

NEURO-INSPIRED DEEP NEURAL NETWORKS WITH SPARSE, STRONG ACTIVATIONS

Metehan Cekic*

Can Bakiskan*

Upamanyu Madhow

Department of Electrical and Computer Engineering
 University of California Santa Barbara, Santa Barbara, CA 93106
 {metehancekic, canbakiskan, madhow}@ucsb.edu

ABSTRACT

While end-to-end training of Deep Neural Networks (DNNs) yields state of the art performance in an increasing array of applications, it does not provide insight into, or control over, the features being extracted. We report here on a promising neuro-inspired approach to DNNs with sparser and stronger activations. We use standard stochastic gradient training, supplementing the end-to-end discriminative cost function with layer-wise costs promoting Hebbian (“fire together,” “wire together”) updates for highly active neurons, and anti-Hebbian updates for the remaining neurons. Instead of batch norm, we use divisive normalization of activations (suppressing weak outputs using strong outputs), along with implicit ℓ_2 normalization of neuronal weights. Experiments with standard image classification tasks on CIFAR-10 demonstrate that, relative to baseline end-to-end trained architectures, our proposed architecture (a) leads to sparser activations (with only a slight compromise on accuracy), (b) exhibits more robustness to noise (without being trained on noisy data), (c) exhibits more robustness to adversarial perturbations (without adversarial training).

Index Terms— Interpretable ML, Hebbian learning, neuro inspired, machine learning

1. INTRODUCTION

Since their original breakthrough in image classification performance, DNNs trained with backpropagation have attained outstanding performance in a wide variety of fields [2, 3, 4, 5]. Yet there remain fundamental concerns regarding their lack of interpretability and robustness (e.g, to noise, distribution shifts, and adversarial perturbations). In this paper, we explore the thesis that a first step to alleviating these problems is to exert more control on the features being extracted by DNNs. Specifically, while standard DNNs produce a large fraction of small activations at each layer, we seek architectures which produce a small fraction of strong activations, while continuing to utilize existing network architectures for

* Equal contribution. Code can be found at: [1]. This work was supported in part by the Army Research Office under grant W911NF-19-1-0053, and by the National Science Foundation under grants CIF-1909320 and CIF-2224263.

feedforward inference and existing software infrastructure for stochastic gradient training.

1.1. Approach and Contributions

In order to attain sparse, strong activations at each layer, we employ the following neuro-inspired strategy for modifying standard DNN training and architecture:

Hebbian/anti-Hebbian (HaH) Training: We supplement a standard end-to-end discriminative cost function with layer-wise costs at each layer which promote neurons producing large activations and demote neurons producing smaller activations. The goal is to develop a neuronal basis that produces a distributed sparse code, without requiring a reconstruction cost as in standard sparse coding [6].

Neuronal Competition via Normalization: We further increase sparsity by introducing Divisive Normalization (DN), which enables larger activations to suppress smaller activations. In order to maintain a fair competition among neurons, we introduce Implicit ℓ_2 Normalization of the neuronal weights, so that each activation may be viewed as a geometric projection of the layer input onto the “direction” of the neuron. (Using implicit rather than explicit weight normalization in our inference architecture simplifies training.)

We report on experiments with CIFAR-10 image classification, comparing a baseline VGG-16 network trained end-to-end against the same architecture with HaH training and DN. Both architectures employ implicit weight normalization, which we have verified does not adversely impact accuracy. We demonstrate that the activations in our proposed architecture are indeed more sparse than for the baseline network. Furthermore, robustness against noise and adversarial perturbations is enhanced, without having used noise augmentation or adversarial training.

