

Can Biases in ImageNet Models Explain Generalization?

Paul Gavrikov

IMLA, Offenburg University

paul.gavrikov@hs-offenburg.de

Janis Keuper

IMLA, Offenburg University and University of Mannheim

keuper@imla.ai

Abstract

The robust generalization of models to rare, in-distribution (ID) samples drawn from the long tail of the training distribution and to out-of-training-distribution (OOD) samples is one of the major challenges of current deep learning methods. For image classification, this manifests in the existence of adversarial attacks, the performance drops on distorted images, and a lack of generalization to concepts such as sketches. The current understanding of generalization in neural networks is very limited, but some biases that differentiate models from human vision have been identified and might be causing these limitations. Consequently, several attempts with varying success have been made to reduce these biases during training to improve generalization. We take a step back and sanity-check these attempts. Fixing the architecture to the well-established ResNet-50, we perform a large-scale study on 48 ImageNet models obtained via different training methods to understand how and if these biases - including shape bias, spectral biases, and critical bands - interact with generalization. Our extensive study results reveal that contrary to previous findings, these biases are insufficient to accurately predict the generalization of a model holistically.

We provide access to all checkpoints and evaluation code at https://github.com/paulgavrikov/biases_vs_generalization/

1. Introduction

Artificial neural networks have achieved outstanding performance in various tasks, particularly excelling in vision tasks like image classification - commonly benchmarked on ImageNet [45]. While accuracy on ImageNet has dramatically improved in the last 15 years [11, 19, 30], with models nearly matching human performance, a critical issue remains: these models often fail to generalize beyond the specific data distribution they were trained on. This fragility becomes evident when faced with distribution shifts, such as changes in weather conditions [55], digital artifacts [21], or even replacing photos objects with artistic renditions

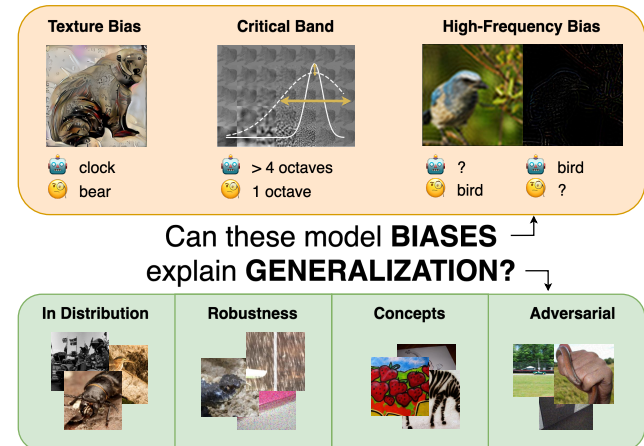


Figure 1. We study the influence of three selected biases that separate models 🤖 from humans 🧑 on the generalization of ImageNet models. Our study suggests that no single bias correlates with generalization in a holistic sense. We measure the texture/shape bias [17], critical band [52], and low/high-frequency spectral biases [61] on 48 models and correlate these biases against generalization that we measure on several benchmarks belonging to four categories: in distribution, robustness, conceptual changes, and adversarial robustness.

[23]. On an abstract level this can be unified by the vulnerability to adversarial attacks [3, 53], where carefully crafted image perturbations can fool models while remaining imperceptible to humans. This raises serious concerns about deploying such models in safety-critical applications where human lives are at stake [4, 13].

Seeking explanations for this lack of generalization, researchers have identified various biases that differentiate model from human perception. Three recent examples that have received significant attention are the *texture/shape bias* (models rely heavily on texture for object recognition, whereas humans prioritize shape information) [17], *critical band bias* (models utilize wider critical bands than humans for recognizing ImageNet objects) [52], and *high-frequency bias* (models exhibit a strong dependence on high-frequency information and perform poorly on low-

pass filtered data) [61].

Following the intuition that aligning these biases with human perception might improve generalization, several studies have explored regularization techniques during training [31, 33, 34, 46]. However, the fundamental question (visualized in Fig. 1) remains: *how well do these biases actually correlate with the ability to generalize?*

In this study, we take a step back and analyze these biases in ImageNet models, investigating their connection to generalization performance under various aspects. A cornerstone of this analysis is that we explore changes *in training* (e.g., augmentation, contrastive learning, and others), *while keeping the model architecture constant* to isolate the effect of inductive biases. We assess generalization by measuring performance on different datasets modeling in-distribution data, corrupted data, conceptual changes, and under adversarial attacks.

This investigation aims to shed light on the true impact of these biases and inform the development of more robust and generalizable models.

We summarize our contributions as follows:

- We systematically evaluate a diverse set of 48 models trained using various techniques for ImageNet classification. We assess their performance across multiple generalization benchmarks, including in-distribution and out-of-distribution tests related to robustness, concept generalization, and adversarial robustness. Then, we examine correlations between these benchmarks and biases like shape bias, spectral biases, and characterizations of the critical band.
- Our analysis highlights that previous studies linking generalization to tested biases often examined only partial aspects (e.g., specific training or benchmarks), introduced confounders by inductive biases, or encountered challenges due to noisy measurements originating from system noise and incompatible data transformation pipelines.
- **Our findings reveal several key insights:** **i)** none of the tested biases can singularly explain generalization - sometimes they are even negatively correlated with human perception; **ii)** most of the previously observed correlations show many outliers; **iii)** in our evaluation, some biases exhibit no correlation at all; **iv)** we find a surprising correlation between a high-frequency bias and generalization.

