

## University of Groningen

### PushPull-Net

Bennabhaktula, Guru Swaroop; Alegre, Enrique; Strisciuglio, Nicola; Azzopardi, George

DOI:  
[10.48550/arXiv.2408.04077](https://doi.org/10.48550/arXiv.2408.04077)

**IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.**

*Document Version*  
Early version, also known as pre-print

*Publication date:*  
2024

[Link to publication in University of Groningen/UMCG research database](#)

*Citation for published version (APA):*  
Bennabhaktula, G. S., Alegre, E., Strisciuglio, N., & Azzopardi, G. (2024). *PushPull-Net: Inhibition-driven ResNet robust to image corruptions*. arXiv. <https://doi.org/10.48550/arXiv.2408.04077>

#### Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

#### Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

*Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.*

# PushPull-Net: Inhibition-driven ResNet robust to image corruptions

Guru Swaroop Bennabhaktula<sup>1</sup>[0000-0002-8434-9271], Enrique Alegre<sup>2</sup>[0000-0003-2081-774X], Nicola Strisciuglio<sup>3</sup>[0000-0002-7478-3509], and George Azzopardi<sup>1</sup>[0000-0001-6552-2596]

<sup>1</sup> University of Groningen, the Netherlands

{g.s.bennabhaktula,g.azzopardi}@rug.nl

<sup>2</sup> University of León, Spain enrique.alegre@unileon.es

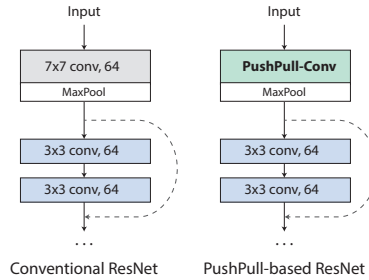
<sup>3</sup> University of Twente, the Netherlands n.strisciuglio@utwente.nl

**Abstract.** We introduce a novel computational unit, termed PushPull-Conv, in the first layer of a ResNet architecture, inspired by the anti-phase inhibition phenomenon observed in the primary visual cortex. This unit redefines the traditional convolutional layer by implementing a pair of complementary filters: a trainable push kernel and its counterpart, the pull kernel. The push kernel (analogous to traditional convolution) learns to respond to specific stimuli, while the pull kernel reacts to the same stimuli but of opposite contrast. This configuration enhances stimulus selectivity and effectively inhibits response in regions lacking preferred stimuli. This effect is attributed to the push and pull kernels, which produce responses of comparable magnitude in such regions, thereby neutralizing each other. The incorporation of the PushPull-Conv into ResNets significantly increases their robustness to image corruption. Our experiments with benchmark corruption datasets show that the PushPull-Conv can be combined with other data augmentation techniques to further improve model robustness. We set a new robustness benchmark on ResNet50 achieving an *mCE* of 49.95% on IMAGENET-C when combining PRIME augmentation with PushPull inhibition.

**Keywords:** Robustness · PushPull Inhibition · ResNet.

## 1 Introduction

One of the key challenges facing ConvNets is their ability to generalize to out-of-distribution images [39]. This distribution shift may manifest itself in several forms, such as image corruption [11], stylized image transformations [17], rendition [10], and adversarial samples [29], among others. Image corruption is frequently encountered and can happen without any deliberate interference or adversary involvement. Corruptions can be of several types. This work focuses on improving the robustness of ResNet to common image corruptions such as noise, blur, weather, and digital.



**Fig. 1.** The proposed approach involves substituting the first convolutional layer with PushPull-Conv, enhancing robustness with minimal computational overhead.

Among a few strategies to tackle this problem [40], a couple of popular ones include (a) data augmentation in the training phase and (b) by design, where intrinsic changes are made to the model’s architecture. Regularization strategies like dropout, batch normalization, and weight decay enhance model generalization and reduce overfitting but help little to improve model robustness against image corruption.

Data augmentation has considerably improved model robustness [11,10,41,36]. However, mere data augmentation cannot address this problem sufficiently [41]. Improvement in the architectural design of ConvNets along with data augmentation strategies have improved model robustness [19]. We contribute to enhancing ConvNet architectural design for improved model robustness. While [19] deals with inhibiting the most significant activations, the proposed push-pull inhibition takes a contrasting approach and reduces signal responses in regions lacking preferred stimuli.

ConvNets trained on IMAGENET [27] are biased towards classifying texture over an object’s shape [8]. This is in contrast to the behaviour of humans, and Geirhos et al. [8] showed that introducing shape bias improves the model’s robustness and accuracy on IMAGENET. Unlike [8], we introduce such an ability using intrinsic changes to the ConvNet rather than by stylized data transforms. Our methodology is biologically inspired by the work of Hubel and Wiesel [15,16] on simple cells in the primary visual cortex of mammals and by the antiphase or push-pull inhibition that they exhibit [2]. We propose a novel inhibition-driven computational unit inspired by the behaviour of these simple cells, employing a push-pull based inhibition mechanism as detailed below.

The proposed computational unit, PushPull-Conv, consists of excitatory and inhibitory convolutions, in other words, push and pull convolutions. The pull kernel is the complement of the push kernel. The PushPull-Conv unit, with its simultaneous use of push (excitatory) and pull (inhibitory) filters, offers a sophisticated approach to visual processing that improves selectivity. In the presence of a preferred stimulus, the PushPull-Conv unit exhibits a pronounced output due to the strong response from the push filter and a weak response from its complementary pull filter, enhancing the detection of specific

features. In contrast, applying a low-pass filter first would attenuate high-frequency details, potentially diminishing critical information before the push filter acts, thereby affecting its effectiveness in identifying the preferred stimulus. Additionally, in the absence of the preferred stimulus, the PushPull-Conv unit ensures minimal output by balancing the similar responses of the push and pull filters, thereby cancelling each other out and enhancing overall selectivity. This mechanism contrasts with the low-pass-then-push approach, which might reduce high-frequency noise but could amplify non-preferred low-frequency stimuli, leading to less specificity. Furthermore, the PushPull-Conv unit dynamically adjusts its response across the entire spectrum of spatial frequencies, providing a high degree of selectivity and sensitivity, while the sequential filtering approach lacks this nuanced control and adaptability, making it less effective in discriminating between preferred and non-preferred stimuli with the same level of precision.