1.2. Related Work

Hebbian learning has a rich history in artificial neural networks, dating back to the neocognitron [7], and including recent attempts at introducing it into deep architectures [8]. However, to the best of our knowledge, ours is the first paper to clearly demonstrate gains in robustness from its incorporation in DNNs. Divisive normalization is a widely accepted

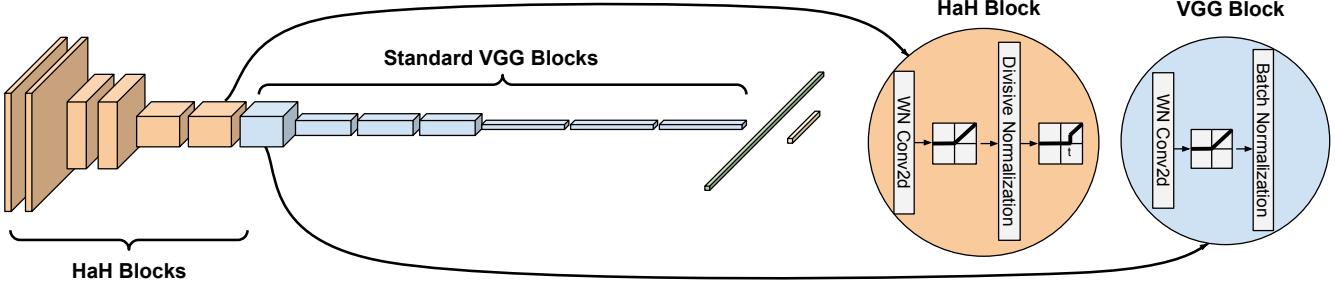


Fig. 1: Our model consists of two different types of blocks: first 6 blocks are Hebbian-anti-Hebbian (HaH) while the rest are regular VGG blocks. HaH blocks use a weight normalized convolutional layer, followed by ReLU, divisive normalization and thresholding. Regular VGG blocks use a weight normalized convolutional layer followed by ReLU and batch norm.

concept in neuroscience [9, 10], and versions of it have been shown to be competitive with other normalization techniques in deep networks [11]. Our novel contribution is in showing that divisive normalization can be engineered to enhance sparsity and robustness. Finally, sparse coding with a reconstruction objective was shown to lead to neuro-plausible outcomes in a groundbreaking paper decades ago [6]. In contrast to the iterative sparse coding and dictionary learning in such an approach, our HaH-based training targets strong sparse activations in a manner amenable to standard stochastic gradient training.

Recent work showing potential robustness gains by directly including known aspects of mammalian vision in DNNs includes [12], which employs Gabor filter blocks and stochasticity, and [13], which employs neural activity measurements from mice for regularization in DNNs. Rather than incorporating specific features from biological vision, we use neuro-inspiration to extract broad principles that can be folded into data-driven learning and inference in DNNs.

2. MODEL

We now describe how we incorporate HaH training and divisive normalization into a standard CNN for image classification. We consider a “classical” CNN for our experiments—VGG-16 [14] applied to CIFAR-10, rather than variants of ResNet [15], because residual connections complicate our interpretation of building models from the bottom-up using HaH learning. Since we wish to build robustness from the bottom up, we modify the first few convolutional blocks to incorporate neuro-inspired principles. We term these modified blocks “HaH blocks.”

Each HaH block employs convolution with implicit weight normalization, followed by ReLU, then divisive normalization, and then thresholding. Implicit weight normalization enables us to interpret the convolution outputs for each filter as projections, and we have verified that employing it in all blocks of a baseline VGG-16 architecture does not adversely impact accuracy (indeed, it slightly improves it). Each standard (non-HaH) block in our architecture therefore also employs convolution with implicit weight normalization,

followed by ReLU, but uses batch norm rather than divisive normalization. Each HaH block contributes a HaH cost for training, so that the overall cost function used for training is the standard discriminative cost and the sum of the HaH costs from the HaH blocks.

We now describe the key components of our architecture, shown in Fig. 1.

2.1. Inference in a HaH block

Implicit weight normalization: Representing the convolution output at a given spatial location from a given filter as a tensor inner product $\langle \cdot, \cdot \rangle$ between the filter weights \mathbf{w} and the input \mathbf{x} , the output of the ReLU unit following the filter is given by

$$y = \text{ReLU} \left(\frac{\langle \mathbf{w}, \mathbf{x} \rangle}{\|\mathbf{w}\|_2} \right) \quad (1)$$

This effectively normalizes the weight tensor of each filter to unit ℓ_2 norm, without actually having to enforce an ℓ_2 norm constraint on the cost.