2. Related Work

Texture/Shape Bias. Geirhos et al. [17] discovered that CNNs primarily rely on texture cues to identify objects (in ImageNet). This is in stark contrast to the strong shape-biased decision observed in human perception. They argue that improving the shape bias may be linked to improved robustness and show that (pre-)training on a dataset without discriminative texture information indeed improves ro-

bustness on some benchmarks including common corruptions [21]. Following this observation, multiple works have attempted to understand origins of this bias [26, 28], improve shape bias in training [31, 34, 48], or study representations beyond ImageNet-CNNs [10, 16, 18, 29, 40]. Prior works have also shown a correlation between adversarial-training (AT) [35] and shape bias [15, 18]. However, an improved shape bias seems not to be the cure-all for generalization: [39] designed an augmentation technique that improves shape bias but not robustness to common corruptions [21] and, thus, shows a counter-example to reject a causal connection at least for this aspect of generalization. In our study, we test the shape bias beyond common corruptions to derive a more holistic understanding of this bias in the context of generalization.

Spectral Bias. Natural signals, such as images [44] and sounds [49], concentrate most information in low-frequency bands. As such, it seems to be reasonable to draw predictive cues from this band. Contrary, Wang et al. [61] demonstrated that image classification models reach high accuracies on high-pass filtered images that mostly resemble noise and should not carry discriminative information. Low-pass filtered images are easily recognizable by humans, on the other hand, are poorly recognized by the same models. While this seems to be more extreme on low-resolution images, higher-resolution models such as ImageNet classifiers also unreasonably heavily rely on cues from high-frequency (HF) bands [1]. [61] also link this *spectral* bias to adversarial robustness [3, 53] and show that adversarial training (AT) [35] results in a reversed bias. Indeed, perturbations generated by adversarial attacks on ImageNet seem to particularly target HF bands [68]. A similar connection to non-adversarial robustness can be made, as most common corruptions [21] primarily target HF bands [68]. This bias has received attention in the robustness community, where some hope to improve robustness by stronger low-frequency biases [33, 34, 46]. We seek to understand how spectral biases, in particular, low/high-frequency biases affect generalization.

Critical Band. Recently, Subramanian et al. [52] have found another bias that separates human from model vision in the *critical band*. This band defines what spatial-frequency channels are used to detect objects. On ImageNet, humans use a one-octave wide channel, yet models measure significantly wider channels, making them more sensitive to noise perturbations in spectral bands that do not affect human vision. The study showed a strong correlation between the parameters of the critical band and shape bias, as well as the adversarial robustness of adversarially-trained networks. Surprisingly, adversarial training (AT) seems to further increase the bandwidth and, thus, increases

misalignment. We propose a few optimizations to the underlying test process and as for the previous biases, aim to understand possible connections to generalization beyond adversarial robustness/training.

3. Methodology

In our study, we measure models’ ability to generalize and correlate this to the biases that separate human and model vision. The following paragraphs show details about how we measure biases, generalization, and correlation, as well as our tested models.

3.1. Measuring Biases

Texture/Shape Bias. We measure the texture/shape bias on the *cue-conflict* dataset [17]. This dataset consists of ImageNet samples where shape and texture cues belong to different (*conflicting*) ImageNet super-classes (e.g., a bear with clock skin as shown in Fig. 1). We report the *shape bias* which is defined by the ratio of shape to all correct decisions on the dataset - with a value of 0/1 indicating a texture/shape bias, respectively. Humans achieve a strong average shape bias of 0.96 but most models are severely texture biased.

Spectral Bias. To understand the spectral bias we measure the performance of bandpass-filtered ImageNet samples, akin to [61]. We differentiate between low and high-frequency biases by filtering the test data at different cutoff frequencies. For instance, a low-pass filter with a cutoff of 30% will contain the lowest 30% of the Fourier spectrum. To compensate for different test performances of the models on the clean data, we report the *ratio* of the accuracy on the filtered data to the clean accuracy. In some cases, this may result in ratios > 1 , i.e., the accuracy on filtered data may exceed the original validation accuracy. As low/high-frequency bands are ill-defined we average the performance observed at 10, 20, 30, and 40% cutoffs along the low/high-pass filter tests to report the *low/high-frequency bias*, respectively.

Critical Band. Following the methodology of Subramanian et al. [52], we measure the critical band by insertions of Gaussian noise to contrast-reduced and gray-scaled ImageNet samples with varying strength at different frequencies. Areas where the accuracy falls below a predefined threshold mark the critical band. The parameters of the band are obtained by the fitting of a Gaussian characterized by the *bandwidth (BW)*, *center frequency (CF)*, and *peak-noise sensitivity (PNS)*. A caveat of the original test is, that the critical band is determined by conditions where the accuracy (on a 16-class ImageNet subset) drops below 50%. However, adversarially-trained models perform poorly on

contrast-reduced images [15, 18]. To compensate for this fact, we *normalize* the accuracy by the performance on the contrast-reduced images before applying noise. Furthermore, the original test only uses a few ImageNet samples for each condition. Instead, we use all ImageNet validation samples. Finally, we determine the accuracy against all 1,000 classes as opposed to 16 super-classes. Taken together this results in a more computationally expensive but also more accurate measurement. In Appendix F, we also show results with the original test routine - with and without normalization.