The main contribution of this work is three-fold. (1) A new computational model for push-pull inhibition is proposed. We call this PushPull-Conv and propose to replace the traditional convolutional unit as the first layer of the ConvNet. This innovative architectural change not only heightens precision in identifying preferred stimuli but also effectively attenuates non-preferred stimuli across all frequency ranges. This introduces a robust alternative to traditional approaches that sequentially apply noise-specific filters followed by push-only filtering, thereby offering a more balanced and comprehensive method for feature detection and image processing. (2) Data augmentation strategies such as AugMix, AutoAug, TrivialAugment, and PRIME are known to improve model robustness. Our research shows that combining PushPull-Conv with these data augmentation techniques further enhances model robustness, yielding additional improvements. (3) Propose a new metric to quantify the trade-off between robustness and accuracy.

## 2 Related Works

Deep networks can be susceptible to image perturbations that humans cannot perceive [26,5], mainly because the model relies on features unrelated to the task. For instance, networks (such as ResNet) trained on IMAGENET for object detection are biased towards texture-based features instead of shape-based, when in fact, humans rely more on the shape for image classification [8]. Although adversarial robustness is not within the scope of this study, it is worth mentioning that this research direction has seen tremendous progress with every adversarial attack [3,4] countered with adversarial defences [25,22], and vice-versa. Some other methods that attempt to improve robustness are model pre-training [12], enhancing performance on in-distribution dataset [34], and using larger models [11].


Data augmentation improves the generalization and robustness of deep models. For instance, image transformations such as flips, crops, resize, affine transforms and other methods are commonly used to augment the training

data. A few popular techniques include, cut-mix [42], AutoAug [6], AugMix [13], PRIME [23], TrivialAugment [24], DeepAugment [10], and adversarial training [9,38], among others. Hendrycks et al. [10] found that using larger models and artificial data augmentation improves model robustness across several corruption types. Interestingly, they showed that out-of-distribution training using the intermediate feature maps from deep networks (artificial data) also improves model robustness to natural image corruptions. Yin et al. [41] performed an analysis of model robustness in the Fourier domain. They showed how models trained with certain augmentations, such as Gaussian noise, and adversarial training improve model robustness against several noise and blur corruptions while reducing robustness to fog and contrast corruptions. Though data augmentation can address model robustness to a certain extent, it can be made effective only with the availability of a lot of data [28].

An orthogonal line of research is to make improvements to the model architecture [37,19,32]. Such methods could be used in conjunction with data augmentation to further enhance model robustness. For instance, Liu et al. [19] proposed a regularization technique called group inhibition, to extract auxiliary features for robust classification. When combined with data augmentation, [37,19] show that their methods improve results in comparison to when only data augmentation is used. Our work also uses a similar strategy to make the models more robust by introducing changes to the architecture. However, our work differs in the specifics of the architectural design. In [37,43], it was proposed to include a strided blur filter following strided convolutions to address the aliasing problem in ConvNets. These modifications were done to several ConvNet layers. In contrast, we propose to replace only the first layer of a ConvNet with the proposed PushPull-Conv layer, resulting in a more computationally efficient solution.

Our architectural design is biologically inspired by the push-pull phenomenon exhibited by simple cells in the mammalian brain [15,16,7,2,14]. In [1], a computational model of this phenomenon was proposed. The resulting model can be used as a contour detection operator that achieves superior robustness against corruption. It was later embedded in the RUSTICO line detector [31] and coupled with surround inhibition [21]. Taking inspiration from simple cells that receive such inhibition, we propose a unit, namely the PushPull Convolutional unit or PushPull-Conv for short. The design of this unit is further elucidated in Sec. 3.1.

While our work might appear akin to that of [32] at first glance, it fundamentally diverges in both methodology and scope. Firstly, we propose a modified design of the push-pull computational unit and make it configurable with learnable inhibition strengths. In the earlier attempt, with learnable inhibition the robustness became inconsistent, slowed down training, and was left for future work. Trainable inhibition in our approach leads to consistent results without slowing down the training and is a default choice for hyperparameters. Secondly, the push-pull approach in [32] is merely an *erosion*



**Fig. 2.** The proposed push-pull computation unit. Refer Eq. (5).

*operation.* Without the upsampling factor in [32], the pull component would be redundant. Our new way of determining the pull kernel results in a push-pull component that is effective even without pooling. Thirdly, we show that our unit complements other data augmentation methods.

### 3 Approach

#### 3.1 The Push Pull unit

The proposed computational design for performing push-pull convolutions is the rectified difference between the responses of the push and the pull convolutions on the input image (Fig. 2). The push kernel is analogous to the traditional convolutional unit in ConvNets and is defined as:

$$\mathbf{x}_{\text{out}} = \mathbf{x} * \mathbf{w} + b \quad (1)$$

where  $\mathbf{x}$  denotes the input signal,  $\mathbf{w}$  denotes the convolutional kernel, and  $b$  is the bias. The pull kernel is derived from the push kernel by taking its complement such that it responds to the same preferred patterns of the push kernel but of opposite contrast. This is inspired by the receptive field with the antiphase of simple neurons that experience this kind of inhibition [16]. Formally, we determine the pull kernel as follows.

$$\mathbf{w}'_{\text{pull}} = 1 - \frac{\mathbf{w} - w_{\min}}{w_{\max} - w_{\min}} \quad (2)$$


where  $\mathbf{w}$  denotes the weights of the push kernel and its maximum and minimum values are represented by  $w_{\max}$  and  $w_{\min}$ , respectively. The values of  $\mathbf{w}'_{\text{pull}}$  are rescaled to match that of  $\mathbf{w}$ , which determines the final pull kernel:

$$\mathbf{w}_{\text{pull}} = \mathbf{w}'_{\text{pull}} \times (w_{\max} - w_{\min}) + w_{\min} \quad (3)$$

$$\mathbf{w}_{\text{pull}} = -\mathbf{w} + w_{\max} + w_{\min} \quad (4)$$

Equation (3) can be simplified to Eq. (4) using Eq. (2). Figure 3 shows a few examples of push and pull kernels.

