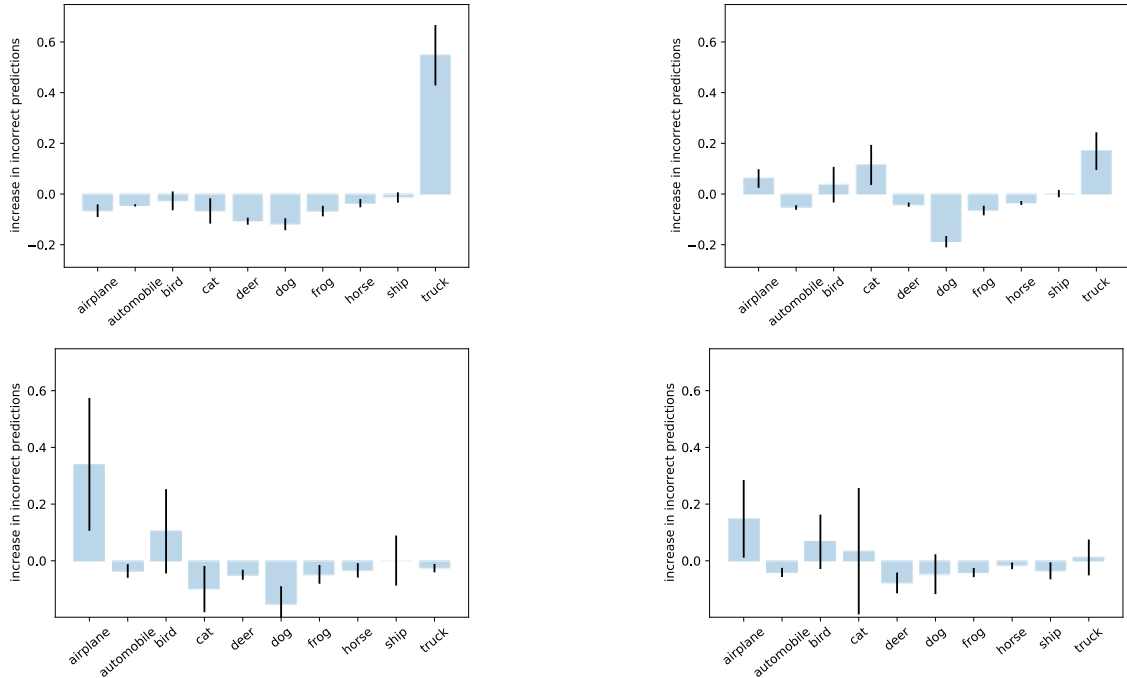


Figure 8. Difference between wrongly predicted classes when testing on original data versus CutOut images. Left to right, top to bottom: basic, MixUp, CutMix, FMix. The evaluated models from left to right, top to bottom are trained on CIFAR-10 with: no mixed-data augmentation (basic), MixUp, CutMix, and FMix.

