

Impact of Aliasing on Generalization in Deep Convolutional Networks

Cristina Vasconcelos^{1†}Hugo Larochelle^{1,2}Vincent Dumoulin¹Rob Romijnders¹Nicolas Le Roux³Ross Goroshin^{1†}¹ Google Research, Brain Team² Mila, Université de Montréal³ Mila, McGill University

Abstract

We investigate the impact of aliasing on generalization in Deep Convolutional Networks and show that data augmentation schemes alone are unable to prevent it due to structural limitations in widely used architectures. Drawing insights from frequency analysis theory, we take a closer look at ResNet and EfficientNet architectures and review the trade-off between aliasing and information loss in each of their major components. We show how to mitigate aliasing by inserting non-trainable low-pass filters at key locations, particularly where networks lack the capacity to learn them. These simple architectural changes lead to substantial improvements in generalization on i.i.d. and even more on out-of-distribution conditions, such as image classification under natural corruptions on ImageNet-C [11] and few-shot learning on Meta-Dataset [26]. State-of-the-art results are achieved on both datasets without introducing additional trainable parameters and using the default hyper-parameters of open source codebases.

1. Introduction

Image analysis in the frequency domain has traditionally played a vital role in computer vision and was even part of the standard pipeline in the early days of deep learning [16]. However, with the advent of large datasets, many practitioners concluded that it was unnecessary to hard-code such priors due to the belief that they can be learned from the data itself.

Deep learning approaches thrive in problem settings where labelled data is abundant and training times are virtually unrestricted – allowing the algorithm to apparently learn all necessary features and priors to achieve impressive performance. Central to the success of deep learning approaches on supervised learning problems is the assumption that the training and test data are sampled from the same distribution (i.i.d. conditions). However, many

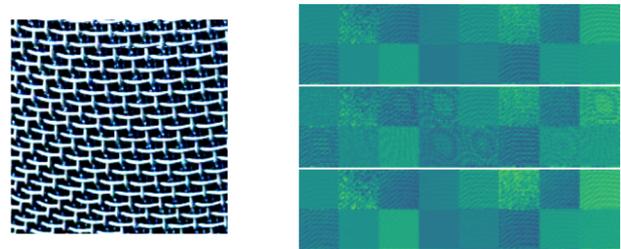


Figure 1: Input image (Left) and feature maps (Right) from a ResNet-50 illustrating features before subsampling (top row), aliasing distortions (middle row) and its prevention using anti-aliasing filters (bottom row).

important problem settings involve out-of-distribution (o.o.d.) data or may restrict the amount of labelled data typically required for training deep networks “from scratch”. In these challenging scenarios, where implicit knowledge cannot be fully obtained from the training datasets, the search for stronger architectural priors as a mechanism to impose explicit knowledge, once again, becomes a critical line of investigation.

We investigate how aliasing impacts robust generalization under shifts in datasets’ spectral distributions and propose simple architectural modifications to mitigate this impact. Our contributions are summarised as follows: (i) A detailed study aimed at more precisely isolating the effects of aliasing in convolutional neural networks (CNNs) is presented. (ii) We use frequency analysis theory to derive the critical network locations where anti-aliasing filters must be added, particularly in the pathways that lack the capacity to learn them. We propose simple architectural modifications to ResNet and EfficientNet based models based on these first principles. (iii) We show that our proposed architecture leads to improved performance on i.i.d. and also on challenging o.o.d. benchmarks, using open source codebases, and their default hyper-parameters. (iv) We show that this architectural improvement complements other techniques for improving generalization, namely data-augmentation and smooth activations. It surpasses the SoTA

† {crisnv, goroshin}@google.com

stand-alone method for ImageNet-C in 9 of 15 categories, while producing the lowest clean error on ImageNet, and achieves SoTA among approaches that learn on all training sources on Meta-Dataset.

2. Spectral Aliasing in Convolutional Networks

Aliasing is a well known phenomenon that may occur when subsampling any signal. It occurs when the sampling rate is too low for the signal’s bandwidth and does not satisfy the Nyquist rate [18]. Any subsampling that violates this condition causes the high frequency components to additively spill into the signal’s low frequency band. This type of distortion can result in noticeable artifacts. In classical image/signal processing, aliasing is prevented by applying a low-pass filter before subsampling. Theoretical fundamentals are presented in Appendix A, where Figure 7 illustrates the frequency leakage caused by aliasing.

In CNNs, any operation that spatially subsamples its input can potentially cause aliasing, if sampling at a rate lower than twice the highest frequency contained in the input (Nyquist rate). Figure 1 illustrates aliasing distortions in CNN’s activation maps. It shows the pre-activations and resulting features from a ResNet-50’s second subsampling layer, with and without anti-aliasing.

We now detail the questions addressed in this work.