Since the pull kernel is a non-linear function of the push kernel, it indirectly introduces non-linearity into the pull response. The pull response can be



**Fig. 3.** (Top row) Five randomly selected filters from the conv1 layer of ResNet50 (trained on IMAGENET). For a visual illustration, only the first filter channel is depicted. (Bottom row) The corresponding pull kernels as determined by Eq. (4).

considered as inhibition and it is subtracted from the response generated by the push kernel. This results in the push-pull computation:

$$\mathbf{x}_{\text{out}} = r(\mathbf{x} * \mathbf{w}) - \alpha \cdot r(\mathbf{x} * \mathbf{w}_{\text{pull}} * \mathbf{1}) + b \quad (5)$$

where  $\alpha$  is the pull inhibition factor,  $r$  is the ReLU activation, and  $\mathbf{1}$  represents the average filter. The average pooling increases the receptive field size of the pull component. It is incorporated in our work because of the great benefit it contributed to the computational model as first proposed in [1]. It is inspired by the neurophysiological discovery that interneurons have broader receptive fields than those of the neurons they inhibit [20]. We extend this idea to any arbitrary (push) kernel learnt by ConvNet, and through push-pull computation to possibly achieve the properties of a band-pass filter. The factor  $\alpha$  controls the degree of pull inhibition and can either be configured as a trainable parameter or set to a fixed value, thereby making it a hyper-parameter. We experimentally show that the latter scheme is useful when manually controlling the trade-off between the model performance for in- and out-of-distribution data (refer to Sec. 5.2 for further details).

Figure 2 depicts the proposed push-pull unit. Note that the bias term is only added at the final stage. Both the push and pull convolutions are performed without any bias. This is done to simplify the push-pull design and to remove any undesired influence of the bias when the pull response is subtracted from the push.

### 3.2 Characteristics of Push Pull convolutions

To illustrate the design of the push-pull kernel, let's consider a simpler form of Eq. (5) without the non-linear activation  $r$ . In this case, Eq. (5) simplifies to:


$$\mathbf{x}_{\text{out}} = \mathbf{x} * \mathbf{w} - \alpha \cdot (\mathbf{x} * \mathbf{w}_{\text{pull}} * \mathbf{1}) + b \quad (6)$$

$$= \mathbf{x} * (\mathbf{w} - \alpha \cdot \mathbf{w}_{\text{pull}} * \mathbf{1}) + b \quad (7)$$


$$= \mathbf{x} * f(\mathbf{w}; \alpha, \mathbf{1}) + b \quad (8)$$

where the push-pull kernel is represented by the non-linear function  $f(\mathbf{w}; \alpha, \mathbf{1})$ .

We conducted an analysis in the Fourier domain to gain deeper insights into the characteristics of push-pull convolutions, by considering all the 64 kernels from the first layer of a ResNet50 trained on IMAGENET. The computed spectra are presented in Fig. 5. On a careful examination, it can be seen that for smaller values of inhibition strength  $\alpha$ , the push-pull filters exhibit the property of band-pass. This trend is observed for both configurations of pooling. Note that, band-pass filtering has been used for edge detection, for instance, using Difference-of-Gaussians (DoG) filters. Therefore, for smaller values of  $\alpha$ , the push-pull mechanism induces a minor bias for extracting shape-based features. Thus, it can assist the model in becoming more robust to high-frequency corruptions. We illustrate this behaviour further with a simple example in Fig. 4, where the SNR improves with the push-pull mechanism.



**Fig. 4.** Illustration of a push-pull filtering with a simulated input corrupted by Gaussian noise. The push kernel is learned by ResNet50. SNR computed as  $20 \log_{10}(A_s/A_n)$ , where  $A_s$  is the average across a 5-pixel vertical edge, and  $A_n$  is the average response in the background.



**Fig. 5.** Fourier spectral analysis of push-pull filters in ResNet50. The averaged push spectrum of the 64 push kernels is presented in the top right, while the corresponding pull spectrum is shown at the bottom right. These push and pull spectrum are then combined with various choices of hyperparameters  $\alpha$  and AvgPool, and the resulting push-pull spectra are displayed on the left. For visual illustration, only the first filter channel is depicted.

### 3.3 Embedding PushPull-Conv in ConvNets

The idea is conceptually simple and easy to implement. We replace the first convolutional layer in a ConvNet with PushPull-Conv, as shown in Fig. 1. The rationale behind using such a computational unit in the first layer of the



network is also biologically inspired. Neurophysiology studies show that this kind of inhibition is found in the early stages of the visual system. Besides simple cells in the primary visual cortex, the lateral geniculate nucleus (the intermediate unit between the retina and the primary visual cortex) contains cells with centre-surround receptive fields. These receptive fields can be modelled by the Laplacian of Gaussian or DoG operators. DoG filters can also be considered a type of push-pull filter, and they achieve band-limited selectivity. As shown in Fig. 5, with appropriately chosen inhibition strength  $\alpha$ , the filters in the PushPull-Conv layer tend to act as band-pass filters and potentially suppress high-frequency corruptions, thereby, correcting the distributional shift for the layers following it in the network.

## 4 Experiments and Results

To evaluate the effectiveness of our proposed approach on out-of-distribution images, we train our models on CIFAR10 [18] and IMAGENET [27] datasets, and test the resulting models on CIFAR10-C and IMAGENET-C [11]. The latter two datasets contain the corrupted versions of images in the former two datasets. The test images are perturbed with 15 types of image corruptions (broadly categorized into noise, blur, weather, and digital) and 5 levels of severity for each corruption type (see Fig. 9 in the [Supplementary Material](#)).