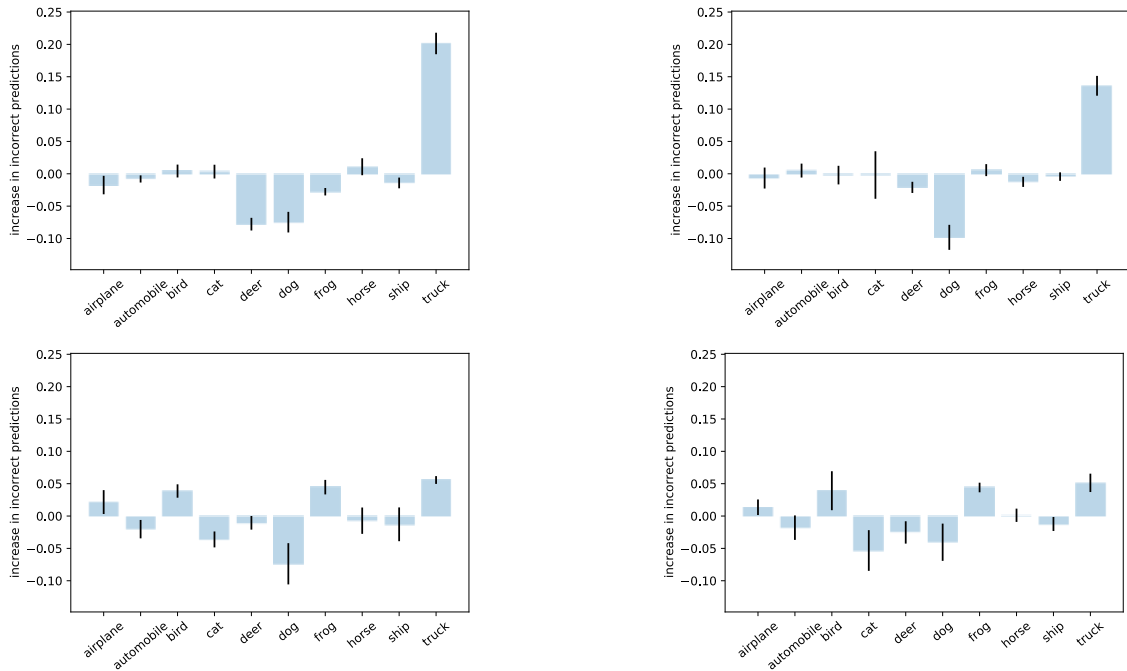


Figure 9. Difference between wrongly predicted classes when testing on original data versus CutMix images. Left to right, top to bottom: basic, MixUp, CutMix, FMix. The evaluated models from left to right, top to bottom are trained on CIFAR-10 with: no mixed-data augmentation (basic), MixUp, CutMix, and FMix.

We then use masks sampled from Fourier space (Table 12) and note that even for these irregularly shaped distortions, we can identify a gap in most cases. The only exception is in the case of Fashion MNIST. It must be stressed that although all the models we experimented with presented Data Interference for this problem, this does not exclude the possibility of constructing a different model that is insensitive to this distortion. For example, we identify a gap for this problem when mix-masking (DI index of 4.09 ± 1.74 for the basic model as opposed to 1.87 ± 0.27 for a model trained on images that were masked out using FMix-like masks). Thus, when occluding with a particular shape we implicitly disfavour models in which learnt representations are related to the features introduced by that shape.

Table 12. DI index for patching using masks sampled from Fourier space.

	basic	MixUp	FMix	CutMix
CIFAR-10	2.08 ± 1.13	1.79 ± 1.09	1.32 ± 0.99	4.21 ± 1.23
CIFAR-100	4.06 ± 01.47	3.11 ± 02.29	9.90 ± 14.32	2.89 ± 05.36
Fashion	49.55 ± 20.4	40.69 ± 21.6	27.87 ± 17.5	61.04 ± 17.9
Tiny	4.37 ± 0.85	6.95 ± 1.84	3.60 ± 1.73	5.92 ± 4.38
ImageNet	3.27	2.24	6.08	—

Figure 10 also gives the results for CutOcclusion and iOcclusion for training with 3 random masks sampled from Fourier space.

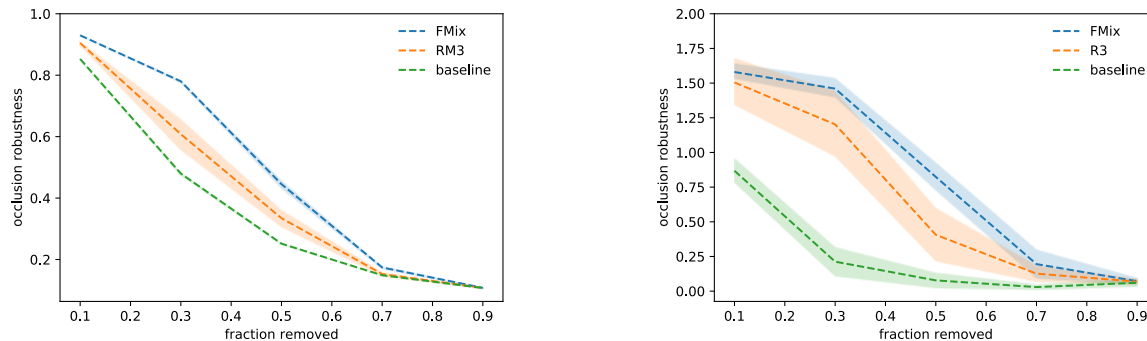


Figure 10. CutOcclusion (left) and iOcclusion (right). Note that there is a difference in scale and the two should not be directly compared. We are rather interested in how the methods situate the different augmentation with respect to each other. It is important to notice that when measuring the robustness with CutOcclusion, RM3 appears significantly less robust than FMix due to its sensitivity to patching with rectangles. On the other hand, iOcclusion highlights the robustness specific to FMix.