Divisive normalization: If we have N filters in a given HaH block, let $y_1(loc), \dots, y_N(loc)$ denote the corresponding activations computed as in (Eq. 1) for a given spatial location loc . Let $M(loc) = \frac{1}{N} \sum_{k=1}^N y_k(loc)$ denote the mean of the activations at a given location, and let $M_{max} = \max_{loc} M(loc)$ denote the maximum of this mean over all locations. We normalize each activation using these terms as follows:

$$z_k(loc) = \frac{y_k(loc)}{\sigma M_{max} + (1 - \sigma)M(loc)}, \quad k = 1, \dots, N \quad (2)$$

where $0 \leq \sigma \leq 1$ is a hyperparameter which can be separately tuned for each HaH block. Thus, in addition to creating competition among neurons at a given location by dividing by $M(loc)$, we also include M_{max} in the denominator in order to suppress contributions at locations for which the input is “noise” rather than a strong enough “signal” well-aligned with one or more of the filters. This particular implementation of divisive normalization ensures that the output of a HaH-block is scale-invariant (i.e., we get the same output if we scale the input to the block by any positive scalar).

Adaptive Thresholding: Finally, we ensure that each neuron is producing significant outputs by neuron-specific thresholding after divisive normalization. The output of the k th neuron at location loc is given by

$$o_k(loc) = \begin{cases} z_k(loc) & \text{if } z_k(loc) \geq \tau_k \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where the threshold τ_k is neuron and image specific. For example, we may set τ_k to the 90th percentile of the statistics of $z_k(loc)$ in order to get an activation rate of 10% for each neuron for every image. Another simple choice that works well, but gives higher activation rates, is to set τ_k to the mean of $z_k(loc)$ for each image.

2.2. HaH Training

For an N -neuron HaH block with activations $y_k(loc)$, $k = 1, \dots, N$ at location loc , the Hebbian/anti-Hebbian cost seeks to maximize the average of the top K activations, and to minimize the average of the remaining $N - K$ activations, where K is a hyperparameter. Thus, sorting the activations $\{y_k(loc)\}$ so that $y^{(1)}(loc) \geq y^{(2)}(loc) \geq \dots \geq y^{(N)}(loc)$, the contribution to the HaH cost (to be maximized) is given by

$$L_{block}(loc) = \frac{1}{K} \sum_{k=1}^K y^{(k)}(loc) - \lambda \frac{1}{N-K} \sum_{k=K+1}^N y^{(k)}(loc) \quad (4)$$

where $\lambda \geq 0$ is a hyperparameter determining how much to emphasize the anti-Hebbian component of the adaptation. The overall HaH cost for the block, L_{block} , which we wish to maximize, is simply the mean over all locations and images.

The overall loss function to be minimized is now given by

$$L = L_{disc} - \sum_{\text{HaH blocks}} \alpha_{block} L_{block} \quad (5)$$

where L_{disc} is the standard discriminative loss, and $\{\alpha_{block} \geq 0\}$ are hyper-parameters determining the relative weight of the HaH costs across blocks.

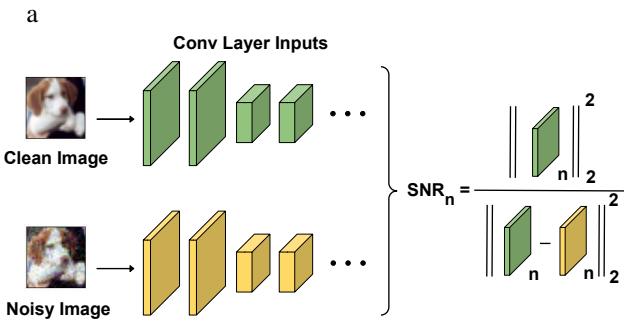


Fig. 3: a: To compute the SNR at the n^{th} block inputs, we divide the ℓ_2 norm of the block input corresponding to clean image by the ℓ_2 norm of the difference of block corresponding to clean and noisy images. **b:** Comparison of SNR values of the block inputs for the standard base model (gray) and ours (red).