To analyze how the robustness of the PushPull-Conv-based model scales with increasing dataset size, two additional subsets of IMAGENET were created with 100 and 200 classes each (IMAGENET100 and IMAGENET200). The classes were randomly chosen and their corresponding corrupted test set versions are referred to as IMAGENET100-C and IMAGENET200-C.

### 4.1 Evaluation metrics

We use the mean corrupted error as a metric to measure robustness [11]. Let the error made by a classifier  $f$ , on a corruption type  $c$ , with a severity level  $s$ , be denoted by  $E_{c,s}^f$ . Then the (absolute) mean corruption error for a classifier  $f$  is defined as:

$$mCE = \frac{1}{|C|} \sum_{c \in C} (CE_c^f) = \frac{1}{|C|} \sum_{c \in C} \left( \frac{1}{S} \sum_{s=1}^S E_{c,s}^f \right) \quad (9)$$

where,  $mCE$  is the mean of corruption errors  $CE_c^f$  over the set of corruptions  $C \in \{\text{Gaussian Noise, Shot Noise, ..., JPEG}\}$ . When comparing the performance of a model  $f$  to a baseline  $b$ , it is useful to represent it as the mean of relative corruption errors, which is given by:

$$mrCE = \frac{1}{|C|} \sum_{c \in C} \left( CE_c^f / CE_c^b \right) \quad (10)$$

## 4.2 Experimental setup

We investigate the model robustness by replacing the first layer of ResNet with the proposed PushPull-Conv. As the original ResNet architecture for IMAGENET had a conv1 layer with 64 filters of size  $7 \times 7$  and a stride of 2, we retain the same configuration for the PushPull-Conv and explore various settings for the hyperparameters - *average kernel size* and the *inhibition strength*. As the *mrCE* scores (refer to Eq. (10)) are always computed with respect to a baseline, we hereby always refer to the vanilla ResNet as the baseline. With this setting, we evaluate the *mrCE* scores of the PushPull variants keeping the corresponding vanilla ResNet as the baseline. For the experiments on the CIFAR10 dataset, we modify the conv1 layer to have 64 filters with size  $3 \times 3$  and with a stride of 1. This change in hyperparameters was made to handle the smaller image dimensions ( $32 \times 32$  pixels) of that dataset.

The classification models were optimized using the Stochastic Gradient Descent (SGD) with a momentum of 0.9 and a weight decay of  $1e-5$ . A cyclic learning rate schedule [30] was employed wherein the learning rate was increased from 0.05 to 1.0 for the first 30 per cent of the steps and decayed thereafter to  $5e-5$ . The learning rate warm-up and decay both followed a cosine policy. The models were trained for 20 epochs with a batch size of 128. The objective function was determined using categorical cross-entropy loss. This training recipe was chosen to enable faster experimentation and validate our hypothesis. The aforementioned recipe was used to train ResNet18 and ResNet50 on IMAGENET. For hyperparameters concerning other models and dataset combinations, we refer the reader to our source code<sup>4</sup>.

## 4.3 Experiments

Experiments were conducted using ResNet18 and ResNet50 models, on 4 different datasets CIFAR10, IMAGENET100, IMAGENET200, and IMAGENET.

We consider the vanilla ResNet as the baseline for each experiment and compare it against two push-pull variants. In particular, we explore the impact of the hyper-parameter *average kernel size* and keep the *inhibition strength* trainable. The classification results presented in Table 1 indicate that the push-pull variants consistently perform better than the respective baselines in terms of robustness. Since all scores are normalized to the baseline, error values less than 1 show an improvement over the baseline’s performance. It is also worth noting that the relative improvement in robustness to the baseline does not diminish with increasing dataset size.

We conducted a preliminary grid search to determine the baseline network’s hyperparameters. The best-performing one was then used to train the push-pull models, ensuring an unbiased selection of learning hyperparameters. Although we conducted experiments with two sizes of the *average filter*, we cannot determine

<sup>4</sup> <https://github.com/bgswaroop/pushpull-conv>

**Table 1.** Classification results. D denotes the dataset,  $E$  is the top-1 error rate on the clean test set. The corruption types are **Noise** - (Ga)ussian, (Sh)ot, and (Im)pulse, **Blur** - (De)focous, (Gl)ass, (Mo)tion, and (Zo)om, **Weather** - (Sn)ow, (Fr)ost, and (Fo)g, and **Digital** - (Br)ightness, (Co)ntrast, (El)astic, (Pi)xelate, and (Jp)eg. The push-pull variants are shown underneath the respective baselines. The baselines are marked with (b#). Scores showing improved robustness are highlighted in green, including those below 1 before rounding to 2 decimal places. The corresponding absolute corruption errors are in Table 6 of the [Supplementary Material](#).