3.2. Measuring Generalization

To benchmark generalization we loosely follow the suggestions of prior work [23] and rely on multiple semantically-different benchmarks to faithfully assess facets. Our benchmarks can be separated into **four categories**: *in-distribution (ID)* datasets with considerable statistical similarity to the ImageNet training set, *robustness* datasets that apply various corruptions to ImageNet samples, datasets challenging models to recognize objects based on *concepts* (e.g. sketches), and *adversarial attacks* to test robustness adaptively. For all benchmarks we evaluate the top-1 accuracy using a resolution of 224^2 px and *consistent* transformation pipelines on the *same system* to avoid the risk of falsifications due to system noise [63]. We report the average over the individual benchmarks for each of the four categories. In the following, we briefly introduce the benchmarks but refer the reader to Appendix C for full details.

In Distribution (ID). We use the *ImageNet (IN)* [45] validation set from the ILSVRC2012 challenge, *ImageNet v2 (IN-v2)* [43] which is a newer test set sampled a decade later following the methodology of the original curation routine, and *ImageNet-Real (IN-Real)* [2] which is a re-annotated version of the original ImageNet validation set. IN-Real assigns multiple labels per image and contains multiple corrections of the original annotations.

Robustness. We benchmark robustness on *ImageNet-C (IN-C)* [21] which applies 19 synthetic corruptions under five levels of severity and its successor *ImageNet- \bar{C} (IN- \bar{C})* [36] which extends this set by 10 additional corruptions which were chosen to be perceptually dissimilar but conceptually similar to the corruptions in IN-C. Additionally, we use *ImageNet-A (IN-A)* [24] which contains samples from the full ImageNet that are naturally hard to classify and, thus, pose natural adversarial examples.

Concepts. To understand model capabilities to recognize objects under conceptual changes, we use *ImageNet-Renditions (IN-R)* [23] representing ImageNet objects

through cartoons, paintings, toys, etc., *ImageNet-Sketch (IN-S)* [60] consisting of hand-drawn sketches of ImageNet objects, and *Stylized ImageNet (SIN)* [17] containing ImageNet validation images that have been stylized using different artistic filters to destroy the correlation between texture information and class labels.

Adversarial Robustness. We apply a Project Gradient Descent (PGD) [35] attack on the ImageNet validation set to adaptively evaluate robustness. Contrary to the previous benchmarks, this benchmark is *adaptive* as it can continuously generate weaknesses for a given model. In theory, a model might overfit a static dataset due to its limited number of test samples; however, an (ideal) adaptive benchmark remains unaffected by this limitation.

3.3. Models

We conduct our main analysis on models based on the ResNet-50 architecture [19] that were trained on ImageNet [45]¹. Fixing the architecture is crucial for our study, as each architecture has its own inductive biases that may influence results. Thus, to disentangle the inductive bias from the influence on generalization we ensure that all models use an identical architecture. This also means avoiding even minor changes like different activation functions as they can drastically affect the performance and introduce confounders to our analysis [66]. Limiting the study to ResNet-50 has additional benefits: the architecture is highly popular and comes without bells and whistles. This status also ensures that a variety of techniques have been evaluated for this architecture, specifically. The architecture also represents architectures used in real-life applications, where training from scratch is necessary but computational resources are limited. Further, we ensured no model was trained with external data to avoid target leakage on our benchmarks. Most models were also not specifically optimized for the tested biases which provides an opportunity to understand if and to what extent (aligned) biases arise through changes in training.

To represent changes in generalization, we use 48 models that differ *in training*. We loosely sort these models into seven categories that we color-code. Throughout the study, we use different marker symbols to differentiate individual models. Marker sizes of adversarially-trained models correlate with the attack budget ϵ during training. A full model legend including all measurements can be found in Appendix B.

Our tested models belong to the following categories:

- a *baseline* model following the original training [19];
- *augmentation* techniques [12, 22, 23, 25, 29, 31, 37, 38];

¹Appendix D also contains results for ViT-B/16 models.

- *(pre)training* on *SIN* which can be seen as an extreme form of augmentation that breaks correlations between textures and class labels (SHAPENET) [18];
- *adversarial training* [35] against a PGD adversary with increasing attack budget ϵ under ℓ_2 and ℓ_∞ -norm [47]. Technically, this too could be seen as augmentation, but we will show significant differences for this group of models;
- *contrastive* learners [5–8] with supervised finetuning of the classification-head;
- improved *training recipes* in TIMM [64] based on the findings of Wightman et al. [65] and PYTORCH [41, 59];
- a model with *randomly weighted* convolution filters [14]

3.4. Measuring correlations

We measure correlations via the Spearman rank-order correlation coefficient r . Due to the low sample size ($N \ll 500$), we also assess the statistical significance by obtaining a p -value from a two-sided permutation test. p is given by the ratio of generated r larger than the baseline r on all samples. We consider results as significant if $p < 0.05$.

4. Limitations of previous studies

Previous studies have already correlated these biases with (some aspects of) generalization [17, 18, 52, 61]. In this section, we discuss limitations in these studies and show why their findings may not transfer to a more holistic view of generalization.

Inductive biases of architectures were not ablated. The architecture plays a significant role in robustness [54] and, thus, can also be expected to affect a more holistic view of generalization. As such, previously discovered correlations may be spurious concerning the bias itself and just an effect of the architecture. To suppress this potential confounder our model zoo is limited to ResNet-50 models without any modifications to the original architecture as defined in [19]. As a matter of fact, we observe that many correlations no longer hold when we fix the architecture.