- *Are anti-aliasing filters necessary, or do convolutional networks implicitly learn to prevent aliasing?*

It is plausible that anti-aliasing filters can be learned by existing, trainable filters of convolutional networks if provided with sufficient spatial support. In Section 4, we review the relation between the bandwidth of well known low-pass filters and their spatial support size, and establish a relation between CNN filter sizes and their intrinsic capacity to learn low-pass filters directly from data. In Section 4.2 we use these principles to identify “aliasing critical paths” in two modern architectures (ResNet [8, 9] and EfficientNet [25]) and identify their architectural components which are particularly prone to cause aliasing. We point to key representation bottlenecks that lack the minimum filter size to represent low-pass filters that prevent these models from learning to prevent aliasing in an end-to-end fashion. We introduce architectural modifications at these bottlenecks which are able to reduce aliasing through low-pass filtering. To prove the contrapositive, we show that introducing these architectural changes where they are not required in theory actually reduces performance in practice. Finally, we experimentally validate the performance improvement of these architectures via ablation studies.

- *Can we separate anti-aliasing from other confounding effects?*

Introducing anti-aliasing filters can affect both prediction and training. Confounding effects with direct influence on performance include: interaction of the anti-aliasing filter with backpropagation dynamics; potential increase in receptive field size and indirect smoothing of non-linearities. In Section 4.1 we formulate a criterion for optimal placement of low-pass filters and a set of anti-aliasing compositions to isolate the effect of anti-aliasing from these other confounding effects. We perform a number of ablation studies that confirm their differences across a variety of experimental settings.

- *Can standard architectures learn anti-aliasing filters via data-augmentation?*

Although rich data augmentation can provide an incentive to learn anti-aliasing filters, we show that the performance improvements occur mainly in the low-frequency bands. In contrast, the improvements obtained by our anti-aliased model are shown to benefit features across all spectral bands. We also show that when combined, our model boosts the benefits of data-augmentation to be sustained across the spectrum. Those results confirm our hypotheses that data-augmentation alone cannot prevent anti-aliasing without additional architectural modifications. Finally, we show that combining our model with data augmentation consistently leads to the best results in all of our experiments. In Section 5.2 we show that this combination extends the invariances induced by data-augmentation across the entire spectrum.

- *Does frequency aliasing impact the generalization performance of convolutional networks?*

Although input and/or feature map aliasing may occur in deep networks, it is not clear a priori if this phenomenon affects their performance on highly abstract, semantic tasks. Arguably, models can learn to ignore aliased features if in practice they do not correlate with training labels. Contradicting this argument, however, we investigate the impact of aliasing on generalization under two hypotheses. First, we claim that a model whose features are susceptible to aliasing may learn brittle correlations that rely on the presence (or absence) of aliasing artifacts in order to generalize. Under this assumption, aliasing impacts out of distribution generalization in settings where test images have different spectral properties, for example, images of varying spatial resolution, compression format, or those affected by natural corruptions. The second, equally important hypothesis, is that aliasing may prevent models from learning useful correlations between features that were corrupted during subsampling. Under this second assumption, we claim that because aliasing can leak frequencies across the entire spectrum, an anti-aliased model has the potential to improve the network’s ability to learn useful features

that use the entire spectrum. In Section 5.2 we analyse the performance gains across different spectral bands, when our anti-aliased model is compared to a baseline, and show that it boosts data-augmentation results, as claimed by our second hypothesis. Finally, our ablation studies (Section 5) evaluate the impact of our proposed model across a large number of datasets (ImageNet, ImageNet-C, ImageNet-R, ImageNet-Vid, ImageNet-V2, ObjectNet, Stylized-ImageNet and another 13 datasets on few-shot-learning) to demonstrate the impact on generalization across a diverse and extremely challenging set of tasks.

3. Related Work

Most relevant to our work is the reformulation of a sub-sampling layer proposed by Zhang [30]. They suggest the use of a low-pass filter after a strided-layer’s non-linearity, that is a strided-convolution followed by ReLU activation is redefined as non-strided convolution followed by ReLU activation, followed by a strided low-pass filter (Figure 2). We point out that this formulation does not fully isolate the aliasing problem from the non-linearity smoothing side-effects, as the filtering operation also smooths the output of the non-linearity. In contrast, our formulation introduces the filters at the location predicted by aliasing theory. Our model delays the removal of high frequencies emerging from non-linear operations until the subsequent sub-sampling operation. Consequently, our formulation maintains high-frequencies throughout the trainable filters that may exist in between the non-linearity and the next sub-sampling, but filters them only at the exact point where they may cause aliasing distortion. Zhang’s formulation also ignores the network’s existing capacity to learn anti-aliasing filters, and because of that, makes use of low-pass filters in layers with sufficient spatial support for learning such filters implicitly end-to-end. A direct comparison to Zhang’s formulation (Appendix E) shows that our model surpasses their results in both i.i.d. and o.o.d. conditions using fewer, and smaller low pass-filters.

Zou et al. [31] propose the use of trainable low-pass filtering layers that operate on feature channel groups and adapt to spatial locations. Their anti-aliasing module augments the model with new trainable and non-linear components, increasing the capacity of the network, thus making it unsuitable for isolating the effects of aliasing from other confounding effects, by the addition of extra convolutional, batch norm, and softmax layers.

