

---

# The Origins and Prevalence of Texture Bias in Convolutional Neural Networks

---

**Katherine L. Hermann**  
Stanford University  
hermannk@stanford.edu

**Ting Chen**  
Google Research, Toronto  
iamtingchen@google.com

**Simon Kornblith**  
Google Research, Toronto  
skornblith@google.com

## Abstract

Recent work has indicated that, unlike humans, ImageNet-trained CNNs tend to classify images by texture rather than by shape. How pervasive is this bias, and where does it come from? We find that, when trained on datasets of images with conflicting shape and texture, CNNs learn to classify by shape at least as easily as by texture. What factors, then, produce the texture bias in CNNs trained on ImageNet? Different unsupervised training objectives and different architectures have small but significant and largely independent effects on the level of texture bias. However, all objectives and architectures still lead to models that make texture-based classification decisions a majority of the time, even if shape information is decodable from their hidden representations. The effect of data augmentation is much larger. By taking less aggressive random crops at training time and applying simple, naturalistic augmentation (color distortion, noise, and blur), we train models that classify ambiguous images by shape a majority of the time, and outperform baselines on out-of-distribution test sets. Our results indicate that apparent differences in the way humans and ImageNet-trained CNNs process images may arise not primarily from differences in their internal workings, but from differences in the data that they see.

## 1 Introduction

Convolutional neural networks (CNNs) define state-of-the-art performance in many computer vision tasks, such as image classification (57), object detection (83; 40), and segmentation (40). Although their performance in several of these tasks approaches that of humans, recent findings show that CNNs differ in intriguing ways from human vision, indicating fundamental deficiencies in our understanding of these models (87; 70; 32; 23; 76; 37; 3; 44; 46; 1; 48). This paper focuses on one such result, namely that CNNs appear to make classifications based on superficial textural features (36; 4) rather than on the shape information preferentially used by humans (61; 60). Building on a long tradition of work in psychology and neuroscience documenting humans’ shape-based object classification, Geirhos et al. (36) compared humans to ImageNet-trained CNNs on a dataset of images with conflicting shape and texture information (e.g. an elephant-textured knife) and found that models tended to classify according to texture (e.g. “elephant”), and humans according to shape (e.g. “knife”). Following their terminology, we call a learner that prefers texture over shape *texture-biased* and one that prefers shape over texture *shape-biased*.

From the point of view of computer vision, texture bias is an important phenomenon for several reasons. First, it may be related to the vulnerability of CNNs to adversarial examples (87), which may exploit features that carry information about the class label but are undetectable to the human visual system (48). Second, a CNN preference for texture could indicate an inductive bias different than that of humans, making it difficult for models to learn human-relevant vision tasks in small-data regimes, and to generalize to different distributions than the distribution on which the model is trained (35).

In addition to these engineering considerations, texture bias raises important scientific questions. ImageNet-trained CNNs have emerged as the model of choice in neuroscience for modeling electrophysiological and neuroimaging data from primate visual cortex (101; 52; 15; 71; 7). Evidence that CNNs are preferentially driven by texture indicates a significant divergence from primate visual processing, in which shape bias is well documented (61; 60; 38). This mismatch raises an important puzzle for human-machine comparison studies. As we show in Section 7, even models specifically designed to match neural data exhibit a strong texture bias.

This paper explores the origins of texture bias in ImageNet-trained CNNs, looking at the effects of data augmentation, training procedure, model architecture, and task. In Section 4, we show that CNNs learn to classify the shape of ambiguous images at least as easily as they learn to classify the texture. So what makes ImageNet-trained CNNs classify images by texture when humans do not? While we find effects of all of the factors investigated, the most important factor is the data itself. Our contributions are as follows:

- We find that naturalistic data augmentation involving color distortion, noise, and blur substantially decreases texture bias, whereas random-crop augmentation increases texture bias. Combining these observations, we train models that classify ambiguous images by shape a majority of the time, without using the non-naturalistic style transfer augmentation of Geirhos et al. (36). These models also outperform baselines on out-of-distribution test sets that exemplify different notions of shape (ImageNet-Sketch (96) and Stylized ImageNet (36)).
- We investigate the texture bias of networks trained with different self-supervised learning objectives. While some objectives decrease texture bias relative to a supervised baseline, others increase it.
- We show that architectures that perform better on ImageNet generally exhibit lower texture bias, but neither architectures designed to match the human visual system nor models that replace convolution with self-attention have texture biases substantially different from ordinary CNNs.
- We separate the extent to which shape information is represented in an ImageNet-trained model from how much it contributes to the model’s classification decisions. We find that it is possible to extract more shape information from a CNN’s later layers than is reflected in the model’s classifications, and study how this information loss occurs as data flows through a network.

## 2 Related work

**Adversarial examples.** Adversarial examples are small perturbations that cause inputs to be misclassified (87; 10; 70). Adversarial perturbations are not entirely misaligned with human perception (29; 106), and reflect true features of the training distribution (48). State-of-the-art defenses on ImageNet use adversarial training (87; 41; 66; 100) or randomized smoothing (17), and images generated by optimizing class confidence under these models are more perceptually recognizable to humans (80; 30; 93; 51). Adversarial training can improve accuracy on out-of-distribution data (99). However, models that are robust to adversarial examples generated with respect to the  $\ell_p$  norm are not necessarily robust to other forms of imperceptible perturbations (98; 92).

**Data augmentation and robustness.** Data augmentation is widely used to improve generalization of CNNs on data drawn from the same distribution as the training data (62; 16; 8; 57; 22; 19; 64; 20; 103; 45). Recent work has shown that data augmentation can also improve out-of-distribution robustness (33; 102; 65; 45; 78). However, most of these studies investigate robustness to simple image corruptions, which may not be correlated with robustness to natural distribution shift (88). As with adversarial training (93; 73), augmentation strategies that improve robustness can degrade clean accuracy (65), and tradeoffs exist between robustness to different corruptions (102). Hosseini et al. (47) found that CNNs can learn to generalize to negative (reversed-brightness) images for a held-out class if the training data is augmented with negative images for the other classes.

**Sensitivity of CNNs to non-shape features.** Our work builds on recent evidence for bias for texture versus shape in neural networks based on classification behavior with ambiguous stimuli (36; 4). Other studies have shown that CNNs are sensitive to a wide range of other image manipulations that have little effect on human judgments (32; 5; 23; 76; 37; 3; 44; 46; 1; 6). Moreover, CNNs are relatively *insensitive* to manipulations such as grid scrambling that make images nearly unrecognizable to humans (11), and are far superior to humans at classifying ImageNet images where the foreground object has been removed (107).