B.7. BagNet Shape and Texture Accuracy

We evaluate on the GST data set BagNet9 models trained on Tiny ImageNet and present the results in Table 13. Despite the basic model displaying a bias towards predicting one of the classes when presented with patch-shuffled images (see Figure 11), once again it does not have a lower texture or higher shape bias than the masked-based augmentations.

Table 13. Shape and texture accuracy of BagNet9 models on the GST data set.

	Shape	Texture
basic	11.29 ± 0.15	18.90 ± 0.66
MixUp	11.04 ± 0.29	12.56 ± 1.26
FMix	11.06 ± 0.48	17.47 ± 1.74
CutMix	10.76 ± 0.27	20.28 ± 0.88

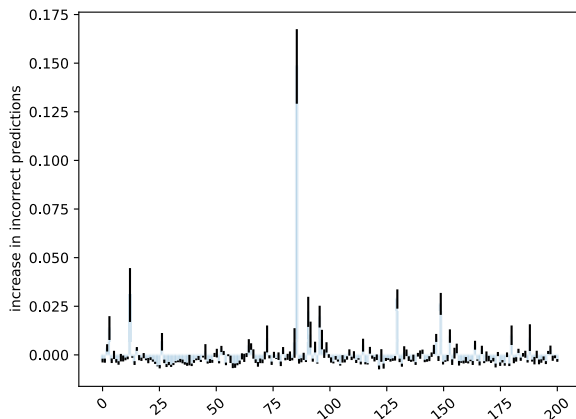


Figure 11. Difference between wrongly predicted classes when testing on original Tiny ImageNet data versus CutOut images.

B.8. Data Interference across Architectures

To verify if a shape bias evaluation based on patch-shuffling would give unfair results when comparing across architectures, we compute the *DI* index for a number of models trained with the basic approach (without mixed data augmentation) on CIFAR-10. The *DI* index of models such as WideResNet and ResNet is high (0.95 ± 0.20 and 1.27 ± 0.39 respectively), while for PyramidNet, BagNet17 and BagNet9 it is small (0.26 , 0.53 ± 0.13 , 0.54 ± 0.16). DenseNet (0.66 ± 0.48) and VGG (0.60 ± 0.15) have comparable *DI* indices.

Thus, we find that intensity with which distortions interfere with learnt representations is different for different architectures. Comparing robustness to occlusion using CutOcclusion would give biased results when comparing architectures trained under the same conditions.

C. Further Results on iOcclusion Experiments

C.1. Alternative CutOcclusion

Table 14 gives the *DI* index when forcing the occluding patch to lie within image boundaries for patch sizes sampled uniformly from $[0.1, 1]$. We also experimented with sampling occluder sizes from a Beta(2,1) distribution and obtained similar results. Naturally, the *DI* index is more pronounced and we once again notice a gap for all data sets.

Table 14. *DI* index for sampling occluder size from a uniform distribution when the patch is restricted to lying within image boundaries and the size is sampled from $[0.1, 1]$ uniformly.

	basic	MixUp	FMix	CutMix
CIFAR-10	5.71 ± 0.18	0.51 ± 0.41	0.77 ± 0.75	2.88 ± 1.31
CIFAR-100	27.81 ± 5.23	4.40 ± 3.25	4.12 ± 3.12	5.13 ± 6.68
Fashion	0.93 ± 0.98	3.46 ± 0.58	0.99 ± 0.99	0.69 ± 0.93
Tiny	13.72 ± 2.46	5.47 ± 1.48	4.10 ± 2.69	6.31 ± 6.76
ImageNet	0.25	0.48	0.14	—

C.2. Occluding with Images from Another Data Set

Since CutOcclusion does not account for the bias introduced by the occluding method, it is expected that changing the patch to a non-uniform one would greatly affect the results. For CIFAR-10 models, Figure 12 presents the results of occluding with uniform patches versus patches extracted from CIFAR-100 images. iOcclusion better rules out the specifics of the occluding patch, its uniform version giving similar results to the non-uniform one, whereas CutOcclusion pushes everything

together.