D	Model	$E$	$mCE$	$mTCE$	Noise			Blur				Weather			Digital																	
					Ga	Sh	Im	De	Gl	Mo	Zo	Sn	Fr	Fo	Br	Co	El	Pi	Jp													
CIFAR10	ResNet18 (b1)	.069	.2853	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00		
	PushPull avg3	.071	.2728	.970	.86	.86	1.01	1.00	1.00	.97	1.01	.98	.90	1.07	1.01	1.03	.99	.94	.90													
	PushPull avg5	.071	.2708	.965	.88	.89	.98	.99	.92	.96	1.00	.98	.96	.99	1.04	1.03	1.00	.94	.93													
	ResNet50 (b2)	.069	.3028	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
	PushPull avg3	.076	.2909	.982	.88	.86	.99	.99	.93	1.04	.97	1.01	.91	1.06	1.07	1.09	.99	.98	.94													
	PushPull avg5	.068	.2866	.959	.89	.89	.94	.97	.95	1.00	.94	.97	.90	.98	1.02	1.06	.97	.97	.95													
IMAGENET100	ResNet18 (b3)	.207	.5812	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00		
	PushPull avg3	.208	.5594	.961	.92	.92	.92	1.00	.95	.99	1.01	1.01	.99	.97	1.01	1.00	.97	.97														
	PushPull avg5	.221	.5662	.976	.92	.91	.92	1.02	.94	1.00	1.00	1.02	1.01	1.02	1.06	1.02	1.00	.85	.95													
	ResNet50 (b4)	.189	.5691	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
	PushPull avg3	.194	.5517	.972	.95	.95	.95	.97	.91	1.01	1.00	.96	.98	1.00	1.04	.99	1.00	.90	.96													
	PushPull avg5	.203	.5452	.962	.91	.92	.91	.95	.88	.97	.95	.99	.98	1.05	1.06	1.02	.99	.92	.93													
IMAGENET200	ResNet18 (b5)	.246	.6424	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00		
	PushPull avg3	.252	.6291	.979	.97	.97	.96	1.00	.97	.98	.98	1.01	.99	1.03	1.02	1.02	.98	.88	.93													
	PushPull avg5	.253	.6280	.976	.97	.97	.96	1.00	.97	.98	.99	.99	.99	1.02	1.01	1.02	.99	.87	.92													
	ResNet50 (b6)	.213	.6183	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
	PushPull avg3	.217	.6076	.979	.99	.98	.99	1.01	.97	.98	.99	1.01	.98	1.01	1.01	1.00	.98	.83	.96													
	PushPull avg5	.222	.6062	.978	.98	.97	.98	.99	.96	.98	.99	1.02	.99	1.05	1.03	1.03	.98	.84	.90													
IMAGENET	ResNet18 (b7)	.325	.7167	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00		
	PushPull avg3	.363	.6776	.949	.95	.94	.93	.95	.88	.96	.98	.98	.97	1.09	1.08	1.01	.93	.67	.91													
	PushPull avg5	.337	.6952	.970	.99	.98	.98	.98	.96	.98	.98	.99	.98	1.01	1.01	.99	.94	.80	.96													
	ResNet50 (b8)	.269	.6674	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
	PushPull avg3	.282	.6454	.967	.99	.98	.98	.96	.96	.98	.97	.99	1.00	1.02	.99	.99	.91	.81	.97													
	PushPull avg5	.276	.6452	.966	.99	.98	.99	.98	.95	.97	.97	.97	.97	.99	.98	.99	.92	.83	1.00													

which filter size is more appropriate for IMAGENET. Perhaps this hyperparameter needs to be explored for other applications and ConvNet architectures. For the models trained on CIFAR10, however, pooling with average filters of size  $5 \times 5$  shows a greater benefit than with  $3 \times 3$  kernels.

Another design choice in the push-pull unit is the *inhibition strength*  $\alpha$ . For the experiments reported in Table 1,  $\alpha$  was kept as a trainable parameter. Figure 6 shows the distribution of the learned  $\alpha$  values. Most of the values are close to zero, indicating the network’s preference to learn filters that are similar to push or with the property of band-pass filtering. With  $3 \times 3$  average pooling, the  $\alpha$  takes on values from a larger range. When  $\alpha$  becomes large in magnitude, the resulting push-pull filter becomes a low-pass filter. We can also keep  $\alpha$  as a fixed hyperparameter which would give us more control over the trade-off between robustness and clean accuracy, discussed further in Sec. 5.1.

#### 4.4 Comparison with the state-of-the-art

For a fair comparison, we select methods that only alter architecture for robustness. [43] proposed using blur filters after strided convolutions to mitigate aliasing impact. [37] observed using blur filters before non-linear



**Fig. 6.** Distribution of learnt inhibition strength  $\alpha$  for PushPull-based ResNets trained on IMAGENET.

activations make ConvNets resilient to aliasing. We trained all models with the same training recipe. Additionally, we compare our approach with the earlier attempt at implementing the push-pull inhibition in a ConvNet [32]. All these results are summarized in Table 2, and later on discussed.

**Table 2.** Comparison with the state-of-the-art. BlurConv1 [43], BlurConv2 [37], the previous attempt of PushPull [32], and with EfficientNet [33]. The corresponding absolute corruption errors are in Table 7 of the [Supplementary Material](#).