Findings on subsets do not generalize. It is well known that correlation does not imply causation and, thus, it may not be surprising that we find that many observed correlations no longer hold for our study. Common causes include (i) benchmarking only a specific type of training, *e.g.*, adversarial training. Yet, correlations for adversarial training may be differently strong or even reversed compared to other forms of training (for example see Fig. 2); and/or (ii) using only one benchmark, *e.g.*, adversarial attacks. Performance on many ImageNet-benchmarks is moderate to strongly correlated [55], which we also confirm in our study (Fig. 3) but clear outliers exist. For instance, adversarial robustness (here: PGD) is not correlated to most other bench-




Figure 2. **Biases often only correlate with specific aspects of generalization or model groups.** We measure Spearman r correlations on all models (*Total*) and separately on adversarially-trained (*AT*), and all *other* models, as there is often a different trend. Non-significant correlations with $p \geq 0.05$ are set to 0. Please note that r does not capture non-monotonic relations.




Figure 3. **Many benchmarks show notable positive correlations between each other - except SIN and PGD (adversarial attack).** Correlations measured by Spearman r . We set non-significant correlations with $p \geq 0.05$ to 0.

marks (Fig. 3) and generally might be divergent to other forms of robustness [32, 42, 51, 58]. Our study focuses on generalization beyond but including adversarial robustness and is performed on various models including adversarially-trained ones.

Inconsistent benchmarking. While the surrounding details of benchmarks in previous studies are not always clear, a common procedure is to simply copy benchmark results reported in publications. However, concerningly, we found that some authors produced results using inconsistent transformation pipelines during testing resulting in non-comparable results. Additionally, even system noise can lead to significant errors in evaluations *e.g.*, due to low-level implementation details in data loaders such as JPEG decompression or subsampling techniques [63]. These issues advocate for benchmarking under comparable condi-




Figure 4. **Shape Bias vs. Generalization.** A value of 0 indicates a texture bias, and 1 is a shape bias.

tions to obtain comparable results. Therefore, we produce all benchmarks on the same system, using the same environment, and identical test transformations for all datasets.

5. Can these biases explain generalization?

In this section, we apply the previously introduced methodology on each bias individually, to understand if and to what extent aligning these biases can help with generalization. Figure 2 shows an overview of the *statistical* correlation for all tested biases. Additionally, we inspect scatterplots (Figs. 4 to 6) to detect non-monotonic correlations missed by Spearman r and pay attention to whether outliers exist.

5.1. Shape Bias

For the shape bias, we observe a strong **negative** correlation to ID tests (Fig. 4) - suggesting that increasing misalignment with human vision increases ID accuracy. While it has been shown before, that ImageNet-CNNs are strongly texture-biased and ImageNet can be well separated by texture alone [17], the monotonic decrease in performance with increasing shape bias without significant outliers is quite surprising. While this does not prove causality, it raises the question of whether strongly shape-biased models can achieve significant performance on ImageNet (without relying on extra data). Similarly, we only note significant improvements in IN-A on strongly texture-biased models (see Appendix E for a plot). Usually, poor performance on IN-A is linked to spurious cues [24] and one may assume that these are related to textures (similar to adversarial attacks on fragile features [27]). However, it seems like these

cues are more shape-related and a certain degree of texture bias is necessary to overcome them. But neither does a texture bias guarantee the best performance, as the strongest texture-biased model (TIMM (A3) [65]) with a shape bias of 0.13 still performs 10.4% worse on IN-A than the best model (TORCHVISION (V2)) with a (higher) shape bias of 0.17.

On other aspects of generalization, we often obtain unrelated trends between adversarially trained and other models. For adversarially-trained (AT) models, we observe a curvilinear (inverted “U”) relation (which is not captured by the Spearman r) to all aspects of generalization which also seems to correlate with the attack budget during training (pay attention to the marker size). For concepts and adversarial robustness, we observe the best performance by AT models that balance shape and texture bias (~ 0.5). There are no notable differences in the trends between models trained on ℓ_2 and ℓ_∞ -norm attacks. Non-AT models only show a (moderate) correlation to concepts.

Our results extend previous findings showing that shape bias is not causally correlated with IN-C performance [39] to a much broader view of generalization, and contrast the claims about improved robustness in [17]. For adversarial robustness, shape bias may define a ceiling where a balanced representation of shapes and textures works best. However, the model set of top performers exclusively consists of adversarially-trained models which may lead to spurious correlations. Shape bias may also improve recognition of concepts - which is intuitive, as these samples do not contain the same textures as ImageNet samples.

In all the models, the only regular pattern we see is that the shape bias is inversely related to ID performance. This means that as ID accuracy increases, alignment with human vision decreases.

5.2. Spectral Bias

It is well known that adversarial training results in models with improved low-frequency but reduced high-frequency bias [15, 18, 61, 68]. Yet, the spectral bias of other methods has been less studied, especially in correlation to other aspects of generalization beyond adversarial robustness.

Low-Frequency Bias. To our surprise, we find no meaningful correlation between a low-frequency bias and most aspects of generalization when considering all models (Fig. 5) - except for adversarial robustness, where a stronger low-frequency bias improves adversarial robustness. This is intuitive, as attacks prefer high-frequency bands on ImageNet [68]. Improved reliance on cues from non-affected bands is therefore a good strategy.

When separating AT from other models we again observe divergent trends. On AT, this bias shows a curvilinear correlation to all aspects of generalization (not captured by




Figure 5. **Spectral Biases vs. Generalization.** We correlate low-frequency (top) and high-frequency bias (bottom).

Spearman r in Fig. 2). For concepts and adversarial robustness, some low-frequency bias helps but eventually starts hurting generalization. For ID and non-adversarial robustness, increasing low-frequency bias seems always to reduce generalization. On non-AT models, it shows a moderate monotonic correlation with all categories except ID. However, there is often a high variance in performance.

Improving the focus to cues from low-frequency bands improves adversarial robustness, but hardly offers any guarantees beyond that. For adversarial training, too much low-frequency bias is not desirable.