Azulay and Weiss [2] show that CNNs are not as robust to small image transformations as commonly assumed. The paper points out that CNN-based models typically ignore the sampling theorem and show large changes in a prediction under small, mostly imperceptible, perturbations of the input. They observe that the improvement in generalization obtained by data augmentation is limited to images that

are similar to those seen during training. Complementary to their findings, we measure the impact of data augmentation in different spectral bands and show that improvements in robustness are concentrated in the low frequencies, but at the cost of reducing the model’s robustness to changes in mid and high frequency bands. We also show that our model is able to boost the gain associated with data-augmentation to extend its beneficial effects across the entire spectrum.

Sophisticated augmentation strategies are currently the state-of-the-art approach to o.o.d. classification under natural corruptions [4, 11, 12, 23]. Rusak et al. [23] currently lead the “standalone leaderboard” on ImageNet-C with a two pass approach. First, they train a generative model to produce additive noise. Next, the classifier and the generative network are jointly trained in an adversarial fashion. Recently, [29] pointed out that their improvements are mainly on corruptions that affect high frequencies, while reducing robustness to corruptions that affect low frequencies and also in detriment of “clean” (uncorrupted) test accuracy.

While [29] pointed out that robustness gains are typically non-uniform across corruption types and that increasing performance in the presence of random noise is often met with reduced performance in corruptions concentrated in different bands, our method overcomes this trade-off (subsection 5.3). *Note that severe aliasing corrupts the entire spectrum, and not only high frequencies.*

Our results, obtained by combining architectural modifications with off-the-shelf data augmentation (Ekin et al. [5]) show improvements in all of the 15 corruption categories and outperforms [23] in 9 of them (including fog and contrast, that are concentrated in low frequencies according to [29]) while obtaining the highest clean accuracy of the benchmark: 78.8% for ImageNet versus 76.1% for [23]. At time of writing, we are not aware of better performing standalone methods on ImageNet-C using a ResNet-50 trained on 224×224 ImageNet examples only.

In contrast to the above methods that settle for a trade-off between clean and corrupted accuracy, our method improves both i.i.d. and o.o.d. accuracy simultaneously.

4. Methods

This section briefly reviews relevant frequency analysis theory, and outlines its implications on CNN architectures. Additional details can be found in Appendix A.

An ideal low-pass filter completely eliminates all frequencies above the cutoff frequency, while allowing those below it to pass without attenuation. It is represented as a rectangular function in the frequency domain and corresponds to the “sinc” function in the discrete spatial-domain with support size equal to the number of elements on the input itself. The spatial support size is defined by the interval with minimum length containing all nonzero filter weights. Learning a good approximation of the ideal low-pass filter

requires much larger filters than those typically found in the individual layers of modern CNNs.

A non-ideal low-pass filter has transition regions, defined as intervals between its passband and stopband responses in the frequency domain. Although they may have much smaller spatial support size, the existence of transition regions imposes a trade-off between preserving information and reducing aliasing.

Any convolutional layer that represents a low-pass filter must have a filter of size of at least 2. A discrete filter with size 1 can only represent an impulse function. Its Fourier transform is constant across the spectrum, which scales all frequencies equally, thus it cannot act as a low pass filter. Convolutional layers with filter size larger than 1 are able to represent low-pass filters, up to varying degrees of approximation error for different bandwidths. However, the uncertainty principle implies that a signal’s spatial support size is inversely proportional to its spectral support size [19]. For this reason, the larger the filter size of a convolutional filter, the better it can approximate a narrow low-pass filter.

Another important property pertains to a stack of convolutional layers which may learn to prevent aliasing in a distributed fashion. More formally, a filter is said to be separable if it can be expressed as the composition of two or more filters. A stack of convolutional layers with small filter sizes (but larger than 1) may learn to represent a larger, separable, low-pass filter.

4.1. Optimal Placement and Confounding effects

The design principle of our anti-aliasing models trade-off between two opposing requirements: band-limiting the spectrum of subsampled feature maps in order to prevent aliasing in layers that subsample, while simultaneously preserving the information encoded in high frequencies as much as possible throughout the other layers.

One naive extreme of this balance is to pre-process the input image with a low-pass filter that has very low cut-off frequency. This solution is sub-optimal for two reasons: (i) it induces early loss of information that may be crucial for solving the task (e.g. classification); (ii) non-linearities may introduce high-frequencies in the internal feature maps, causing them to be susceptible to aliasing in subsequent subsampling layers.