Figure 1: Example items from the three datasets labeled according to shape (top) and texture (bottom). GST items reproduced with permission of (36; 91). ImageNet-C items from (44).

**Similarity of human and CNN perceptual biases.** Despite these differences, ImageNet-trained CNNs appear to share some perceptual biases and representational characteristics with humans. Previous studies found preferences for shape over color (77) and perceptual shape over physical shape (58). Euclidean distance in the representation space of CNN hidden layers correlates well with human perceptual similarity (50; 105), although, when used to *generate* perceptual distortions, simple biologically inspired models are a better match to human perception (9). CNN representations also provide an effective basis for modeling the activity of primate visual cortex (101; 52; 12), even though CNNs’ image-level confusions differ from those of humans (74).

### 3 Methods

**Datasets.** In addition to using ImageNet ILSVRC-2012 (79), our experiments used datasets consisting of colored  $224 \times 224$  images with independent sets of shape and texture labels. Instead of adopting a single parametric model (e.g. (43; 72; 28; 34)), we construe texture broadly to include natural textures (e.g. dog fur), small units repeating across space, and surface-level noise patterns. Each dataset (Figure 1) captures a possible interpretation of texture. See Appendix E.1.1 for additional discussion of the datasets and their limitations, as well as Figure C.1.

*Geirhos Style-Transfer (GST) dataset.* Introduced by (36), this dataset contains images generated using neural style transfer (34), which combines the content (shape) of one target natural image with the style (texture) of another. The dataset consists of 1200 images rendered from 16 shape classes, with 10 exemplars each, and 16 texture classes, with 3 exemplars each (91).

*Navon dataset.* Introduced by psychologist David Navon in the 1970s to study how people process global versus local visual information (69), Navon figures consist of a large letter (“shape”) rendered in small copies of another letter (“texture”). Unlike the GST stimuli, the primitives for shape and texture are identical apart from scale, allowing for a more direct comparison of the two feature types. We rendered each possible shape-texture combination ( $26 \times 26$  letters) at 5 positions, yielding a total of 3250 items after excluding items with matching shape and texture. Each image was rotated with an angle drawn from  $[-45, 45]$  degrees. In his experiments, Navon found that humans process the shape of these stimuli more rapidly than texture.

*ImageNet-C dataset.* ImageNet-C (44) consists of ImageNet images corrupted by different kinds of noise (e.g. shot noise, fog, motion blur). Here, we take noise type as “texture” and ImageNet class as “shape”. The original dataset contains 19 textures each at 5 levels, with 1000 ImageNet classes per level and 50 exemplars per class. To balance shape and texture, for each of 5 subsets of the data (dataset “versions”), we subsampled 19 shapes, yielding a total of 90,250 items per version.

**Shape bias evaluation.** Following (36) we define the *shape bias* of a model as the percentage of the time it classified images from the GST dataset according to shape, provided it classified either shape or texture correctly (see Appendix E.2 for additional details). We call a model *shape-biased* if its shape bias is  $> 50\%$ , and *texture-biased* if it is  $< 50\%$ . Throughout, *shape match* and *texture match* indicate the percentage of the time the model chose the image’s correct shape or texture label, respectively, as opposed to choosing some other label.

### 4 CNNs can learn shape as easily as texture

Do the texture-driven classifications documented by (36) come from the model, as an inherent property of the CNN architecture or learning process, or from the data, as an accidentally useful




Figure 2: **Task performance of AlexNet and ResNet-50 models as a function of data (first column) and time efficiency (second column), for each dataset (rows).** (**Data efficiency**). Performance is the maximum classification accuracy for shape (blue) or texture (orange) over training. All plots show the mean  $\pm$  SD across 5 splits. Dashed line indicates chance performance. (**Learning dynamics**). Accuracy over training time on 100% of the training data (GST: 716 items, Navon: 2875 items, ImageNet-C: 80750 items).

regularity that CNNs can exploit? We compared the learning dynamics of models trained to classify ambiguous non-ImageNet images according to either shape or texture: examining the number of training epochs and amount of training data required for models to perform well on these tasks is a way to determine whether shape or texture information is more easily exploited by CNNs. For each of the Geirhos-Style Transfer, Navon, and ImageNet-C datasets, we trained AlexNet (57) and ResNet-50 (42) models to classify images according to (a) shape, or (b) texture, using [5%, 10%, 20%, 30%, ..., 90%, 100%] of the training data. We compared validation accuracies for the two tasks.<sup>1</sup> For details, see Appendix E.1.3.

As shown in Figure 2, nothing prevents AlexNet and ResNet-50 from learning to classify ambiguous images by their shape. We should be cautious in comparing shape and texture performance on GST and ImageNet-C given the unknown in-principle recoverability of shape and texture from these datasets. However, best-possible error rates are known and equal for the synthetic Navon tasks, and we note that in that setting, both models learn to classify by shape at least as accurately as by texture, and, in the case of AlexNet, do so with greater efficiency in time and data. Overall, then, the source of

<sup>1</sup>Note that in the experiments detailed here, we trained models to classify ambiguous images, using either shape or texture labels as targets. In the other experiments in the paper, we trained models on ordinary ImageNet and evaluated shape bias on ambiguous images.

Table 1: **Random-crop augmentation biases models towards texture.** Characteristics of ImageNet-trained models with random-crop (Random) versus with center-crop (Center) preprocessing. For each model and metric, the preprocessing that achieved a higher value is boldfaced.

Model	Shape Bias		Shape Match		Texture Match		ImageNet Top-1 Acc.	
	Random	Center	Random	Center	Random	Center	Random	Center
AlexNet	28.2%	<b>37.5%</b>	16.4%	<b>19.3%</b>	<b>41.8%</b>	32.1%	<b>56.4%</b>	50.7%
VGG16	11.2%	<b>15.8%</b>	7.6%	<b>10.7%</b>	<b>60.1%</b>	57.1%	<b>71.8%</b>	62.5%
ResNet-50	19.5%	<b>28.4%</b>	11.7%	<b>16.3%</b>	<b>48.4%</b>	41.1%	<b>76.6%</b>	70.7%
Inception-ResNet v2	23.1%	<b>27.9%</b>	15.1%	<b>19.8%</b>	50.2%	<b>51.2%</b>	<b>80.3%</b>	77.3%

Table 2: **Color distortion, Gaussian blur, Gaussian noise, and Sobel filtering reduce texture bias.** ResNet-50 models were trained on ImageNet with random crops for 90 epochs. Augmentations were applied with 50% probability. Bolding indicates significantly greater shape bias than the baseline model ( $p < 0.05$ , permutation test).