High-Frequency Bias. A high-frequency bias, on the other hand, shows surprisingly moderate to strong correlations to all aspects of generalization (Fig. 2). For adversarial robustness, it is negatively correlated, but positive for all others. Upon manual inspection of the scatterplots (Fig. 5), we find that this correlation is slightly deceptive. There is a strong monotonic increase in generalization as models start to classify some cues in high frequency, but after some threshold, the gains stagnate or even fall. A common outlier is SHAPENET [17], which shows some ability to recognize high-frequency cues despite being trained on data without discriminative texture cues that typically reside in these high-frequency bands. For adversarial robustness, we completely reject a connection due to the variance for the lowest measurements - but it seems like only models without any high-frequency bias can reach peak adversarial robustness. Unsurprisingly, AT models show next to no high-frequency bias as they are desensitized to this frequency band.

Overall, a slight capability to detect high-frequency cues seems to improve generalization except for adversarial robustness. In this context, it also seems like more high-frequency bias is better than no bias. For adversarial robustness, even the slightest increase in high-frequency bias decreases performance rapidly.

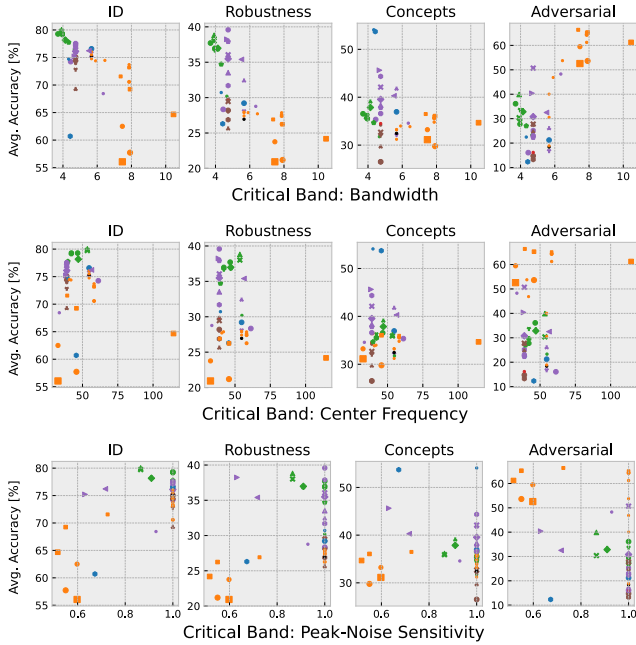


Figure 6. **Critical Band vs. Generalization.** We correlate three properties: the bandwidth (top), center frequency (center), and peak-noise sensitivity (bottom).

5.3. Critical Band

The critical band parameterized by its bandwidth, center frequency, and peak-noise sensitivity has been recently shown to correlate with adversarial robustness (and shape bias) [52]. Akin to what we have observed on other biases, results show different trends between adversarially trained (AT) and other models. For AT, a higher bandwidth of the critical band correlated positively with adversarial robustness, but the correlation was negative for other models. We repeat the measurements using a more careful collection of models and benchmarks. Additionally, we tweak the original test by using normalization, more samples, and more classes. Our normalization allows us to include models trained under ℓ_∞ -norm and larger attack budgets (ϵ) that couldn't be tested with the original methodology and allows us to scale the evaluation against all 1,000 classes. In both cases, unnormalized measurements will break the detection of the critical band as all regions would be classified as critical. Overall, our measurements are more accurate, but may not be directly comparable to the original study. Appendix F contains results using the original methodology (and in combination with normalization).

Bandwidth. The bandwidth not only moderately correlates with adversarial robustness (similarly to the findings in [52]), but also strongly correlates with ID and robustness (Fig. 2). However, there is an inverted trend and the correla-

tion to adversarial robustness is negative. For generalization to concepts we only see a weak positive correlation. Generally, we observe many outliers in the scatterplots (Fig. 6) and often the best performance is not achieved by the models with the lowest bandwidth. Further, some trends only become noticeable or accelerated by adversarially-trained models which accumulate for models with the lowest bandwidth. It remains to be discovered whether low-critical-bandwidth non-AT ResNet-50s exist and how they perform on generalization tests.

Center Frequency. In contrast to the original study, we find *no correlation* between the center frequency and any of our benchmarks (Fig. 6). Even when we single out adversarially trained models, we still observe no statistically significant correlation.

Peak-Noise Sensitivity. We measure a slight statistical correlation among both adversarial training (AT) and non-AT models (Fig. 2). However, the significant concentration (Fig. 6) around the value of 1 indicates an issue in the test and suggests that *this metric might not effectively capture trends in extensive evaluations such as ours*. As for the rest, AT models don't exhibit a distinct clustering pattern, yet display a notably strong positive correlation with both ID and non-adversarial robustness.

For the most part, the critical band seems to poorly explain generalization. However, a lower bandwidth seems to be necessary for generalization except for adversarial robustness where the opposite holds.

6. Other observations

We also notice a couple of other observations that may not directly fit into our storyline but may still be interesting.

Improved low-frequency performance does not guarantee a shape bias. A shape bias causes a model to prefer global over local information as textures are mostly local information found in higher frequency bands. Indeed, we find a moderately strong positive correlation between low-frequency and shape bias ($r = 0.66, p = 2e - 4$), but also clear outliers in both directions: SHAPENET (trained on SIN) [17] achieves a high shape bias of 0.81 but only a low-frequency bias of 0.56; the baseline [19] achieves a low-frequency bias of 0.63 but only a shape bias of 0.21.