We claim that the solution proposed in [30] is sub-optimal. It suggests the use of a low-pass filter just after a subsampling layer non-linearity. We observe that the high frequencies created by this non-linearity are harmless up to the next subsampling operator, and thus this solution prematurely discards high frequency information. To satisfy the two requirements mentioned above, an optimal anti-aliased model should be able to preserve high-frequencies up to the subsequent subsampling operation. To allow the network’s features to make use of high frequency information while

avoiding aliasing, *it is ideal to apply low-pass filtering immediately before subsampling.*

Optimizing these criteria does not uniquely specify an anti-aliasing architecture. All four variations in Figure 2 satisfy the stated optimality criteria. Nevertheless, each variation introduces its own side-effects on training, both in the forward and backward pass. Although convolutions *per se* are commutative (see [14] for a detailed proof), note that this property no longer holds when back-propagating gradients or when the layers include subsampling/non-linearities. That is, a change in the order of the fixed and trainable filters induces different priors. For example, the *backward* pass in the *post-filter* variant induces the trainable convolution filter to be influenced by a larger neighborhood of gradients than the *pre-filter* variant. Similarly, the decision of where to include subsampling affects the *resolution* of the feature maps and also the gradients that influence the trainable layer. In our variant named *enlarged receptive field (e.r.f.)*, we explore the case where subsampling is part of the low-pass filter that precedes the trainable convolution, which potentially increases the receptive field size of the trainable filter as a side effect.

Among the four variations, the *post-filter* (outlined in red in Figure 2) has the benefit of smoothing gradients with the low-pass filter in the backward pass, in a way that better aligns the low-pass filter with the positions where gradients are up-sampled. This potentially explains why this variant produces superior empirical results, as confirmed in Table 1.

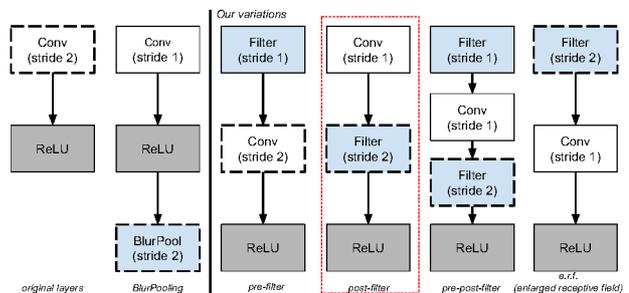


Figure 2: Low-pass filters can be added at various locations with potentially different side-effects. *BlurPooling* [30] suggest their addition after the ReLU non-linearity, whereas our variants apply low-pass filtering before the original subsampling and preserve any high frequencies produced by the non-linearity. The optimal “post-filter” variation is outlined in red.

4.2. Aliasing critical paths

In this section we characterize a network’s capacity to learn anti-aliasing filters end-to-end. Although low-pass filtering can occur in any of the trainable convolution layers, in order for low-pass filters to prevent aliasing they must

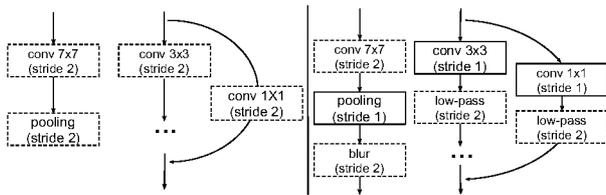


Figure 3: ResNets components that include subsampling. Strided-skip connections filters (1×1) lack the minimum size necessary to represent a low pass filter. Left: Original composition. Right: Our anti-aliased model based on applying the *post-filter* composition only at key places. Non-linearities are omitted for clarity. Dashed lines denote layers with subsampling.

be located between the operation that produces high frequencies and the subsampling operation. For this reason, we define the aliasing critical path as the set of operations preceding subsampling that maintain or reduce the range of frequencies contained in the signal and lack the capacity to produce new high frequency content. The burden of preventing aliasing, in accordance with the sampling theorem, falls on the operations found in the critical path. Given a sequence of layers, the aliasing critical path for a certain subsampling operation traces back to the last non-linearity that precedes it (or the input image itself if none exist). Non-linearities may produce high-frequencies (see appendix B), but these are harmless up until the next down-sampling operation. Skip connections are defined as the summation of a feature map with a downstream feature map. They can act as pathways for replicating frequencies through the network, bypassing any low-pass filters that may have been learned in the main path¹.

In summary, we define the critical path for learning low-pass filters as the sequence of operations between a subsampling operation and either the skip connection or a non-linearity preceding it, whichever is closest to the subsampling operation. Note that our definition of critical path for learning anti-aliasing does not preclude trainable layers further upstream (i.e. preceding the skip-connection or non-linearity) from serving as anti-aliasing filters by aggressively low-pass filtering their input. Rather, we define a critical set of layers that must act as “last chance” low-pass filters, whenever the incoming signal does not satisfy the Nyquist rate.

4.3. ResNets and EfficientNets

ResNet’s aliasing critical paths can be differentiated according to their composition into the following: (1) the ini-

¹In theory, the summation at the skip connection can act as a low-pass filter through phase cancellation. Phase cancellation occurs when two (or more) signals of the same frequency but with inverted phases (a phase difference of π) cancel each other when summed together.