Augmentation	Shape Bias	Shape Match	Texture Match	ImageNet Top-1 Acc.
Baseline	19.5%	11.7%	48.4%	76.6%
Rotate 90°, 180°, 270°	19.4%	10.8%	45.1%	75.7%
Cutout	<b>21.4%</b>	12.3%	45.2%	76.9%
Sobel filtering	<b>24.8%</b>	12.8%	38.9%	71.2%
Gaussian blur	<b>25.2%</b>	14.1%	41.7%	75.8%
Color distort.	<b>25.8%</b>	15.3%	44.2%	76.9%
Gaussian noise	<b>30.7%</b>	17.2%	38.8%	75.6%

texture bias seems to lie in training on ImageNet rather than on CNNs’ inductive biases; the remainder of the paper attempts to determine which aspects of ImageNet training are most responsible.

## 5 The role of data augmentation in texture bias

**Random-crop data augmentation increases texture bias.** In their experiments showing a texture bias in ImageNet-trained CNNs, Geirhos et al. (36) followed the standard practice of doing random-crop augmentation: crop shapes are sampled as random proportions of the original image size from  $[0.08, 1.0]$  with aspect ratio sampled from  $[0.75, 1.33]$ , and then resized to  $224 \times 224$ px (90). We hypothesized that random-crop augmentation might remove global shape information from the image, since for large central objects, randomly varying parts of the object’s shape may appear in the crop, rendering shape a less reliable feature relative to texture. We tested whether this was the case by comparing shape bias in random-crop and center-crop versions of AlexNet (57), VGG16 (86), ResNet-50 (42; 89; 49), and Inception-ResNet v2 (84). We evaluated shape bias using the GST images as described in Section 3 (see Appendix E.3 for details and additional discussion).

Center-crop augmentation reduced texture bias relative to random-crop augmentation (Table 1), and center-crop models had higher shape bias than random-crop models throughout the training process (Appendix C.2). Further, texture bias decreased as minimum crop area used in random-crop augmentation increased (Appendix B).

**Appearance-modifying data augmentation reduces texture bias.** The majority of photographs in major computer vision datasets are taken in well-lit environments (94). Humans encounter much more varied illumination conditions. We hypothesized that the development of human-like shape representations might require more diverse training data than what is present in ImageNet. Geirhos et al. (91) previously showed that neural style transfer data augmentation induces a model to classify images according to shape. Here, we tested whether a similar effect could be achieved with more naturalistic forms of data augmentation that do not require training on stylized images, and that more closely resemble the variation encountered by humans. Following (13), we tested six individual augmentations as well as their compositions (see Appendix E.3 for details). We applied augmentations to a randomly selected 50% of the total examples in each mini-batch.

Table 3: **The effect of augmentations that reduce texture bias is additive.** ResNet-50 models were trained on ImageNet for 90 epochs with random-crop augmentation and augmentation  $p = 0.5$  unless noted. Augmentations are cumulative across the rows (e.g. the “+ Gaussian blur” model used color distortion and Gaussian blur augmentation). “Stronger” augmentation used  $p = 0.75$ ; “Longer” training is 270 epochs. The final two models are shape-biased ( $>50\%$  shape-based classifications). Bolding indicates the highest value across augmentation types within a given column.

Augmentation(s)	Shape Bias	Shape Match	Texture Match	ImageNet		IN-Sketch		SIN	
				top-1	top-5	top-1	top-5	top-1	top-5
Baseline	19.5%	11.7%	<b>48.4%</b>	76.6%	<b>93.3%</b>	22.4%	39.3%	7.7%	17.0%
+ Color distortion	25.8%	15.3%	44.2%	<b>76.9%</b>	<b>93.3%</b>	28.1%	46.6%	9.9%	20.5%
+ Gaussian blur	30.7%	17.2%	38.8%	76.8%	<b>93.3%</b>	29.0%	47.9%	11.1%	21.9%
+ Gaussian noise	36.1%	20.1%	35.5%	75.9%	92.8%	29.8%	48.9%	12.6%	24.3%
+ Min. crop of 64%	48.7%	29.1%	30.7%	73.5%	91.5%	<b>30.9%</b>	<b>51.4%</b>	14.5%	<b>28.2%</b>
+ Stronger aug.	55.2%	33.3%	27.1%	72.0%	90.7%	30.4%	50.5%	<b>15.1%</b>	<b>28.8%</b>
+ Longer training	<b>62.2%</b>	<b>38.3%</b>	23.3%	71.1%	90.0%	30.5%	50.4%	14.9%	28.4%

We found that many simple augmentations decreased texture bias relative to baseline (Table 2), and these effects were stronger when paired with center crops as opposed to random crops (Table D.1). In addition, as shown in Table 3, the effect of augmentations that reduce texture bias was additive. A model trained for 90 epochs with strong color distortion, Gaussian blur, Gaussian noise, and less aggressive random crop augmentation was actively shape-biased (55.2%), and one trained for longer (270 epochs) was even more shape-biased (62.2%). These results show that it is possible to induce a shape bias using naturalistic augmentation that does not include stylization. Augmentation that reduced texture bias according to the GST dataset also improved accuracy on the ImageNet-Sketch (IN-Sketch) (96) dataset, which consists of human-drawn sketches of each ImageNet class, and Stylized ImageNet (SIN) (36), which was generated by processing the ImageNet test set with a style transfer algorithm. These datasets privilege shape by ablating texture information in other ways.

In general, we observed a trade-off between ImageNet top-1 and shape bias (Figure 3). In Appendix A, we show that manipulating learning rate and weight decay during training produced a similar trade-off.

## 6 Effect of training objective

A hypothesis consistent with the results above is that texture bias is driven by the joint image-label statistics of the ImageNet dataset. To correctly label the many dog breeds in the dataset, for instance, a model would have to make texture-based distinctions between similarly shaped objects. To test this hypothesis, we compared the shape bias of models trained with standard supervised learning to models trained with self-supervised objectives different from supervised classification. To explore the interaction of objective and model architecture, we used each objective to train models with both AlexNet and ResNet-50 base architectures.

**Self-supervised losses.** We paired each self-supervised objective with up to two architecture backbones: AlexNet and ResNet-50 v2 (Appendix E.4.1), allowing for comparison with architecture-matched supervised counterparts.

*Rotation classification.* Input images are rotated 0, 90, 180, or 270 degrees, and the task is to predict which rotation was applied (chance =  $1/4$ ) (39; 54). In their original presentation of this loss, Gidaris et al. (39) argued that to do well on this task, a model must understand which objects are present in an image, along with their location, pose, and semantic characteristics.