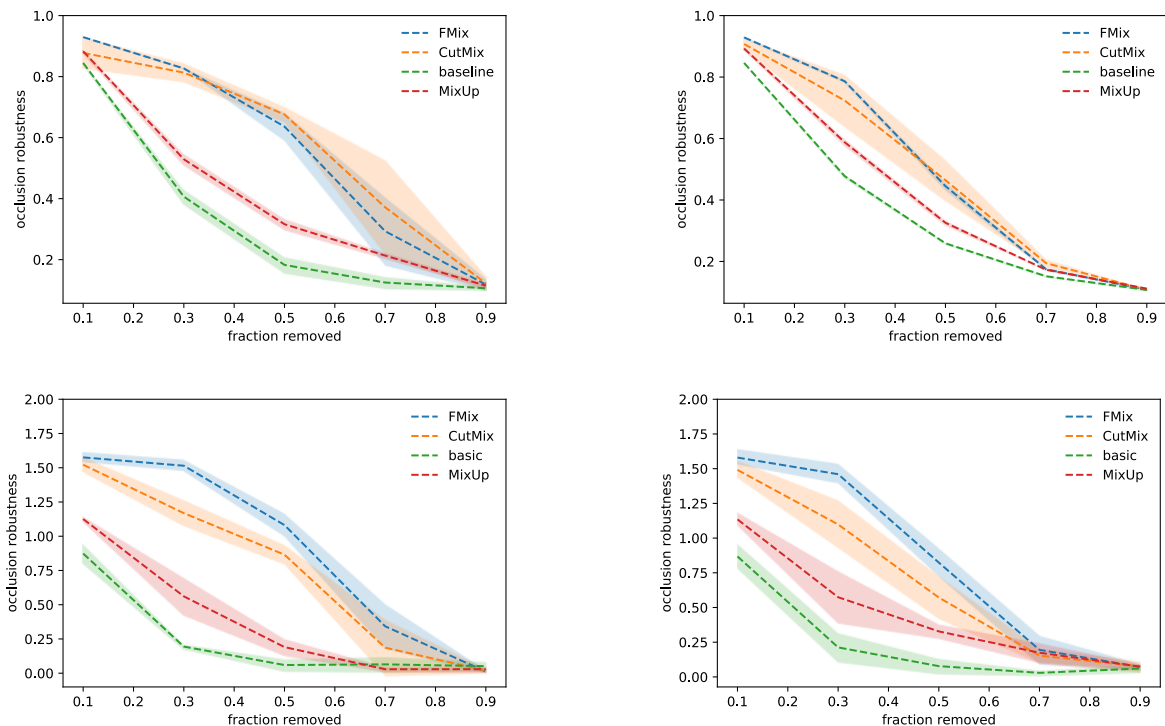


Figure 12. Comparison of metric sensitivity to textured occlusion. CutOcclusion (top) versus iOcclusion (bottom) when occluding with uniform patches (left hand-side) and non-uniform patches (right hand-side). Uniform occlusion refers to superimposing uniform patches over CIFAR-10 images, while nonuniform refers to superimposing part of CIFAR-100 samples. Nonuniform CutOcclusion provides significantly different results to its regular counterpart, whereas iOcclusion provides more consistent results.

C.3. Randomising Labels

To assess the sensitivity of CutOcclusion and iOcclusion to the overall performance of the model, we also experiment with randomising all the labels of the CIFAR-10 data set. When evaluated on the unaugmented training data, all the basic models achieve 100% accuracy, while the FMix models reach 99.99 ± 0.01 . Since all labels are corrupted, the accuracy on the test set before and after occlusion is no greater than random. However, the robustness of the augmentation-trained model can be seen on the training data, as captured by our metric (see Figure 13). On the other hand, CutOcclusion makes no distinction between learning with regular and augmented data (Table 15). Despite being such a peculiar case, it shows the comprehensiveness gained by accounting for the degradation on test data in relation to that on train.

Table 15. Robustness to occluding with patches covering 50% of each image. The models are trained with and without masking augmentation on data with randomised labels. CutOcclusion makes no difference between regular and augmented training.

	basic random	FMix random	FMix clean
CutOcclusion	10.24 ± 0.27	9.78 ± 0.18	63.63 ± 4.54
iOcclusion	14.63 ± 1.12	47.94 ± 19.84	82.36 ± 10.06

C.4. iOcclusion Numerator and Denominator

In this section we give the values of the composing elements of iOcclusion for models trained on CIFAR-10. While the numerator depends on the size of the occluder, the denominator is a fixed term: the generalisation gap. Results are presented in Table 16.