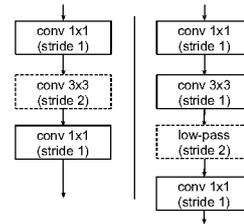


Figure 4: EfficientNet layers that include subsampling. Left: Original composition. Right: Our anti-aliased model which uses the *post-filter* variation at key locations. Non-linearities are omitted for clarity. Dashed lines indicate layers with subsampling.

tial strided-convolutional layer; (2) the initial strided-max-pooling (absent in some implementations such as [6]); (3) the sequence of convolutions within a strided-residual block main path up to its subsampling operation; and (4) the sequence of convolutions within a strided-skip-connection up to its spatial subsampling. Figure 3 illustrates the components in these critical paths and their filter sizes which impact their capacity to prevent aliasing. The right-hand side of the figure illustrates our proposed model, respecting the optimal placement criteria (Section 4.1). Note that the 7×7 trainable convolution has sufficient capacity to act as an anti-aliasing filter, therefore our model does not include an additional anti-aliasing filter in this critical path.

EfficientNet is currently the best-performing architecture in several resource constrained settings [25]. Its baseline architecture was found via neural architecture search, which optimizes both accuracy and efficiency (FLOPS), while balancing the network’s depth, width, and input and feature map resolution as hyper-parameters. It is interesting to analyse the EfficientNet architecture from the perspective of aliasing because the hyper-parameters that were optimized are directly related to the aliasing problem.

EfficientNets’s aliasing critical paths can be differentiated according to their composition into the following (Figure 4): (1) initial strided-convolutional layer; (2) strided-blocks with a sequence of 1×1 and 3×3 convolutions. Layers that perform subsampling do not contain a parallel path with skip connections, while other blocks do. EfficientNets use the Swish activation function [20] instead of ReLU. Smooth functions have rapidly decaying Fourier transforms [15, 3]. This property reduces the likelihood, assuming random inputs, of introducing high frequency content as compared to non-smooth activation functions (see Appendix B).

5. Experimental Results

Our experiments investigate different aspects of the impact of aliasing on performance. We begin by evaluating the i.i.d. generalization performance of our proposed mod-

Location	Pre-filter (filter size k)			Post-filter (filter size k)			Pre-post-filter (filter size k)			Enlarge receptive field (filter size k)		
	$k = 3$	$k = 5$	$k = 7$	$k = 3$	$k = 5$	$k = 7$	$k = 3$	$k = 5$	$k = 7$	$k = 3$	$k = 5$	$k = 7$
all layers	74.20	71.06	68.74	73.91	70.71	68.00	70.92	65.80	61.52	73.06	71.09	68.67
conv1	76.18	75.81	75.45	76.10	75.72	75.47	75.76	75.07	74.63	75.27	75.12	75.03
max pool	76.00	75.75	75.40	76.56	76.43	76.35	76.03	75.54	75.24	74.88	74.87	74.60
block-conv	76.87	76.74	76.83	76.88	76.83	76.74	76.87	76.73	76.63	76.64	76.71	76.61
strided-skip	77.05	77.14	77.07	77.15	77.12	76.83	77.02	76.90	76.62	-	-	-

Table 1: Imagenet results: rows demonstrate the impact of anti-aliasing the model’s components from Figure 3 individually while columns show blur variations from Figure 2. Results show a significant accuracy increase when anti-aliasing the strided-skip connections of a Resnet-50 and the negative impact of information loss when blur is applied to all layers indistinctly and also to the first convolutional layer that contain large trainable kernels, capable of learning anti-aliasing filters. Note that baseline accuracy is 76.49%. The values correspond to the mean accuracy over 3 runs with different seeds.

els on the ImageNet benchmark in Sections 5.1-5.2. The experiments in subsection 5.1 are designed to validate the hypotheses presented in section 4 such as the optimal placement of low-pass filters. These experiments are also designed to disambiguate the improvements obtained by anti-aliasing from other confounding effects.

Next, the experiments in subsection 5.2 show that anti-aliasing improves the results obtained by data-augmentation by extending its effects across a larger range of frequencies.

Finally, we evaluate our models on challenging o.o.d. benchmarks and show that our proposed modifications lead to even more striking performance improvement – achieving state-of-the-art performance without any additional hyper-parameter sweeps. The detailed description of these datasets is presented in the appendices (see Appendix D for ImageNet-C, Appendix F for Meta-Dataset, and Appendix E for extra results on other datasets used in robustness analysis). Our experiments use TensorFlow’s official ResNet-50 and EfficientNet-B0 models². The o.o.d. experiments on the Meta-Dataset benchmark use the publicly available SUR codebase [6] – a recent few-shot classification model. On the following tables, results on bold implies “better” with statistical significance.

5.1. ImageNet

To investigate the effects of aliasing in the ResNet family of architectures we initially evaluate their performance on ImageNet [24]. Our results show that ResNets are most severely impacted by aliasing in their strided skip connections (which are preceded by 1×1 convolutions).