*Exemplar.* This objective, first introduced by (26), learns a representation where different augmentations of the same image are close in the embedding space. Our implementation follows (54). At

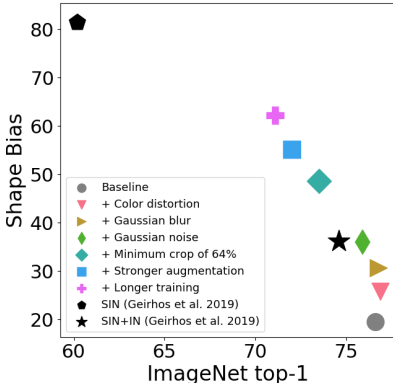


Figure 3: **Tradeoff between ImageNet top-1 and shape bias for models trained with different augmentations.** ResNet-50 models with naturalistic data augmentation achieve comparable tradeoffs to those from Geirhos et al. (36).

Table 4: **Both objective and base architecture affect shape bias.** We trained AlexNet and ResNet-50 base architectures on five objectives (rows). For SimCLR, we additionally include a baseline with the same augmentation. We froze the convolutional layers, reinitialized and retrained the fully connected layers, and evaluated the models using the GST stimuli. Shape match is generally higher for models with an AlexNet than ResNet-50 base architecture, while the reverse is true for texture. Rank order of shape bias across tasks is largely preserved across base architectures.

Objective	Shape Bias		Shape Match		Texture Match		ImageNet Top-1 Acc.	
	AlexNet	ResNet-50	AlexNet	ResNet-50	AlexNet	ResNet-50	AlexNet	ResNet-50
Supervised	29.8%	21.9%	17.5%	13.5%	41.2%	48.2%	57.0%	75.8%
Rotation	47.0%	32.3%	21.6%	14.2%	24.3%	29.8%	44.8%	44.4%
Exemplar	29.9%	14.4%	12.6%	7.5%	29.5%	44.7%	37.2%	41.8%
BigBiGAN	-	31.9%	-	17.7%	-	37.7%	-	55.4%
SimCLR	-	37.0%	-	17.3%	-	29.4%	-	69.2%
Supervised w/ SimCLR aug.	-	40.4%	-	23.1%	-	34.0%	-	76.3%

training time, each batch consists of 8 copies each of 512 dataset examples with different augmentations (random crops from the original image, converted to grayscale with a probability of  $2/3$ ). A triplet loss encourages distances between augmentations of the same example to be smaller than distances to other examples.

*BigBiGAN.* The BiGAN framework (24; 27) jointly learns a generator  $\mathcal{G}(\mathbf{z})$  that converts latent codes  $\mathbf{z}$  to images and an encoder  $\mathcal{E}(\mathbf{x})$  that converts dataset images  $\mathbf{x}$  to latent codes. At training time, the encoder and generator are optimized adversarially with a discriminator. The discriminator is optimized to distinguish pairs of sampled latents with their corresponding generator output ( $\mathbf{z}, \mathcal{G}(\mathbf{z})$ ) from pairs of dataset images with their corresponding latents ( $\mathcal{E}(\mathbf{x}), \mathbf{x}$ ), whereas the generator and encoder are optimized to minimize the discriminator’s performance. We use the representation from the penultimate layer of a ResNet-50 v2 encoder (25).

*SimCLR.* This model learns representations by maximizing agreement between differently augmented views of the same data example via a contrastive loss in the latent space. SimCLR uses a composition of augmentations, namely random crop, color distortion, and Gaussian blur, and leverages a non-linear transformation (projection head) to transform the representation before applying the contrastive loss (based on softmax cross-entropy). We follow the same setting as in (13), using the ResNet-50 base architecture, and training the model for 1000 epochs with a batch size of 4096. To test the contribution of augmentation alone, we created a supervised baseline with the same augmentation and training (“Supervised w/ SimCLR aug.” in Table 4). See also Appendix E.4.1.

**Evaluation of shape bias.** To facilitate comparison across supervised and self-supervised tasks, we trained classifiers on top of the learned representations using the standard ImageNet dataset and softmax cross-entropy objective, freezing all convolution layers and reinitializing and retraining later layers. We provide full training details in Appendix E.4.2. See also Appendix Tables D.2 and D.3.

**Results.** We found effects of both objective and base architecture on texture bias (Table 4). The rotation model had significantly lower texture bias than supervised models for both architectures; this may be because rotationally-invariant texture features are not useful for performing the rotation classification task. In general, shape match was higher for models with AlexNet than ResNet-50 architecture (log odds 1.04, 95% CI [0.83, 1.24], logistic regression; see Appendix E.4.3 for details), whereas the reverse was true for texture match (log odds 0.78, 95% CI [0.67, 0.92]). The effects of architecture and task appear to be largely independent: the rank order of shape bias across tasks was similar for the two model architectures. Amongst the unsupervised ResNet-50 models, SimCLR had the lowest texture bias. However, a supervised baseline with the same augmentation had a still-lower texture bias, suggesting that augmentation rather than objective was the more important factor.

## 7 Effect of architecture

**Shape bias and ImageNet accuracy.** Figure 4 shows the shape bias, shape match, and texture match of 16 high-performing ImageNet models trained with the same hyperparameters (Appendix E.5.1 for details). Both shape bias and shape match correlated with ImageNet top-1 accuracy, whereas texture match had no significant relationship. These results suggest that models selected for high ImageNet

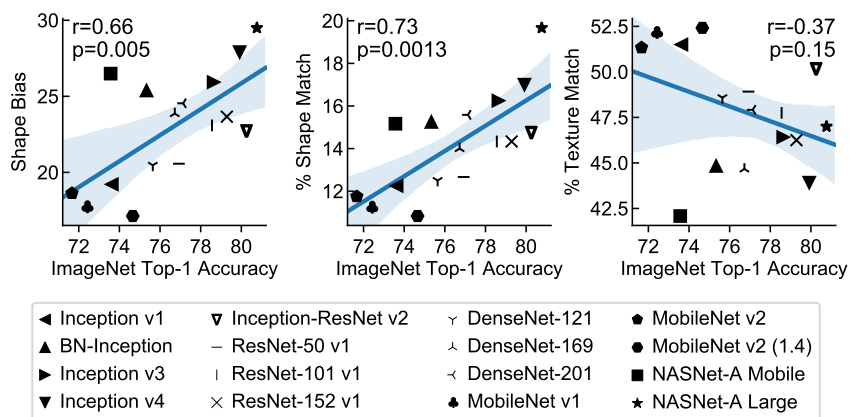


Figure 4: **Among high-performing ImageNet models, shape bias and accuracy correlate with ImageNet accuracy.**  $p$ -values indicate significance according to a  $t$  distribution. Blue line is a least squares fit to the plotted points; shaded area reflects 95% bootstrap confidence interval.

top-1 are more effective at extracting shape information. However, the asymmetry with Figure 3 shows that models selected for high shape bias do not necessarily have high ImageNet top-1.