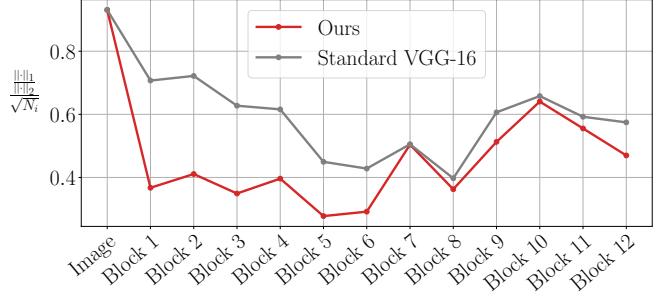


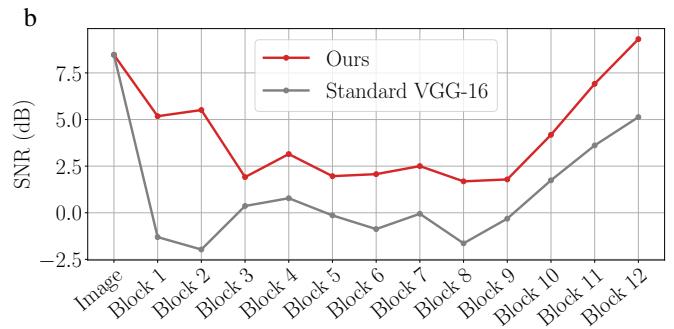
Fig. 2: HaH blocks yield sparser activations than baseline. The measure of sparsity is the Hoyer ratio [16] of ℓ_1 norm to ℓ_2 norm of activations across channels, averaged across spatial locations, and then normalized to lie in $[0,1]$ (lower values correspond to more sparsity).

3. EXPERIMENTS

We consider VGG-16 with the first 6 blocks (each block includes conv, ReLU, batch norm) replaced by HaH blocks (each block includes conv, ReLU, divisive norm, thresholding). In our training, we use Adam optimizer [17] with an initial learning rate of 10^{-3} , multiplied by 0.1 at epoch 60 and again at epoch 80. We train all models for 100 epochs on CIFAR-10. We choose τ_k in Eq. 3 to keep 20% of activations. We use $[4.5 \times 10^{-3}, 2.5 \times 10^{-3}, 1.3 \times 10^{-3}, 1 \times 10^{-3}, 8 \times 10^{-4}, 5 \times 10^{-4}]$ for α in Eq. 5. We use 0.1 for λ and set K to 10% of number of filters in each layer in Eq. 4 and set $\sigma = 0.1$ in Eq. 2. Details about other hyper-parameters can be found in our code repository [1].

Sparser activations: Fig. 2 shows that the activations in these first 6 blocks are indeed more sparse for our architecture than for baseline VGG.

Enhanced robustness to noise: We borrow the concept of signal-to-noise-ratio (SNR) from wireless communication to obtain a block-wise measure of robustness. Let $f_n(x)$ denote the input tensor at block n in response to clean image x , and $f_n(x + w)$ the input tensor when the image is corrupted by



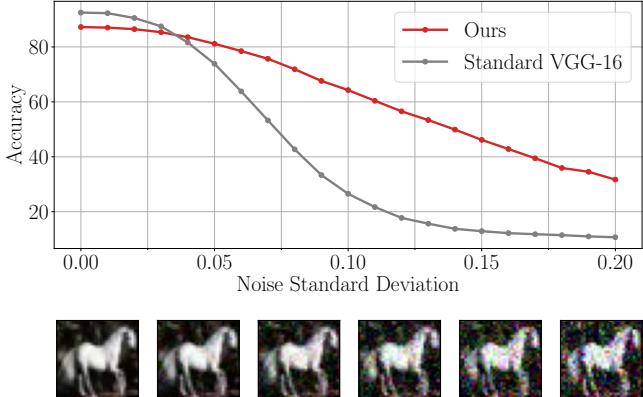


Fig. 4: Comparison of classification accuracies as a function of noise σ . To provide a concrete sense of the impact of noise, noisy images at increasing values of σ are shown below the graph.

noise w . As illustrated in Fig. 3a, we define SNR as

$$\text{SNR}_n = 10 \log_{10} \left(\mathbb{E}_{x \sim \mathcal{D}_{test}} \left[\frac{\|f_n(x)\|_2^2}{\|f_n(x + w) - f_n(x)\|_2^2} \right] \right) dB \quad (6)$$

converting to logarithmic decibel (dB) scale as is common practice. Fig. 3b shows that the SNR for our model comfortably exceeds that of the standard model, especially in the first 6 HaH blocks.