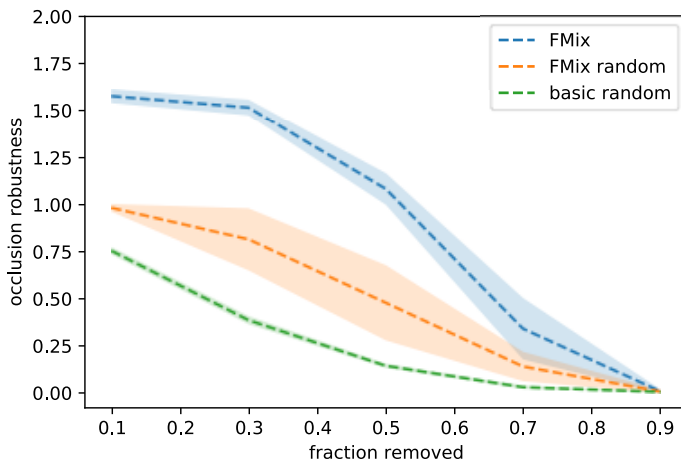


Figure 13. iOcclusion results for training with clean and corrupted labels for basic and FMix augmentation.

Table 16. Value of the numerator (accuracy gap when occluding a proportion of the image) and denominator (generalisation gap) for models trained on CIFAR-10.

proportion	numerator					denominator
	10%	30%	50%	70%	90%	
basic	4.66 ± 0.38	1.06 ± 0.10	0.32 ± 0.24	0.35 ± 0.30	0.29 ± 0.19	5.47 ± 0.10
CutMix	6.96 ± 0.24	5.49 ± 0.44	3.78 ± 0.30	0.86 ± 0.94	0.09 ± 0.07	4.60 ± 0.12
FMix	7.04 ± 0.16	7.32 ± 0.20	5.07 ± 0.39	1.51 ± 0.71	0.05 ± 0.04	4.58 ± 0.15
MixUp	5.43 ± 0.09	2.68 ± 0.65	0.86 ± 0.27	0.13 ± 0.08	0.14 ± 0.10	4.70 ± 0.12

C.5. Approximating iOcclusion

As alternative methods for computing iOcclusion we experiment both with masks sampled from Fourier space and randomly positioned square patches. Although using this type of random masking methods for computing \mathcal{D}_{train}^p and \mathcal{D}_{test}^p in Equation (1) gives less precise results, it has the advantage of incurring less computation and can be used for rapid model analysis. For assessing a model across 5 runs for 6 different levels of occlusion, this method leads to a carbon footprint of 0.05 kgCO₂eq as opposed to 1.04 using Grad-CAM. In Figure 14 we present the results obtained with these alternatives. We expect both methods to provide overoptimistic results for small patches, while Fourier sampling is expected to give more truthful scores as the size of the patch increases. On the other hand, the contiguity of CutOut-based occlusion comes at the cost of not determining the robustness to multiple simultaneous occluders. This seems to play a role especially in the case of CutMix augmentation. Indeed, when superimposing a rectangular patch, it is difficult to differentiate CutMix from FMix-trained models. To confirm that this is caused by the granularity of the occluders and not the shape, we also experiment with occluding using multiple rectangular patches. We split the images in a 4×4 grid and occlude $i\%$ of the tiles, obtaining results that are more similar to those obtained when occluding with Fourier-sample patches. Thus, while significantly noisier, using randomly positioned occluders can provide an alternative for computing iOcclusion given that one takes into account the number of occluders.

C.5.1. REMOVING THE DOMINANT CLASS

We remove the 10th class from the CIFAR-10 data set and retrain on the remaining classes. We then evaluate the basic and MixUp models on data distorted with non-uniform patches. Just as in the other experiments, we obtain non-uniform patches by randomly sampling rectangular regions from the CIFAR-100 images. We once again find that the basic model show a bias towards one of the classes, namely the “ship” class (see Figure 15). The basic model has a DI index of 1.11 ± 0.28 , while CutMix 0.16 ± 0.04 . Thus, when measuring robustness using CutOcclusion, the basic model will be disadvantaged.