**Shape bias in neurally motivated models.** Human visual judgments are well known to be shape-biased. Would a model explicitly built to match the primate visual system display lower texture bias than standard CNNs? We tested the shape bias of the CORNet models introduced by Kubilius et al. These models have architectures specifically designed to better match the primate ventral visual pathway both structurally (depth, presence of recurrent and skip connections, etc.) and functionally (behavioral and neural measurements) (59). We computed the shape bias of CORNet-Z, -R, and -S, using the publicly available trained models (18). The simplest of these models, CORNet-Z, had a shape bias of 14.9%, shape match of 9.2%, and texture match of 52.2%. CORNet-R, which incorporates recurrent connections, had a shape bias of 36.7%, and shape and texture accuracies of 19.6% and 33.8%, respectively. CORNet-S, the model with the highest BrainScore (81), had a shape bias of 20.3%, and shape and texture accuracies of 13.3% and 51.9%. Taken together, these models did not exhibit a greater shape bias than other models we tested. Perhaps surprisingly, the texture match was in the high range of those we observed.

**Shape bias of attention vs. convolution.** We wondered whether convolution itself could be a cause of texture bias. Ramachandran et al. (75) recently proposed a novel approach to image classification that replaces every convolutional layer in ResNet-50 v1 with local attention, where attention weights are determined based on both relative spatial position and content. We compared this model (with spatial extent  $k = 7$ , and 8 attention heads) against a baseline ResNet-50 v1 model trained with the same hyperparameters. The attention model had a shape bias of 20.2% (shape match: 12.8%; texture match: 50.7%), similar to the baseline’s shape bias of 23.2% (shape match: 14.4%; texture match: 47.7%). Thus, use of attention in place of convolution appears to have little effect upon texture bias.

## 8 Degree of representation of shape and texture in ImageNet models

Standard ImageNet-trained CNNs are biased towards texture in their classification decisions, but this does not rule out the possibility that shape information is still represented in layers of the model prior to the output. We tested the extent to which it is possible to decode stimulus shape versus texture in the final layers of AlexNet and ResNet-50. We trained linear multinomial logistic regression classifiers on two classification tasks. Taking as input activations from a layer of a frozen, ImageNet-trained model, the classifier predicted either (i) the shape of a GST image or (ii) its texture. We used center-crop AlexNet and ResNet-50, considering AlexNet “pool3” (final convolutional layer, including the max pool), “fc6” (first linear layer of the classifier, including the ReLU), and “fc7” (second linear layer), and ResNet-50 “pre-pool” (output of the final bottleneck layer) and “post-pool” (following the global average pool). See Section E.6 for additional details, and Figure C.4 for decoding results from all AlexNet layers.



**Results.** We found that, despite the high texture bias of AlexNet and ResNet-50, shape could nonetheless be decoded with high accuracy (Figures 5 and C.4). In fact, it was possible to classify shape (77.9%) more accurately than texture (65.6%) throughout the convolutional layers of AlexNet. In ResNet-50, while decoding accuracy for texture (80.9%) was higher than for shape (66.2%) for pre-pool, shape accuracy was still high. Interestingly, shape accuracy decreased through the fully-connected layers of AlexNet’s classifier, and from ResNet pre-pool to post-pool, suggesting that these models’ classification layers remove shape information.

## 9 Conclusion

From the perspective of the computer vision practitioner, our results indicate that models that prefer to classify images by shape rather than texture outperform baselines on some out-of-distribution test sets (IN-Sketch and SIN). We suggest practical ways to reduce texture bias, for example using additive augmentation (color distortion, blur, noise) and center crops. At the same time, we show that there is often a trade-off between shape bias and ImageNet top-1 accuracy.

On the scientific side, our findings flesh out the picture of how and why ImageNet-trained CNNs may differ from human vision. While both architecture and training objective have an effect on the level of texture bias in a model, the statistics of the training dataset are the most important factor. Changing these statistics using data augmentations qualitatively similar to those induced by the human visual system and visual environment is the most effective way to instill in CNNs shape-biased representations like those documented in the human psychological literature.

## Broader Impact

People who build and interact with tools for computer vision, especially those without extensive training in machine learning, often have a mental model of computer vision models as similar to human vision. Our findings contribute to a body of work showing that this view is actually far from correct, especially for ImageNet, one of the datasets most commonly used to train and evaluate models. Divergences between human and machine vision of the kind we study could cause users to make significant errors in anticipating and reasoning about the behavior of computer vision systems. Our findings contribute to a body of work delineating divergences between human and machine vision, and suggesting avenues for bringing the two systems closer together. Allowing people from a wide range of backgrounds to make safe, predictable, and equitable models requires vision systems to perform at least roughly in accordance with their expectations. Making computer vision models that share the same inductive biases as humans is an important step towards this goal. At the same time, we recognize the possible negative consequences of blindly constraining models’ judgments to agree with people’s: human visual judgments display forms of bias that should be kept out of computer models. More broadly, we believe that work like ours can have a beneficial impact on the internal sociology of the machine learning community. By identifying connections to developmental psychology and neuroscience, we hope to enhance interdisciplinary connections across fields, and to encourage people with a broader range of training and backgrounds to participate in machine learning research.

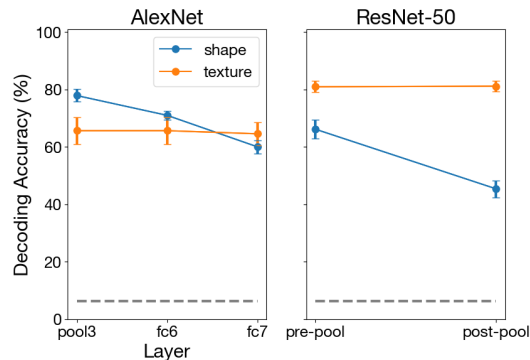


Figure 5: **Both the shape and texture of ambiguous images is decodable from the hidden representations of texture-biased ImageNet-trained models.** Performance of linear classifiers trained to classify shape (blue) or texture (orange) of the GST stimuli given layer activations from frozen ImageNet-trained AlexNet (left) and ResNet-50 (right). Performance is the maximum classification accuracy over the training period (mean  $\pm$  SD across 5 splits). Chance is 6.25%. At the final convolutional layer, both models contain substantial shape information, but it decreases subsequently.