Scaling data does not guarantee a shape bias. Previous works have observed a trend where upscaling (pre-)training data results in higher shape bias (and better generalization) [10, 18]. However, when comparing a ResNet-50 trained with self-supervised learning [67] on *YFCC-100M* [56] containing 100 million additional samples, we notice

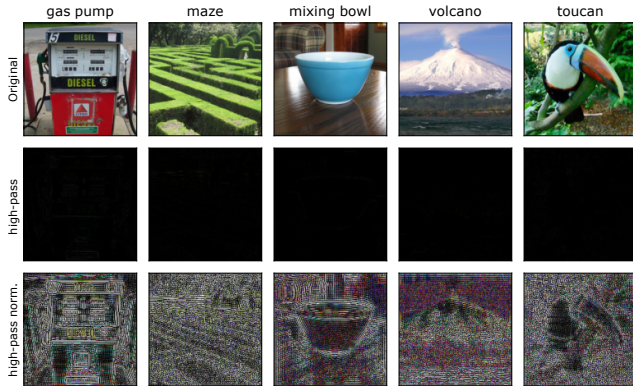


Figure 7. **Examples of high-pass (30%) filtered ImageNet samples.** High-pass filtered samples (center) are correctly classified by DEEPAUGMENT + AUGMIX [23] but not by most other models and unlikely by humans. Only *std* normalized samples (bottom) reveal some similarity to the original samples (top).

that it obtains the same 0.21 shape bias as the baseline despite generally better performance. This shows that the shape bias is not simply determined by the dataset scale alone.

Performance on IN- \bar{C} correlates well with most other benchmarks. Analogous to [55], we notice that ImageNet accuracy strongly correlates with the average ($\bar{r} = 0.79$) of the other benchmarks (excluding SIN and PGD) (Fig. 3) which suggest that ImageNet performance is key to most aspects of generalization. However, we notice that IN- \bar{C} shows an even stronger correlation with the remaining benchmarks ($\bar{r} = 0.83$). Even its weakest correlation (IN-A) is still moderately strong. Yet, just like ImageNet, it does not correlate with PGD and SIN. Future studies may prefer this benchmark to study robustness if they are restricted in their compute budget - IN- \bar{C} is a faster benchmark than IN-C due to fewer corruption classes.

7. Discussion and Conclusion

We have identified numerous shortcomings in previous studies showing correlations between shape bias, spectral biases, and the critical band to generalization. Our holistic and more carefully designed study shows that while correlations between some of these biases and *some* benchmarks exist, most of them fail to correlate to a broader view of generalization.

This leaves us with the question *Is it useful to directly optimize for these biases? I.e., should we use these biases as priors?* Ideally, a prior would *monotonically* and *causally* correlate with all aspects of generalization. However, for the tested biases, we have seen many outliers, non-monotonic trends, and/or only correlations with some as-

pects of generalization. It is important to understand, that in this case the utility of the bias concerning holistic generalization is significantly lowered, but it does not necessarily mean that the bias itself is not useful. Instead, it only tells us that such a bias alone is not sufficient to improve generalization but might be still helpful in combination with other aspects. Beyond that, there might be other valid reasons (*e.g.*, interpretability or alignment) to optimize these biases.

Additionally, we have seen that most biases show strong correlations on adversarially-trained models (*e.g.*, shape bias, low-frequency bias), but not on the entire model zoo. On the one hand, this means that the bias is not useful independent of the training mode, *i.e.*, it does not correlate if we discard this metadata. On the other hand, this might also be a limitation of our study as our AT models are extremely homogenous compared to the other models. All of the models were trained with the same hyperparameters except for the adversary parameters (see [47]). While we cannot reject causal correlations based on our model zoo, we cannot prove them either and look forward to future studies that study adversarial training more closely.

Finally, we have seen some counter-intuitive trends where models misaligned with human perception perform better - *e.g.*, a stronger texture bias increases ID performance and a high-frequency bias seems to improve generalization (except for adversarial robustness) despite containing next to no discriminative cues (see Fig. 7 for examples). For the latter, it is worth noting that even the strongest high-frequency biased model only detects an average of 9% of the samples based on this band alone. As such, we cannot claim that a high-frequency bias is all you need, but drawing some predictive cues from this band seems to be reasonable. In general, we want to issue a word of caution as we are concerned that the reason for improvements we have seen might be due to dataset biases, *e.g.* frequency shortcuts [62] and not an intrinsically better representation. A counter-argument to that theory is that we have seen correlations on many semantically different datasets and it seems unlikely that they all contain the same shortcuts. Yet, we are curious to see if future work can derive a better understanding of generalization and design better benchmarks.

In summary, it seems that the generalization in neural networks - even when fixing the architecture - is too complex to be explained by a single bias. However, certain biases may be necessary (*e.g.*, a low bandwidth of the critical band, or a bit of high-frequency bias) to achieve generalization. It remains to be shown if a single bias or a combination of them exists that can explain generalization better.

Acknowledgements

We thank Robert Geirhos and Ross Wightman for their help regarding model details and Ajay Subramanian for his help in replicating the initial study results.