D Model	$E$	$mCE$	$mvCE$	Noise			Blur			Weather			Digital						
				Ga	Sh	Im	De	Gl	Mo	Zo	Sn	Fr	Fo	Br	Co	El	Pi	Jp	
CIFAR10	ResNet50 (b2)	.069	.2809	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
	PushPull avg3 (ours)	.076	.2681	.982	0.88	0.86	0.99	0.99	0.93	1.04	0.97	1.01	0.91	1.06	1.07	1.09	0.99	0.98	0.94
	PushPull avg5 (ours)	.068	.2808	.959	0.89	0.89	0.94	0.97	0.95	1.00	0.94	0.97	0.90	0.98	1.02	1.06	0.97	0.97	0.95
	PushPull in [32]	.069	.2984	.973	0.94	0.92	0.93	0.97	0.94	1.01	0.97	0.99	0.91	1.02	0.99	1.07	0.99	0.97	0.97
	BlurConv1	.069	.2831	.948	0.85	0.86	1.00	0.95	0.89	0.96	0.99	0.95	0.93	0.98	0.96	0.94	0.96	0.98	0.99
	BlurConv2	.098	.3648	1.24	1.06	1.10	1.07	1.34	1.13	1.46	1.34	1.27	1.14	1.35	1.41	1.37	1.29	1.26	1.08
IMAGENET	ResNet50 (b8)	.269	.6670	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
	PushPull avg3 (PP3)	.282	.6450	.967	0.99	0.98	0.98	0.96	0.96	0.98	0.97	0.99	1.00	1.02	0.99	0.99	0.91	0.81	0.97
	PushPull avg5 (PP5)	.276	.6450	.966	0.99	0.98	0.99	0.98	0.95	0.97	0.97	0.97	0.97	0.99	0.98	0.99	0.92	0.83	1.00
	PushPull in [32]	.269	.6650	.996	1.00	1.00	1.02	0.99	1.00	0.99	0.99	0.99	0.98	1.00	0.98	0.99	0.99	0.99	1.03
	BlurConv1	.265	.6260	.962	0.97	0.96	0.97	0.97	0.97	0.98	0.97	0.97	0.96	0.95	0.96	0.97	0.95	0.94	0.95
	BlurConv2	.276	.6210	.958	0.90	0.89	0.89	0.99	0.97	0.99	0.96	1.00	0.96	0.98	1.00	1.00	0.98	0.70	0.87
	BlurConv2+PP3	.307	.6166	.924	0.91	0.89	0.89	0.95	0.89	0.96	0.92	1.01	0.97	1.09	1.08	1.03	0.91	0.55	0.87
	Res50+AugMix (b9)	.269	.6121	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	Res50+AugMix+PP3	.287	.5900	.968	0.95	0.94	0.93	0.98	0.94	0.99	0.95	1.01	0.99	1.06	1.02	1.00	0.89	0.84	1.02
	Res50+AutoAug (b10)	.269	.6298	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	Res50+AutoAug+PP3	.283	.6041	.961	0.95	0.95	0.94	0.97	0.97	0.99	0.98	1.00	0.98	1.02	1.01	0.95	0.93	0.82	0.96
	Res50+PRIME (b11)	.298	.5223	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Res50+PRIME+PP3	.306	.4995	.954	0.91	0.90	0.89	0.99	0.97	1.00	1.01	0.99	0.98	1.00	1.00	0.96	0.96	0.84	0.90	
Res50+TrivialAug (b12)	.269	.5961	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
Res50+TrivialAug+PP3	.286	.5867	.985	1.02	1.02	0.99	1.00	0.96	1.01	0.99	1.00	1.00	1.07	1.03	0.96	0.92	0.83	0.97	
EffNet-b0 (b13)	.286	.6818	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	
EffNet-b0 + PP3	.288	.6660	.976	0.98	0.99	0.97	0.99	0.98	1.00	1.00	1.00	0.99	1.00	0.99	1.01	0.98	0.91	0.86	

The most effective standalone method on CIFAR10 is identified as BlurConv1 [43], while on IMAGENET, BlurConv2 [37] leads. Our method not only outperforms earlier approaches to push-pull mechanisms but also approaches the effectiveness of BlurConv. Notably, our method enhances robustness by altering only one aspect of the ConvNet, unlike BlurConv, which modifies all layers with strided convolutions and pooling. Hence, our method achieves robustness with minimal computational overhead. Furthermore, our method enhances the robustness provided by BlurConv2 when coupled together, albeit with a reduction in clean accuracy. For every type of corruption where BlurConv2 is effective, our PushPull-Conv method has


increased robustness, with a particularly notable enhancement against pixelate corruption.

The proposed approach complements the robustness offered by data augmentation schemes. We trained ResNet50 with data augmentation schemes including AutoAug, AugMix, PRIME, and TrivialAugment. On adding the PushPull layer to these models, the robustness improved further (see Table 2). This demonstrates our method’s compatibility and effectiveness in conjunction with current schemes. Moreover, in terms of absolute  $mCE$ , push-pull with PRIME augmentations achieves the best-ever  $mCE$  score of 49.95% with ResNet50. This is a new benchmark when considering models with about 25.5M parameters.

Finally, the proposed method was evaluated with the EfficientNet-b0 [33] architecture. We integrated the push-pull layer in the first convolutional layer, that is, in the stem of EfficientNet. The results presented in Table 2 clearly demonstrate the robustness improvements achieved with PushPull-Conv. While further investigation into various ConvNet-based architectures is necessary, the evidence from ResNets and EfficientNet implementations shows promising outcomes. In fact, with EfficientNet we achieve robustness with negligible cost to clean error. We present a metric to quantify such a trade-off in the following section.

## 5 Discussion

### 5.1 Clean error vs robustness



**Fig. 7.** Comparison of PushPull-Conv configurations with vanilla ResNet18. The PushPull variants improve mean corruption error (mCE) over the baseline on IMAGENET-C, with adjustable hyper-parameters to control the trade-off between robustness and clean error. (Best viewed in colour)

Tsipras et. al [35] explained the trade-off between robustness and generalization (clean accuracy) and concluded that the interplay between robustness and standard accuracy might be more nuanced than one might expect. We investigate this tradeoff with PushPull-Conv for various settings of hyperparameters. The inhibition strength  $\alpha$  plays a key role and can be used as

**Table 3.** Summary of experiments using various push-pull configurations with ResNet18 trained on IMAGENET100. The results of configurations labelled PushPull\* are averaged over experiments with values of  $\alpha \in \{0.5, 1, 2, 3, 4, 5, 6\}$ .

Configuration	Clean error	mrCE
ResNet18 (b3)	0.207	1.000
PushPull*: no pooling	$0.249 \pm 0.012$	$0.932 \pm 0.012$
PushPull*: avg 3	$0.251 \pm 0.015$	$0.929 \pm 0.015$
PushPull*: avg 5	$0.260 \pm 0.020$	$0.928 \pm 0.017$

a knob. As depicted in Fig. 7, with increasing values of  $\alpha$  the models become progressively resilient to image corruption while taking a dip in performance on clean data. We propose a metric to quantify the trade-off between clean error and robustness (see Sec. 5.2). The best trade-off is achieved when  $\alpha$  is set as a trainable parameter and AvgPool  $3 \times 3$ .

We evaluate the contribution of the average pooling to the overall robustness. To this end, we conducted three sets of experiments. The first set had no pooling and the other two were conducted with average pooling of sizes  $3 \times 3$  and  $5 \times 5$ . The results are summarized in Table 3. It can be seen that the robustness improves with higher levels of pooling while it has the opposite effect on the clean error.

## 5.2 A Quantitative Metric for Assessing the Trade-off Between Robustness and Clean Accuracy

It is important to quantify the trade-off between clean error and robustness. Figure 7 illustrates this trade-off when considering ResNet18 models trained with several configurations of hyperparameters on IMAGENET100. To identify which configuration of hyperparameters results in the best trade-off, we need to quantify it. Therefore, we introduce a metric, which quantifies the net reduction in error rate of a model with respect to a baseline.