## Acknowledgments and Disclosure of Funding

We thank Jay McClelland, Andrew Lampinen, Akshay Jagadeesh, and Chengxu Zhuang for useful conversations, and Guodong Zhang and Lala Li for comments on an earlier version of the manuscript. KLH was supported by NSF GRFP grant DGE-1656518.

## References

- [1] ALCORN, M. A., LI, Q., GONG, Z., WANG, C., MAI, L., KU, W.-S., AND NGUYEN, A. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 4845–4854.
- [2] AUTHORS, T. Tensorflow tpu implementation of resnet-50. <https://github.com/tensorflow/tpu/tree/master/models/official/resnet>.
- [3] AZULAY, A., AND WEISS, Y. Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177* (2018).
- [4] BAKER, N., LU, H., ERLIKHMAN, G., AND KELLMAN, P. J. Deep convolutional networks do not classify based on global object shape. *PLoS computational biology* 14, 12 (2018), e1006613.
- [5] BALLESTER, P., AND ARAUJO, R. M. On the performance of googlenet and alexnet applied to sketches. In *Thirtieth AAAI Conference on Artificial Intelligence* (2016).
- [6] BARBU, A., MAYO, D., ALVERIO, J., LUO, W., WANG, C., GUTFREUND, D., TENENBAUM, J., AND KATZ, B. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. In *Advances in Neural Information Processing Systems* (2019), pp. 9448–9458.
- [7] BASHIVAN, P., KAR, K., AND DICARLO, J. J. Neural population control via deep image synthesis. *Science* 364, 6439 (2019), eaav9436.
- [8] BENGIO, Y., BASTIEN, F., BERGERON, A., BOULANGER-LEWANDOWSKI, N., BREUEL, T., CHHERAWALA, Y., CISSE, M., CÔTÉ, M., ERHAN, D., EUSTACHE, J., ET AL. Deep learners benefit more from out-of-distribution examples. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (2011), pp. 164–172.
- [9] BERARDINO, A., LAPARRA, V., BALLÉ, J., AND SIMONCELLI, E. Eigen-distortions of hierarchical representations. In *Advances in neural information processing systems* (2017), pp. 3530–3539.
- [10] BIGGIO, B., CORONA, I., MAIORCA, D., NELSON, B., ŠRNDIĆ, N., LASKOV, P., GIACINTO, G., AND ROLI, F. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases* (2013), Springer, pp. 387–402.
- [11] BRENDDEL, W., AND BETHGE, M. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. *arXiv preprint arXiv:1904.00760* (2019).
- [12] CADIEU, C. F., HONG, H., YAMINS, D. L., PINTO, N., ARDILA, D., SOLOMON, E. A., MAJAJ, N. J., AND DICARLO, J. J. Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS computational biology* 10, 12 (2014), e1003963.
- [13] CHEN, T., KORNBLITH, S., NOROUZI, M., AND HINTON, G. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709* (2020).
- [14] CHILEY, V., SHARAPOV, I., KOSSON, A., KOSTER, U., REECE, R., DE LA FUENTE, S. S., SUBBIAH, V., AND JAMES, M. Online normalization for training neural networks. *arXiv preprint arXiv:1905.05894* (2019).

- [15] CICHY, R. M., KHOSLA, A., PANTAZIS, D., TORRALBA, A., AND OLIVA, A. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific reports* 6 (2016), 27755.
- [16] CIREŞAN, D. C., MEIER, U., GAMBARDELLA, L. M., AND SCHMIDHUBER, J. Deep, big, simple neural nets for handwritten digit recognition. *Neural computation* 22, 12 (2010), 3207–3220.
- [17] COHEN, J., ROSENFELD, E., AND KOLTER, Z. Certified adversarial robustness via randomized smoothing. In *Proceedings of the 36th International Conference on Machine Learning* (Long Beach, California, USA, 09–15 Jun 2019), K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97 of *Proceedings of Machine Learning Research*, PMLR, pp. 1310–1320.
- [18] <https://github.com/dicarlolab/CORnet>.
- [19] CUBUK, E. D., ZOPH, B., MANE, D., VASUDEVAN, V., AND LE, Q. V. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2019), pp. 113–123.
- [20] CUBUK, E. D., ZOPH, B., SHLENS, J., AND LE, Q. V. Randaugment: Practical data augmentation with no separate search. *arXiv preprint arXiv:1909.13719* (2019).
- [21] DEEPMIND. TensorFlow Hub: bigbigan-resnet50. <https://tfhub.dev/deepmind/bigbigan-resnet50/1>.
- [22] DEVRIES, T., AND TAYLOR, G. W. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* (2017).
- [23] DODGE, S., AND KARAM, L. A study and comparison of human and deep learning recognition performance under visual distortions. In *2017 26th international conference on computer communication and networks (ICCCN)* (2017), IEEE, pp. 1–7.
- [24] DONAHUE, J., KRÄHENBÜHL, P., AND DARRELL, T. Adversarial feature learning. *arXiv preprint arXiv:1605.09782* (2016).
- [25] DONAHUE, J., AND SIMONYAN, K. Large scale adversarial representation learning. *arXiv preprint arXiv:1907.02544* (2019).
- [26] DOSOVITSKIY, A., SPRINGENBERG, J. T., RIEDMILLER, M., AND BROX, T. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in neural information processing systems* (2014), pp. 766–774.
- [27] DUMOULIN, V., BELGHAZI, I., POOLE, B., MASTROPIETRO, O., LAMB, A., ARJOVSKY, M., AND COURVILLE, A. Adversarially learned inference. *arXiv preprint arXiv:1606.00704* (2016).
- [28] EFROS, A. A., AND FREEMAN, W. T. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (2001), ACM, pp. 341–346.
- [29] ELSAYED, G., SHANKAR, S., CHEUNG, B., PAPERNOT, N., KURAKIN, A., GOODFELLOW, I., AND SOHL-DICKSTEIN, J. Adversarial examples that fool both computer vision and time-limited humans. In *Advances in Neural Information Processing Systems* (2018), pp. 3910–3920.
- [30] ENGSTROM, L., ILYAS, A., SANTURKAR, S., TSIPRAS, D., TRAN, B., AND MADRY, A. Adversarial Robustness as a Prior for Learned Representations. *arXiv e-prints* (Jun 2019), arXiv:1906.00945.
- [31] FAHLMAN, S. E. An empirical study of learning speed in back-propagation networks. Tech. rep., Carnegie Mellon University, Computer Science Department, 1988.
- [32] FAWZI, A., AND FROSSARD, P. Manitest: Are classifiers really invariant? *arXiv preprint arXiv:1507.06535* (2015).