References

- [1] Antonio A. Abello, Roberto Hirata, and Zhangyang Wang. Dissecting the high-frequency bias in convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2021, virtual, June 19-25, 2021*, pages 863–871. Computer Vision Foundation / IEEE, 2021. [2](#)
- [2] Lucas Beyer, Olivier J. Hénaff, Alexander Kolesnikov, Xiao-hua Zhai, and Aäron van den Oord. Are we done with imagenet?, 2020. [3](#), [1](#)
- [3] Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Šrđić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Machine Learning and Knowledge Discovery in Databases*, pages 387–402, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. [1](#), [2](#)
- [4] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Z. Morley Mao. Adversarial sensor attack on lidar-based perception in autonomous driving. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, page 2267–2281, New York, NY, USA, 2019. Association for Computing Machinery. [1](#)
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*, pages 9912–9924. Curran Associates, Inc., 2020. [4](#), [5](#), [6](#)
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. [2](#), [5](#), [6](#)
- [7] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. In *Advances in Neural Information Processing Systems*, pages 22243–22255. Curran Associates, Inc., 2020. [5](#), [6](#)
- [8] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9640–9649, 2021. [4](#), [5](#), [6](#)
- [9] Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets without pre-training or strong data augmentations. In *International Conference on Learning Representations*, 2022. [2](#)
- [10] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, Rodolphe Jenatton, Lucas Beyer, Michael Tschanen, Anurag Arnab, Xiao Wang, Carlos Riquelme Ruiz, Matthias Minderer, Joan Puigcerver, Utku Evci, Manoj Kumar, Sjoerd Van Steenkiste, Gamaleldin Fathy Elsayed, Aravindh Mahendran, Fisher Yu, Avital Oliver, Fantine Huot, Jasmijn Bastings, Mark Collier, Alexey A. Gritsenko, Vighnesh Birodkar, Cristina Nader Vasconcelos, Yi Tay, Thomas Mensink, Alexander Kolesnikov, Filip Pavetic, Dustin Tran, Thomas Kipf, Mario Lucic, Xiaohua Zhai, Daniel Keysers, Jeremiah J. Harmsen, and Neil Houlsby. Scaling vision transformers to 22 billion parameters. In *Proceedings of the 40th International Conference on Machine Learning*, pages 7480–7512. PMLR, 2023. [2](#), [7](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [1](#), [2](#)
- [12] N. Benjamin Erichson, Soon Hoe Lim, Winnie Xu, Francisco Utrera, Ziang Cao, and Michael W. Mahoney. Noisymix: Boosting model robustness to common corruptions, 2022. [4](#), [5](#), [6](#)
- [13] Samuel G. Finlayson, John D. Bowers, Joichi Ito, Jonathan L. Zittrain, Andrew L. Beam, and Isaac S. Kohane. Adversarial attacks on medical machine learning. *Science*, 363(6433):1287–1289, 2019. [1](#)
- [14] Paul Gavrikov and Janis Keuper. The power of linear combinations: Learning with random convolutions, 2023. [4](#), [5](#), [6](#)
- [15] Paul Gavrikov, Janis Keuper, and Margret Keuper. An extended study of human-like behavior under adversarial training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2360–2367, 2023. [2](#), [3](#), [6](#)
- [16] Paul Gavrikov, Jovita Lukasik, Steffen Jung, Robert Geirhos, Bianca Lamm, Muhammad Jehanzeb Mirza, Margret Keuper, and Janis Keuper. Are vision language models texture or shape biased and can we steer them?, 2024. [2](#)
- [17] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [18] Robert Geirhos, Kantharaju Narayanappa, Benjamin Mitzkus, Tizian Thieringer, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Partial success in closing the gap between human and machine vision. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 2021. [2](#), [3](#), [4](#), [6](#), [7](#)
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015. [1](#), [4](#), [7](#), [5](#), [6](#)
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2022. [2](#)
- [21] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019. [1](#), [2](#), [3](#)

- [22] Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 4, 5, 6
- [23] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8320–8329, 2021. 1, 3, 4, 8, 5, 6
- [24] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3, 5, 1
- [25] Dan Hendrycks, Andy Zou, Mantas Mazeika, Leonard Tang, Bo Li, Dawn Song, and Jacob Steinhardt. Pixmix: Dreamlike pictures comprehensively improve safety measures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16783–16792, 2022. 4, 5, 6
- [26] Katherine Hermann, Ting Chen, and Simon Kornblith. The origins and prevalence of texture bias in convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 19000–19015. Curran Associates, Inc., 2020. 2
- [27] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. 5
- [28] Md Amirul Islam, Matthew Kowal, Patrick Esser, Sen Jia, Björn Ommer, Konstantinos G. Derpanis, and Neil Bruce. Shape or texture: Understanding discriminative features in CNNs. In *International Conference on Learning Representations*, 2021. 2
- [29] Priyank Jaini, Kevin Clark, and Robert Geirhos. Intriguing properties of generative classifiers. In *The Twelfth International Conference on Learning Representations*, 2024. 2, 4, 5, 6
- [30] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2012. 1
- [31] Yingwei Li, Qihang Yu, Mingxing Tan, Jieru Mei, Peng Tang, Wei Shen, Alan Yuille, and cihang xie. Shape-texture debiased neural network training. In *International Conference on Learning Representations*, 2021. 2, 4, 5, 6
- [32] Chang Liu, Yinpeng Dong, Wenzhao Xiang, Xiao Yang, Hang Su, Jun Zhu, Yuefeng Chen, Yuan He, Hui Xue, and Shibao Zheng. A comprehensive study on robustness of image classification models: Benchmarking and rethinking. *arXiv preprint arXiv:2302.14301*, 2023. 5
- [33] Raphael Gontijo Lopes, Dong Yin, Ben Poole, Justin Gilmer, and Ekin D. Cubuk. Improving robustness without sacrificing accuracy with patch gaussian augmentation, 2020. 2
- [34] Jovita Lukasik, Paul Gavrikov, Janis Keuper, and Margret Keuper. Improving native CNN robustness with filter frequency regularization. *Transactions on Machine Learning Research*, 2023. 2
- [35] Aleksander Madry, Aleksandar Makelev, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. 2, 4, 5, 6
- [36] Eric Mintun, Alexander Kirillov, and Saining Xie. On interaction between augmentations and corruptions in natural corruption robustness. In *Advances in Neural Information Processing Systems*, pages 3571–3583. Curran Associates, Inc., 2021. 3, 1
- [37] Apostolos Modas, Rahul Rade, Guillermo Ortiz-Jiménez, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. Prime: A few primitives can boost robustness to common corruptions. In *Computer Vision – ECCV 2022*, pages 623–640, Cham, 2022. Springer Nature Switzerland. 4, 5, 6
- [38] Patrick Müller, Alexander Braun, and Margret Keuper. Classification robustness to common optical aberrations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 3632–3643, 2023. 4, 5, 6
- [39] Chaithanya Kumar Mummadi, Ranjitha Subramaniam, Robin Huttmacher, Julien Vitay, Volker Fischer, and Jan Hendrik Metzen. Does enhanced shape bias improve neural network robustness to common corruptions? In *International Conference on Learning Representations*, 2021. 2, 6
- [40] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. In *Advances in Neural Information Processing Systems*, pages 23296–23308. Curran Associates, Inc., 2021. 2
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 4, 5, 6
- [42] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. Adversarial training can hurt generalization. In *ICML 2019 Workshop on Identifying and Understanding Deep Learning Phenomena*, 2019. 5
- [43] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet classifiers generalize to ImageNet? In *Proceedings of the 36th International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019. 3, 1
- [44] Daniel L. Ruderman. The statistics of natural images. *Network: Computation In Neural Systems*, 5:517–548, 1994. 2
- [45] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 1, 3, 4