We used the official TensorFlow [1] public code for training a ResNet-50 architecture, yielding a top-1 accuracy of 76.49%. The codebase reproduces the training pipeline and hyper-parameters from [7], in which models are trained for 90 epochs. Table 1 shows the effect on top-1 accuracy of adding anti-aliasing filters before (*pre*), after (*post*), and before-and-after (*pre-post*), various operations in the net-

²https://github.com/tensorflow/models/tree/master/official/vision/image_classification/resnet

work. Recall that adding these fixed filters may also affect the back-propagation of gradients, as well as increase the receptive field size. Our ablation studies are designed to disambiguate these effects from anti-aliasing. Table 1 reports accuracies as a result of adding a *pre-filter* (77.14%) or *post-filter* (77.15%) to the skip connections that include *subsampling*, and are immediately preceded by a trainable 1×1 convolutional layer. Note that inserting low-pass filters into all convolutional layers (“all” row) degrades the performance to as low as 61.00% accuracy. The second row shows that anti-aliasing the “conv1” layer degrades performance because it already has the capacity to learn a low-pass filter in its 7×7 kernel. The last column, ERF (enlarged receptive field), is aimed to disambiguate whether the improved performance is truly due to anti-aliasing or merely the enlarged receptive field caused as a side-effect.

Finally, by combining anti-aliasing with *post-filter* variants at all strided layers except the first strided-convolution which has large filter size (Figure 3), we are able to report a top-1 accuracy of (77.47%). These results support our hypothesis that anti-aliasing in networks must trade-off between mitigating the effects of aliasing while preserving high-frequency information as much as possible. This is achieved by adding anti-aliasing filters only to those critical paths that lack the capacity to learn them.

Finally we note that analogous improvements were also obtained with EfficientNets using the same design principles. These results can be found in Appendix C.

5.2. Data-Augmentation

The ablation studies presented in this section are designed to differentiate between performance gains obtained by data-augmentation from those obtained from our anti-aliased model. The four models evaluated in this section and their respective accuracies on ImageNet are: (i) “Baseline”-baseline model (77.36%), (ii) “Anti-Aliasing”-our proposed model containing non-trainable anti-aliasing filters (77.76%), (iii) “Rand Augmentation”- model trained with random data augmentation [5] (77.32%), (iv) “Rand

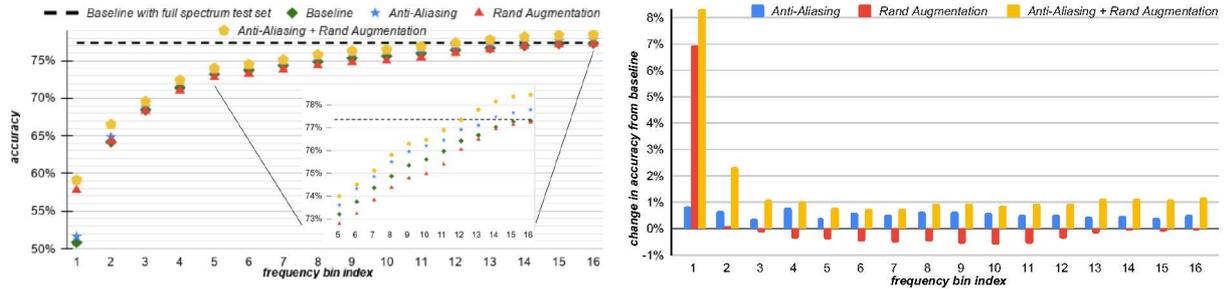


Figure 5: The impact on pre-trained models when tested on images that have 1/16 of their spectral band removed. *Left*: Degradation Curve. All models are more robust to the lack of high-frequency than low frequency content. *Right*: Difference between the baseline performance and the remaining three models. Our anti-aliased model performance is higher than baseline in all spectral bands, presenting positive differences across the spectrum. Data augmentation model performance is higher when removing content from the lowest frequencies, but presents the worst performance of all models for the remaining bins. The combined model (AA+DA) boosts their individual advantages.



Figure 6: Random subset of ImageNet wrongly classified by both baseline and data-augmented models, but correctly classified by our Anti-Aliased model. The images have different spectral characteristics, from those that include high frequencies (fine structures and texture) to those with low frequency content (regions of constant color).

Augmentation + Anti-Aliasing” combines (ii) and (iii) (78.45%). These models were trained with the same number of epochs (180) on ImageNet only, as the public source code from [5] (when data-augmentation is enabled).

Figure 6 illustrates images correctly classified by our anti-aliased model, but not by the baseline nor the model trained with data-augmentation. Note that these images contain different spectral characteristics suggesting that anti-aliased models benefit features across the entire spectrum. To further explore this observation we investigate the reliance of trained networks on various frequency bands of input images. This can be done by measuring the impact of removing frequency intervals from the input images at inference time only. To accomplish this, we divide the input spectrum into 16 bands and apply a “notch filter”, zeroing out one frequency interval at a time. We then evaluate the performance of the same four models (trained on

unfiltered images) on the resulting filtered image test sets. The results are presented in Figure 5. The left plot shows that for all models, the performance impact is larger when lower frequency bands are filtered out. The right plot shows the change in accuracy from the baseline accuracy over the same bin. Notice a striking characteristic pattern of decreased performance around the mid-band frequencies of the “Rand Augmentation” model, if the mid-range band is missing it performs significantly worse than the baseline. This relative decrease in performance is in agreement with the work of Hermann et al. [13], who noted that models trained with data augmentation tend to rely on image textures – which occupy the mid-range frequency bands.