- [33] FORD, N., GILMER, J., CARLINI, N., AND CUBUK, D. Adversarial examples are a natural consequence of test error in noise. *arXiv preprint arXiv:1901.10513* (2019).
- [34] GATYS, L. A., ECKER, A. S., AND BETHGE, M. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 2414–2423.
- [35] GEIRHOS, R., JACOBSEN, J.-H., MICHAELIS, C., ZEMEL, R., BRENDEL, W., BETHGE, M., AND WICHMANN, F. A. Shortcut learning in deep neural networks. *arXiv preprint arXiv:2004.07780* (2020).
- [36] 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).
- [37] GEIRHOS, R., TEMME, C. R., RAUBER, J., SCHÜTT, H. H., BETHGE, M., AND WICHMANN, F. A. Generalisation in humans and deep neural networks. In *Advances in Neural Information Processing Systems* (2018), pp. 7538–7550.
- [38] GERSHKOFF-STOWE, L., AND SMITH, L. B. Shape and the first hundred nouns. *Child development* 75, 4 (2004), 1098–1114.
- [39] GIDARIS, S., SINGH, P., AND KOMODAKIS, N. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* (2018).
- [40] GIRSHICK, R., DONAHUE, J., DARRELL, T., AND MALIK, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2014), pp. 580–587.
- [41] GOODFELLOW, I. J., SHLENS, J., AND SZEGEDY, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [42] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778.
- [43] HEEGER, D. J., AND BERGEN, J. R. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (1995), Citeseer, pp. 229–238.
- [44] HENDRYCKS, D., AND DIETTERICH, T. G. Benchmarking neural network robustness to common corruptions and surface variations. *arXiv preprint arXiv:1807.01697* (2018).
- [45] HENDRYCKS, D., MU, N., CUBUK, E. D., ZOPH, B., GILMER, J., AND LAKSHMINARAYANAN, B. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781* (2019).
- [46] HOSSEINI, H., AND POOVENDRAN, R. Semantic adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (2018), pp. 1614–1619.
- [47] HOSSEINI, H., XIAO, B., JAISWAL, M., AND POOVENDRAN, R. Assessing shape bias property of convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (2018), pp. 1923–1931.
- [48] ILYAS, A., SANTURKAR, S., TSIPRAS, D., ENGSTROM, L., TRAN, B., AND MADRY, A. Adversarial examples are not bugs, they are features. *arXiv preprint arXiv:1905.02175* (2019).
- [49] ImageNet training script for torchvision model implementations (used with modification).

<https://github.com/pytorch/examples/blob/master/imagenet/main.py>.

- [50] JOHNSON, J., ALAHI, A., AND FEI-FEI, L. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision* (2016), Springer, pp. 694–711.
- [51] KAUR, S., COHEN, J., AND LIPTON, Z. C. Are perceptually-aligned gradients a general property of robust classifiers? *arXiv preprint arXiv:1910.08640* (2019).
- [52] KHALIGH-RAZAVI, S.-M., AND KRIEGESKORTE, N. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology* 10, 11 (2014), e1003915.
- [53] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [54] KOLESNIKOV, A., ZHAI, X., AND BEYER, L. Revisiting self-supervised visual representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019).
- [55] KORNBLITH, S., SHLENS, J., AND LE, Q. V. Do better imagenet models transfer better? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 2661–2671.
- [56] KRIZHEVSKY, A. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997* (2014).
- [57] KRIZHEVSKY, A., SUTSKEVER, I., AND HINTON, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (2012), pp. 1097–1105.
- [58] KUBILIUS, J., BRACCI, S., AND DE BEECK, H. P. O. Deep neural networks as a computational model for human shape sensitivity. *PLoS computational biology* 12, 4 (2016), e1004896.
- [59] KUBILIUS, J., SCHRIMPF, M., NAYEBI, A., BEAR, D., YAMINS, D. L., AND DICARLO, J. J. Cornet: modeling the neural mechanisms of core object recognition. *BioRxiv* (2018), 408385.
- [60] KUCKER, S. C., SAMUELSON, L. K., PERRY, L. K., YOSHIDA, H., COLUNGA, E., LORENZ, M. G., AND SMITH, L. B. Reproducibility and a unifying explanation: Lessons from the shape bias. *Infant Behavior and Development* (2018).
- [61] LANDAU, B., SMITH, L. B., AND JONES, S. S. The importance of shape in early lexical learning. *Cognitive development* 3, 3 (1988), 299–321.
- [62] LECUN, Y., BOTTOU, L., BENGIO, Y., AND HAFNER, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86, 11 (1998), 2278–2324.
- [63] LI, Y., WEI, C., AND MA, T. Towards explaining the regularization effect of initial large learning rate in training neural networks. *arXiv preprint arXiv:1907.04595* (2019).
- [64] LIM, S., KIM, I., KIM, T., KIM, C., AND KIM, S. Fast autoaugment. In *Advances in Neural Information Processing Systems* (2019), pp. 6662–6672.
- [65] LOPES, R. G., YIN, D., POOLE, B., GILMER, J., AND CUBUK, E. D. Improving robustness without sacrificing accuracy with patch gaussian augmentation. *arXiv preprint arXiv:1906.02611* (2019).
- [66] MADRY, A., MAKELOV, A., SCHMIDT, L., TSIPRAS, D., AND VLADU, A. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations* (2018).
- [67] MEHTA, D., KIM, K. I., AND THEOBALT, C. On implicit filter level sparsity in convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 520–528.
- [68] MÜLLER, R., KORNBLITH, S., AND HINTON, G. When does label smoothing help? In *Advances in Neural Information Processing Systems* (2019).