- [46] Tonmoy Saikia, Cordelia Schmid, and Thomas Brox. Improving robustness against common corruptions with frequency biased models. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 10191–10200. IEEE, 2021. [2](#)
- [47] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? In *Advances in Neural Information Processing Systems*, pages 3533–3545. Curran Associates, Inc., 2020. [4](#), [8](#), [5](#), [6](#)
- [48] Baifeng Shi, Dinghuai Zhang, Qi Dai, Zhanxing Zhu, Yadong Mu, and Jingdong Wang. Informative dropout for robust representation learning: A shape-bias perspective. In *Proceedings of the 37th International Conference on Machine Learning*, pages 8828–8839. PMLR, 2020. [2](#)
- [49] Nandini C. Singh and Frédéric Theunissen. Modulation spectra of natural sounds and ethological theories of auditory processing. *The Journal of the Acoustical Society of America*, 114:3394–411, 2004. [2](#)
- [50] Andreas Peter Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *Transactions on Machine Learning Research*, 2022. [2](#)
- [51] David Stutz, Matthias Hein, and Bernt Schiele. Disentangling adversarial robustness and generalization. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6969–6980, 2019. [5](#)
- [52] Ajay Subramanian, Elena Sizikova, Najib Majaj, and Denis Pelli. Spatial-frequency channels, shape bias, and adversarial robustness. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [1](#), [2](#), [3](#), [4](#), [7](#), [8](#)
- [53] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. [1](#), [2](#)
- [54] Shiyu Tang, Ruihao Gong, Yan Wang, Aishan Liu, Jiakai Wang, Xinyun Chen, Fengwei Yu, Xianglong Liu, Dawn Song, Alan Yuille, Philip H. S. Torr, and Dacheng Tao. Robustart: Benchmarking robustness on architecture design and training techniques, 2021. [4](#)
- [55] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. In *Advances in Neural Information Processing Systems*, pages 18583–18599. Curran Associates, Inc., 2020. [1](#), [4](#), [8](#)
- [56] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Commun. ACM*, 59(2):64–73, 2016. [7](#)
- [57] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers and distillation through attention. In *Proceedings of the 38th International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. [2](#)
- [58] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2019. [5](#)
- [59] Vasilis Vryniotis. How to Train State-Of-The-Art Models Using TorchVision’s Latest Primitives, 2023. [Online; accessed 15. Nov. 2023]. [4](#), [5](#), [6](#)
- [60] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P. Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pages 10506–10518, 2019. [4](#), [1](#)
- [61] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P. Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [1](#), [2](#), [3](#), [4](#), [6](#)
- [62] Shunxin Wang, Raymond Veldhuis, Christoph Brune, and Nicola Strisciuglio. What do neural networks learn in image classification? a frequency shortcut perspective. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1433–1442, 2023. [8](#), [4](#)
- [63] Yan Wang, Yuhang Li, Ruihao Gong, Tianzi Xiao, and Fengwei Yu. Real world robustness from systematic noise. In *Proceedings of the 1st International Workshop on Adversarial Learning for Multimedia*, page 42–48, New York, NY, USA, 2021. Association for Computing Machinery. [3](#), [5](#)
- [64] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. [4](#), [5](#), [6](#)
- [65] Ross Wightman, Hugo Touvron, and Herve Jegou. Resnet strikes back: An improved training procedure in timm. In *NeurIPS 2021 Workshop on ImageNet: Past, Present, and Future*, 2021. [4](#), [6](#), [5](#)
- [66] Cihang Xie, Mingxing Tan, Boqing Gong, Alan Yuille, and Quoc V. Le. Smooth adversarial training, 2021. [4](#)
- [67] I. Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification, 2019. [7](#)
- [68] Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier perspective on model robustness in computer vision. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. [2](#), [6](#)