On the other hand, our anti-aliased model obtained a consistent improvement over the baseline model in all of the 16 bins. This effect is further boosted when combined with data augmentation. Similar results were also observed with 8 and 32 bins. Appendix B extends these experiments to our models using smooth activation functions. We show that their benefit is complementary and combining them with anti-aliasing produces the best results.

5.3. ImageNet-C

Table 2 presents a comparison of our models on the ImageNet-C leaderboard. It uses the corruption error (CE) measurement proposed by [11] that is defined as $CE_c^f = \frac{\sum_{s=1}^5 E_{s,c}^f}{\sum_{s=1}^5 E_{s,c}^{AlexNet}}$ where $E_{s,c}^f$ is the top-1 error of a classifier f for a corruption c with severity s . The mean Corruption Error (mCE) is taken by averaging over all the 15 corruptions. The ‘Clean err.’ column shows the classifier’s top-1 error on the original ImageNet, included in the table for contrasting gains on clean and corrupted test sets.

Our anti-aliased model, described in subsection 5.1, achieves mCE of 70.0% on ImageNet-C and 22.5% “Clean err.” on ImageNet, which represent a significant improve-

Method	Noise			Blur				Weather				Digital			mCE Clean err.		
	Gauss.	Shot	Impulse	Defocus	Glass	Motion	Zoom	Snow	Frost	Fog	Bright	Contrast	Elastic	Pixel	JPEG		
ResNet-50 Published [11]	80	82	83	75	89	78	80	78	75	66	57	71	85	77	77	76.7	23.9
BlurPool [30]	73	74	76	74	86	78	77	77	72	63	56	68	86	71	71	73.4	23.0
Ours (90 epochs)	68	70	70	72	85	82	75	72	62	50	52	66	81	81	66	70.0	22.5
AutoAugment**[4]	69	68	72	77	83	80	81	79	75	64	56	70	88	57	71	72.7	22.8
Rand AutoAugm.**[12]	70	71	72	80	86	82	81	81	77	72	61	75	88	73	72	76.1	23.6
AUGMIX [12]	65	66	67	70	80	66	66	75	72	67	58	58	79	69	69	68.4	22.4
ANT (3x3) [23]	39	40	39	68	78	73	77	71	66	68	55	69	79	63	64	63.0	23.9
RandAugment** [5]	60	58	60	70	90	76	80	70	67	44	50	57	80	86	64	67.4	22.8
Ours + RandAugm.**	59	58	61	70	84	75	76	69	65	41	48	55	80	82	61	65.5	21.6
Swish	71	72	74	69	88	80	76	74	69	51	54	68	81	80	67	71.6	22.4
Swish + Rand Augm.**	61	61	62	69	88	73	78	69	67	42	49	55	81	87	63	66.9	21.9
Ours + Swish + Rand Augm.	60	59	61	69	85	71	75	68	64	41	47	55	78	81	61	64.9	21.2

Table 2: Corruption Error (CE), mCE, and Clean Error values when including our anti-aliasing variations on top of ResNet-50. Adding anti-aliasing leads to a lower error (mCE) than all existing models with the exception of ANT. ANT uses adversarial training and has an extra generative network, is significantly more expensive to train, has a higher clean error and only performs better than ours in 6 of the 15 corruptions. Our model uses fewer and smaller filters than [30] but at precise locations. Models with data augmentation were trained for longer. For all columns, lower is better.

Method	Average rank
RelationNet [26]	9.80
k-NN [26]	8.95
MatchingNet [26]	8.60
fo-MAML [26]	8.15
Finetune [26]	7.10
ProtoNet [26]	6.70
fo-Proto-MAML [26]	4.80
CNAPs [22]	4.45
SUR [6]	3.40
URT [17]	2.40
SUR + Anti-aliased + GELU (ours)	1.65

Table 3: Average rank over all Meta-Dataset test sources for approaches that learn on all Meta-Dataset training sources. We recompute ranks after including our proposed approach into Meta-Dataset’s public leaderboard.

ment over the results reported in [30]. Also note that our model uses fewer and smaller filters than [30] but at key positions discussed in section 4. Zhang’s best model uses 7x7 filters at every sub-sampling operation. Appendix E presents a direct comparison to [30].

Next we show that performance can be further improved by the use of data-augmentation and smooth activation functions (Swish [20]). These models were trained for longer (180 epochs) in order to replicate training conditions of the remaining references mentioned in the table. Similar to [12], our data-augmented models do not use augmentations such as contrast, color, brightness, sharpness, as they overlap with the ImageNet-C test set corruptions. We achieve a “Clean” top-1 error of 21.2% and an mCE of 64.9%. Note that unlike [23] we demonstrate that mCE can be improved *without sacrificing* clean error.