- [69] NAVON, D. Forest before trees: The precedence of global features in visual perception. *Cognitive psychology* 9, 3 (1977), 353–383.
- [70] NGUYEN, A., YOSINSKI, J., AND CLUNE, J. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2015).
- [71] PONCE, C. R., XIAO, W., SCHADE, P. F., HARTMANN, T. S., KREIMAN, G., AND LIVINGSTONE, M. S. Evolving images for visual neurons using a deep generative network reveals coding principles and neuronal preferences. *Cell* 177, 4 (2019), 999–1009.
- [72] PORTILLA, J., AND SIMONCELLI, E. P. A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision* 40, 1 (2000), 49–70.
- [73] RAGHUNATHAN, A., XIE, S. M., YANG, F., DUCHI, J. C., AND LIANG, P. Adversarial training can hurt generalization. *arXiv preprint arXiv:1906.06032* (2019).
- [74] RAJALINGHAM, R., ISSA, E. B., BASHIVAN, P., KAR, K., SCHMIDT, K., AND DICARLO, J. J. Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience* 38, 33 (2018), 7255–7269.
- [75] RAMACHANDRAN, P., PARMAR, N., VASWANI, A., BELLO, I., LEVSKAYA, A., AND SHLENS, J. Stand-alone self-attention in vision models. *arXiv preprint arXiv:1906.05909* (2019).
- [76] RICHARDWEBSTER, B., ANTHONY, S., AND SCHEIRER, W. Psyphy: A psychophysics driven evaluation framework for visual recognition. *IEEE transactions on pattern analysis and machine intelligence* (2018).
- [77] RITTER, S., BARRETT, D. G., SANTORO, A., AND BOTVINICK, M. M. Cognitive psychology for deep neural networks: A shape bias case study. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (2017), ICML’17, JMLR.org, pp. 2940–2949.
- [78] RUSAK, E., SCHOTT, L., ZIMMERMANN, R., BITTERWOLF, J., BRINGMANN, O., BETHGE, M., AND BRENDEL, W. Increasing the robustness of dnns against image corruptions by playing the game of noise. *arXiv preprint arXiv:2001.06057* (2020).
- [79] RUSSAKOVSKY, O., DENG, J., SU, H., KRAUSE, J., SATHEESH, S., MA, S., HUANG, Z., KARPATY, A., KHOSLA, A., BERNSTEIN, M., ET AL. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.
- [80] SANTURKAR, S., TSIPRAS, D., TRAN, B., ILYAS, A., ENGSTROM, L., AND MADRY, A. Computer vision with a single (robust) classifier. *CoRR abs/1906.09453* (2019).
- [81] SCHRIMPF, M., KUBILIUS, J., HONG, H., MAJAJ, N. J., RAJALINGHAM, R., ISSA, E. B., KAR, K., BASHIVAN, P., PRESCOTT-ROY, J., SCHMIDT, K., ET AL. Brain-score: which artificial neural network for object recognition is most brain-like? *BioRxiv* (2018), 407007.
- [82] SEABOLD, S., AND PERKTOLD, J. Statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference* (2010).
- [83] SERMANET, P., EIGEN, D., ZHANG, X., MATHIEU, M., FERGUS, R., AND LECUN, Y. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229* (2013).
- [84] SILBERMAN, N., AND GUADARRAMA, S. TensorFlow-Slim implementation of Inception-ResNet v2. <https://github.com/tensorflow/models/tree/master/research/slim>.
- [85] <https://github.com/google-research/simclr>.
- [86] SIMONYAN, K., AND ZISSERMAN, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

- [87] SZEGEDY, C., ZAREMBA, W., SUTSKEVER, I., BRUNA, J., ERHAN, D., GOODFELLOW, I., AND FERGUS, R. Intriguing properties of neural networks. In *International Conference on Learning Representations* (2014).
- [88] TAORI, R., DAVE, A., SHANKAR, V., CARLINI, N., RECHT, B., AND SCHMIDT, L. When robustness doesn't promote robustness: Synthetic vs. natural distribution shifts on imagenet, 2020.
- [89] TORCHVISION CONTRIBUTORS. PyTorch implementations of AlexNet, VGG16, ResNet-50. <https://pytorch.org/docs/stable/torchvision/models.html>.
- [90] TORCHVISION CONTRIBUTORS. transforms.RandomResizedCrop. <https://github.com/pytorch/vision/blob/7c9bbf5bdf68564a511578177dc8054e3c66fdf3/torchvision/transforms/transforms.py#L601-L604>.
- [91] <https://github.com/rgeirhos/texture-vs-shape>.
- [92] TRAMÈR, F., AND BONEH, D. Adversarial training and robustness for multiple perturbations. In *Advances in Neural Information Processing Systems* (2019).
- [93] TSIPRAS, D., SANTURKAR, S., ENGSTROM, L., TURNER, A., AND MADRY, A. Robustness may be at odds with accuracy. In *International Conference on Learning Representations* (2019).
- [94] TSOTSOS, J. K., KOTSERUBA, I., ANDREOPOULOS, A., AND WU, Y. A possible reason for why data-driven beats theory-driven computer vision. *arXiv preprint arXiv:1908.10933* (2019).
- [95] VISUAL TASK ADAPTATION BENCHMARK. sup-100, rotation, and exemplar tf-hub modules. <https://tfhub.dev/s?publisher=vtab>.
- [96] WANG, H., GE, S., LIPTON, Z., AND XING, E. P. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems* (2019), pp. 10506–10518.
- [97] 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* (2017), pp. 4148–4158.
- [98] XIAO, C., ZHU, J.-Y., LI, B., HE, W., LIU, M., AND SONG, D. Spatially transformed adversarial examples. In *International Conference on Learning Representations* (2018).
- [99] XIE, C., TAN, M., GONG, B., WANG, J., YUILLE, A., AND LE, Q. V. Adversarial examples improve image recognition. *arXiv preprint arXiv:1911.09665* (2019).
- [100] XIE, C., WU, Y., MAATEN, L. V. D., YUILLE, A. L., AND HE, K. Feature denoising for improving adversarial robustness. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 501–509.
- [101] YAMINS, D. L., HONG, H., CADIEU, C. F., SOLOMON, E. A., SEIBERT, D., AND DICARLO, J. J. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences* 111, 23 (2014), 8619–8624.
- [102] YIN, D., LOPES, R. G., SHLENS, J., CUBUK, E. D., AND GILMER, J. A fourier perspective on model robustness in computer vision. In *Neural Information Processing Systems* (2019).
- [103] 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* (2019), pp. 6023–6032.
- [104] ZHAI, X., PUIGSERVER, J., KOLESNIKOV, A., RUYSSSEN, P., RIQUELME, C., LUCIC, M., DJOLONGA, J., PINTO, A. S., NEUMANN, M., DOSOVITSKIY, A., ET AL. The visual task adaptation benchmark. *arXiv preprint arXiv:1910.04867* (2019).

- [105] ZHANG, R., ISOLA, P., EFROS, A. A., SHECHTMAN, E., AND WANG, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 586–595.
- [106] ZHOU, Z., AND FIRESTONE, C. Humans can decipher adversarial images. *Nature communications* 10, 1 (2019), 1334.
- [107] ZHU, Z., XIE, L., AND YUILLE, A. L. Object recognition with and without objects. *arXiv preprint arXiv:1611.06596* (2016).