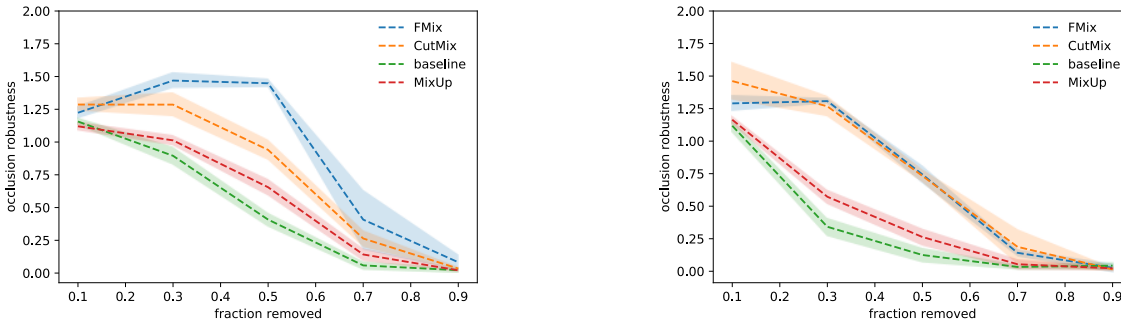


Figure 14. iOcclusion when using Fourier-sampled masks (left) and rectangular masks (right). Approximating iOcclusion with random masking can provide a first intuition, but is noisier than using a saliency method.

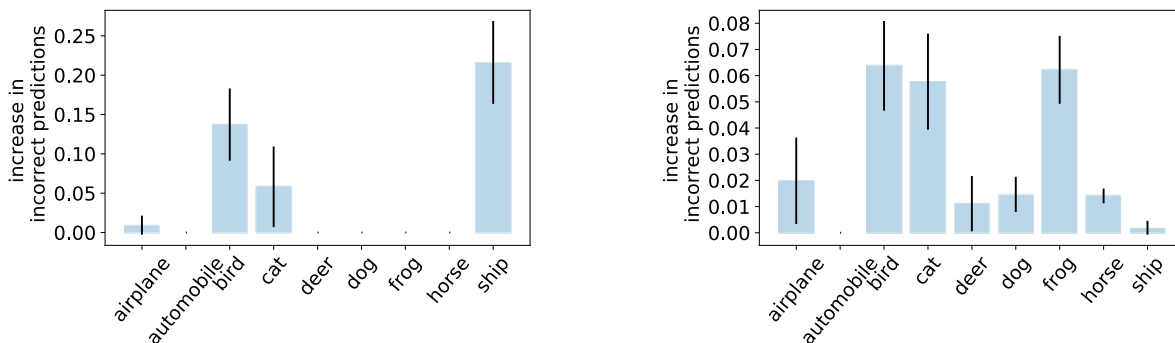


Figure 15. Increase in incorrect in incorrect predictions for the basic (left) and CutMix (right) models when removing the dominant class (“truck”) and occluding with non-uniform patches.

C.6. Results Obtained on a Data Set with Small Object Size

All the benchmark data sets we evaluate robustness on have large, centred subjects. To provide a more comprehensive evaluation, we included a data set where the subject differs in scale and location for a classification task. Since we were unable to find any publicly available data sets with these characteristics, we adapted Xie (2019)’s data set to a binary classification task to identify the presence or absence of hard hats in the image samples. It must be noted that in the case of binary classification the bias is less clearly identifiable. In the multiclass case, we are looking for a consistent association with a specific class across runs and distortion intensities. However in the binary case, any misclassification would naturally be assigned to the opposite class.

We evaluate the CutOcclusion and iOcclusion robustness levels for a level of occlusion of 30% and present results in Table 17. The CutOcclusion robustness is massively overlapping , while iOcclusion reflects some of the added robustness of FMix and CutMix.

Table 17. Robustness to occlusion as measured using CutOcclusion and iOcclusion for the hard hat data set when occluding 30% of the pixels in each image.

	basic	FMix	CutMix
CutOcclusion	67.95 \pm 10.07	71.02 \pm 8.49	69.59 \pm 11.95
iOcclusion	0.39 \pm 0.37	1.47 \pm 0.76	1.42 \pm 0.51