5.4. Meta-Dataset with SUR

Meta-Dataset is a challenging few-shot image classification benchmark which samples heterogeneous learning episodes from a diverse collection of datasets, two of which (Traffic Signs, MSCOCO) are used exclusively for evaluation and are therefore considered out-of-domain. We apply anti-aliasing and smooth-activation (GELU) to SUR [6], which is a competitive few-shot classification approach that trains one feature extractor per training domain and combines their representations during inference. We obtain SoTA average rank across all test datasets among approaches which learn from all available training classes (Table 3). Additional experimental details and accuracy breakdowns are presented in Appendix F.

6. Conclusion

Drawing from the classical sampling theorem from signal processing, we proposed simple architectural improvements to convolutional architectures to counter aliasing occurring at various stages. These changes lead to substantial performance gains on both i.i.d. and o.o.d. generalization, and were shown to boost the impact of data augmentation and smooth activation functions by extending their effect across the spectrum. Compared to other performance enhancement techniques, anti-aliasing is simple to implement, computationally inexpensive, and does not require additional trainable parameters. In all our experiments, we could not find a setting where it degraded the performance which leads us to recommend their use as a standard component of convolutional architectures.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. **6**
- [2] Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *Journal Machine Learning Research*, 2019. **3**
- [3] Guy Bresler and Dheeraj Nagaraj. A corrective view of neural networks: Representation, memorization and learning, 2020. **5**
- [4] Ekin Dogus Cubuk, Barret Zoph, Dandelion Mané, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation policies from data. *arXiv:1805.09501*, 2018. **3, 8**
- [5] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data augmentation with a reduced search space, 2019. **3, 6, 7, 8, 14**
- [6] Nikita Dvornik, Cordelia Schmid, and Julien Mairal. Selecting relevant features from a universal representation for few-shot classification. *arXiv:2003.09338*, 2020. **5, 6, 8, 16**
- [7] Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: training imagenet in 1 hour. *arXiv:1706.02677*, 2017. **6**
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, 2016. **2**
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. *arXiv:1603.05027*, 2016. **2**
- [10] 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. *arXiv:2006.16241*, 2020. **15**
- [11] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *International Conference on Learning Representations*, 2019. **1, 3, 7, 8, 13, 14**
- [12] 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. *International Conference on Learning Representations*, 2020. **3, 8, 14**
- [13] Katherine L Hermann, Ting Chen, and Simon Kornblith. The origins and prevalence of texture bias in convolutional neural networks. *arXiv:1911.09071*, 2019. **7**
- [14] Julius O. Smith III. *Mathematics of the Discrete Fourier Transform (DFT): with Audio Applications (2nd edition)*. W3K Publishing, 2007. **4**
- [15] Yitzhak Katznelson. *An Introduction to Harmonic Analysis*. Cambridge Mathematical Library. Cambridge University Press, 3 edition, 2004. **5, 12**
- [16] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Master’s thesis, Department of Computer Science, University of Toronto, 2009. **1**
- [17] Lu Liu, William Hamilton, Guodong Long, Jing Jiang, and Hugo Larochelle. A universal representation transformer layer for few-shot image classification. In *International Conference on Learning Representations*, 2021. **8**
- [18] A.V. Oppenheim, A.S. Willsky, S.H. Nawab, W. Hamid, and I.T. Young. *Signals & Systems*. Signal Processing Series. Prentice-Hall, 1997. **2, 11**
- [19] Reza Parhizkar, Yann Barbotin, and Martin Vetterli. Sequences with minimal time–frequency uncertainty. *Applied and Computational Harmonic Analysis*, 2015. **4**
- [20] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions, 2017. **5, 8**
- [21] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, 2019. **15**
- [22] James Requeima, Jonathan Gordon, John Bronskill, Sebastian Nowozin, and Richard E Turner. Fast and flexible multi-task classification using conditional neural adaptive processes. In *Advances in Neural Information Processing Systems*, 2019. **8**
- [23] Evgenia Rusak, Lukas Schott, Roland S. Zimmermann, Julian Bitterwolf, Oliver Bringmann, Matthias Bethge, and Wieland Brendel. A simple way to make neural networks robust against diverse image corruptions, 2020. **3, 8, 14**
- [24] 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*, 2015. **6**
- [25] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 2019. **2, 5**
- [26] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples, 2020. **1, 8, 15, 16**
- [27] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016. **16**
- [28] Cihang Xie, Mingxing Tan, Boqing Gong, Alan Yuille, and Quoc V Le. Smooth adversarial training. *arXiv:2006.14536*, 2020. **12**

- [29] 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*, 2019. [3](#)
- [30] Richard Zhang. Making convolutional networks shift-invariant again. In *International Conference on Machine Learning*, 2019. [3](#), [4](#), [8](#), [14](#), [15](#)
- [31] Xueyan Zou, Fanyi Xiao, Zhiding Yu, and Yong Jae Lee. Delving deeper into anti-aliasing in convnets. In *British Machine Vision Conference*, 2020. [3](#)