For a given model  $m$ , let the clean error be denoted by  $E$  and the mean corruption error as  $mCE$ . The corresponding errors for the baseline are denoted by  $E_0$  and  $mCE_0$ . On considering clean error, the net reduction in error rate for a model  $m$  w.r.t a baseline is given by:

$$R_E = \frac{E_0 - E}{E_0} \quad (11)$$

Likewise, the reduction in error rate for  $mCE$  is given by:

$$R_{mCE} = \frac{mCE_0 - mCE}{mCE_0} \quad (12)$$

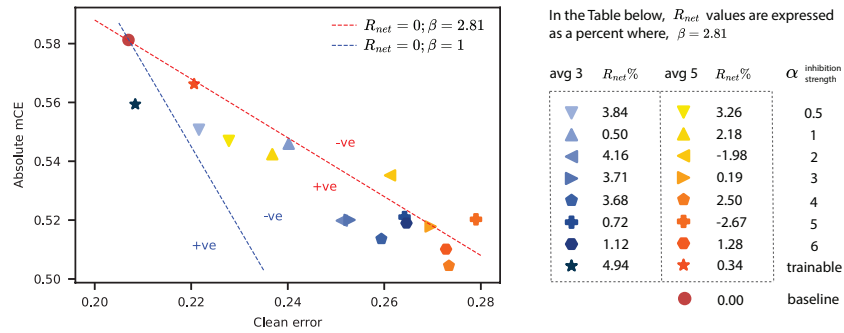
The net reduction in error rate is then determined by:

$$R_{net} = \frac{1}{2}(R_E + \beta \times R_{mCE}) \quad (13)$$

where the factor  $\beta$  determines the difficulty in reducing  $mCE$  over  $E$ . When  $\beta = mCE_0/E_0$ , it quantifies the difficulty faced by the baseline to achieve the

same level of robustness as the clean accuracy. Setting  $\beta = 1$  implies giving equal weight to both robustness and accuracy, without accounting for the difficulty of the problem. For a given model  $m$ , when  $R_{net} = 0$  it implies the model is equivalent in terms of net performance to the baseline. When  $R_{net} > 0$ , the performance tradeoff favours model  $m$  compared to the baseline. Conversely,  $R_{net} < 0$  indicates a negative tradeoff with respect to the baseline.

Figure 8 illustrates the robustness of the push-pull variants compared to the baseline. It is interesting to note that setting  $\beta = mCE_0/E_0 = 2.81$ , accounts for the difficulty of improving robustness over clean accuracy. In this scenario, almost all variants of push-pull positively reduce the error over the baseline. Even when considering equal difficulty for both problems ( $\beta = 1$ ), the push-pull variant with AvgPool  $3 \times 3$  and trainable inhibition performs better than the baseline.



**Fig. 8.** Comparison of PushPull-Conv configurations with baseline vanilla ResNet18 trained on IMAGENET100. The dotted lines indicate the points at which the models do not make any improvements over the baseline in terms of net reduction in error rate. Each line splits the error space into two regions (+ve and -ve). The models in the +ve regions are the favourable ones with  $R_{net} > 0$  and vice-versa. (Best viewed in colour)

### 5.3 Robustness to high-frequency corruptions

In most experiments, a trend was observed where the PushPull models achieved higher robustness to high- and mid-frequency corruptions compared to their low-frequency counterparts. To illustrate this, we categorized the 15 corruption types from the CIFAR10-C dataset into low-, mid-, or high-frequency corruptions based on the analysis of [41]. Frost, fog, brightness, and contrast are classified as low-frequency corruptions; glass-blur, motion-blur, zoom, snow, and elastic are classified into mid-frequency corruptions, and the rest as high-frequency. Table 4 shows the  $mrCE$  scores when the corruption types are grouped based on their frequency. The results indicate that PushPull offers more robustness to high-frequency corruptions than low-frequencies. This trend is consistent with the other datasets and ResNet models.

**Table 4.** The averaged *mrCE* scores for CIFAR10-C corruption types, grouped by corruption frequency. All PushPull models are trained with trainable inhibition  $\alpha$ .

Model	Low freq.	Mid freq.	High freq.
ResNet18 (b1)	1.00 $\pm$ 0.000	1.00 $\pm$ 0.000	1.00 $\pm$ 0.000
PushPull avg3	1.00 $\pm$ 0.074	0.99 $\pm$ 0.013	0.93 $\pm$ 0.069
PushPull avg5	1.01 $\pm$ 0.036	0.97 $\pm$ 0.032	0.93 $\pm$ 0.046
ResNet50 (b2)	1.00 $\pm$ 0.000	1.00 $\pm$ 0.000	1.00 $\pm$ 0.000
PushPull avg3	1.03 $\pm$ 0.025	0.99 $\pm$ 0.026	0.94 $\pm$ 0.032
PushPull avg5	0.99 $\pm$ 0.019	0.97 $\pm$ 0.013	0.93 $\pm$ 0.040

## 6 Conclusion

PushPull-based convolutions improve the robustness of ResNet architectures to out-of-distribution image corruptions, especially those with high-frequency corruptions. This robustness is achieved without data augmentation and with a relatively small computational cost. Furthermore, this approach can be combined with data augmentation. In particular, when combined with PRIME augmentation, it achieves the best-ever reported *mCE* of 49.95% using ResNet50 on IMAGENET-C.

## References

1. Azzopardi, G., Rodríguez-Sánchez, A., Piater, J., Petkov, N.: A push-pull CORF model of a simple cell with antiphase inhibition improves SNR and contour detection. *PloS one* **9**(7), e98424 (2014)
2. Borg-Graham, L.J., Monier, C., Fregnac, Y.: Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature* **393**(6683), 369–373 (1998)
3. Carlini, N., Wagner, D.: Defensive distillation is not robust to adversarial examples. *arXiv preprint arXiv:1607.04311* (2016)
4. Carlini, N., Wagner, D.: Adversarial examples are not easily detected: Bypassing ten detection methods. In: *Proceedings of the 10th ACM workshop on artificial intelligence and security*. pp. 3–14 (2017)
5. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: *2017 IEEE Symposium on Security and Privacy (SP)*. pp. 39–57. *Ieee* (2017)
6. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: *CVPR* (June 2019)
7. Ferster, D.: Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. *Journal of Neuroscience* **8**(4), 1172–1180 (1988)
8. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: *Proceedings of the International Conference on Learning Representations* (2019)
9. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *ICLR* (2015)
10. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical



- analysis of out-of-distribution generalization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8340–8349 (2021)
11. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: International Conference on Learning Representations (2019)
  12. Hendrycks, D., Lee, K., Mazeika, M.: Using pre-training can improve model robustness and uncertainty. In: International Conference on Machine Learning. pp. 2712–2721. PMLR (2019)
  13. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple data processing method to improve robustness and uncertainty. In: International Conference on Learning Representations (2019)
  14. Hirsch, J.A., Alonso, J.M., Reid, R.C., Martinez, L.M.: Synaptic integration in striate cortical simple cells. *Journal of neuroscience* **18**(22), 9517–9528 (1998)
  15. Hubel, D.H., Wiesel, T.N.: Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology* **160**(1), 106 (1962)
  16. Hubel, D.H., Wiesel, T.N.: Brain mechanisms of vision. *Scientific American* **241**(3), 150–163 (1979)
  17. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8110–8119 (2020)
  18. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
  19. Liu, H., Wu, H., Xie, W., Liu, F., Shen, L.: Group-wise inhibition based feature regularization for robust classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 478–486 (2021)
  20. Martinez, L.M., Wang, Q., Reid, R.C., Pillai, C., Alonso, J.M., Sommer, F.T., Hirsch, J.A.: Receptive field structure varies with layer in the primary visual cortex. *Nature neuroscience* **8**(3), 372–379 (2005)
  21. Melotti, D., Heimbach, K., Rodríguez-Sánchez, A., Strisciuglio, N., Azzopardi, G.: A robust contour detection operator with combined push-pull inhibition and surround suppression. *Information sciences* **524**, 229–240 (2020)
  22. Metzen, J.H., Genewein, T., Fischer, V., Bischoff, B.: On detecting adversarial perturbations. arXiv preprint arXiv:1702.04267 (2017)
  23. Modas, A., Rade, R., Ortiz-Jiménez, G., Moosavi-Dezfooli, S.M., Frossard, P.: Prime: A few primitives can boost robustness to common corruptions. In: European Conference on Computer Vision (ECCV) (2022)
  24. Müller, S.G., Hutter, F.: Trivialaugment: Tuning-free yet state-of-the-art data augmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 774–782 (October 2021)
  25. Papernot, N., McDaniel, P., Wu, X., Jha, S., Swami, A.: Distillation as a defense to adversarial perturbations against deep neural networks. In: 2016 IEEE symposium on security and privacy (SP). pp. 582–597. IEEE (2016)
  26. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do cifar-10 classifiers generalize to cifar-10? arXiv preprint arXiv:1806.00451 (2018)
  27. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* **115**(3), 211–252 (2015)

28. Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., Madry, A.: Adversarially robust generalization requires more data. *Advances in neural information processing systems* **31** (2018)
29. Shiva Kasiviswanathan, N., et al.: Simple black-box adversarial attacks on deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pp. 6–14 (2017)
30. Smith, L.N., Topin, N.: Super-convergence: Very fast training of neural networks using large learning rates. In: *Artificial intelligence and machine learning for multi-domain operations applications*. vol. 11006, pp. 369–386. SPIE (2019)
31. Strisciuglio, N., Azzopardi, G., Petkov, N.: Robust inhibition-augmented operator for delineation of curvilinear structures. *IEEE Transactions on Image Processing* **28**(12), 5852–5866 (2019)
32. Strisciuglio, N., Lopez-Antequera, M., Petkov, N.: Enhanced robustness of convolutional networks with a push–pull inhibition layer. *Neural Computing and Applications* **32**(24), 17957–17971 (2020)
33. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International conference on machine learning*. pp. 6105–6114. PMLR (2019)
34. Taori, R., Dave, A., Shankar, V., Carlini, N., Recht, B., Schmidt, L.: When robustness doesn’t promote robustness: Synthetic vs. natural distribution shifts on imagenet (2019)
35. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. In: *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=SyxAb30cY7>
36. Vaish, P., Wang, S., Strisciuglio, N.: Fourier-basis functions to bridge augmentation gap: Rethinking frequency augmentation in image classification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 17763–17772 (June 2024)
37. Vasconcelos, C., Larochelle, H., Dumoulin, V., Roux, N.L., Goroshin, R.: An effective anti-aliasing approach for residual networks. *arXiv preprint arXiv:2011.10675* (2020)
38. Wang, J., Zhang, H.: Bilateral adversarial training: Towards fast training of more robust models against adversarial attacks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6629–6638 (2019)
39. Wang, S., Veldhuis, R., Brune, C., Strisciuglio, N.: A survey on the robustness of computer vision models against common corruptions (2024), <https://arxiv.org/abs/2305.06024>
40. Wang, S., Veldhuis, R., Strisciuglio, N.: The robustness of computer vision models against common corruptions: a survey. *arXiv preprint arXiv:2305.06024* (2023)
41. Yin, D., Gontijo Lopes, R., Shlens, J., Cubuk, E.D., Gilmer, J.: A fourier perspective on model robustness in computer vision. *Advances in Neural Information Processing Systems* **32** (2019)
42. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: *ICCV*. pp. 6023–6032 (2019)
43. Zhang, R.: Making convolutional networks shift-invariant again. In: *International conference on machine learning*. pp. 7324–7334. PMLR (2019)