These higher SNR values also translate to gains in accuracy with noisy images: Fig. 4 compares the accuracy of our model and the base model for different levels of Gaussian noise. There are substantial accuracy gains at high noise levels: 64% vs. 26% at a noise standard deviation of 0.1, for example.

Enhanced robustness to adversarial attacks: While we have not trained with adversarial examples, we find that, as expected, the noise rejection capabilities of the HaH blocks also translates into gains in adversarial robustness relative to the baseline VGG model. This holds for state-of-the-art gradient-based attacks [18, 19], as well as AutoAttack, an ensemble of parameter-free attacks suggested by RobustBench [20]. We observe no additional benefit of using gradient-free attacks, and conclude that the robustness provided by our scheme is not because of gradient-masking. Because of space constraints, we only report on results from minimum-norm adversarial attacks and AutoAttack.

Fig. 5 shows that the minimum distortion needed to flip

Table 1: Enhanced accuracy against noise and adversarial attacks

	Clean	Noisy ($\sigma = 0.1$)	Adv (ℓ_∞) ($\epsilon = 2/255$)	Adv (ℓ_2) ($\epsilon = 0.25$)
Standard	92.5%	26.6%	10.4%	13.9%
Ours	87.3%	64.0%	21.5%	27.6%

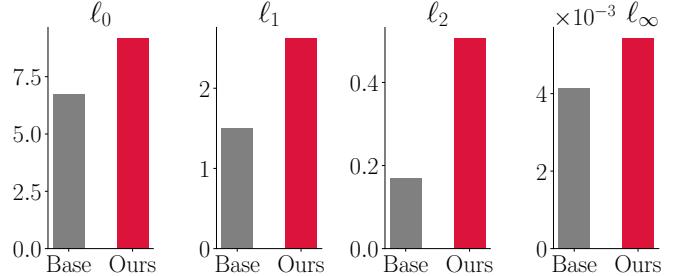


Fig. 5: The average norm of minimum-norm adversarial attacks is higher for our model for all ℓ_p norms considered.

the prediction of our model (computed using the recently proposed fast minimum norm computation method [19]) is higher for our model for all the ℓ_p attacks considered.

We have also obtained substantial gains in adversarial accuracy against all four ℓ_p norm attacks ($p = 0, 1, 2, \infty$) used as benchmarks in adversarial machine learning. Table 1 displays a subset of results demonstrating accuracy gains against noise and adversarial perturbations, at the expense of a slight decrease in clean accuracy.

Ablation: Since we have different components in our HaH blocks, we explore the effectiveness of each component by doing an ablation study. Table 2 summarizes the contribution from each of the components. We see that all of the components (HaH training, divisive normalization, adaptive thresholding) play an important role in obtaining the reported gains in robustness to noise and adversarial attacks.

4. CONCLUSION

Our results indicate the promise of incorporating appropriately engineered neuro-inspired principles into DNN architectures and training. We have chosen supervised learning without augmentation for this initial exposition, but we hope these results motivate further exploration in developing a fundamental understanding of HaH training and inference, as well as in extensive experimentation with a variety of architectures, training techniques (including unsupervised and semi-supervised learning, and data augmentation) and applications.

Table 2: Accuracies for ablation study

	Clean	Noisy ($\sigma = 0.1$)	Adv (ℓ_∞) ($\epsilon = 2/255$)	Adv (ℓ_2) ($\epsilon = 0.25$)
All included	87.3%	64.0%	21.5%	27.6%
No HaH loss	89.7%	50.4%	8.8%	11.7%
Batch norm instead of divisive norm		90.4%	46.7%	12.3%
No thresholding		89.9%	37.5%	3.7%
				2.5%

5. REFERENCES

- [1] “Our github repository,” github.com/metehancekic/SparseStrongActivations.
- [2] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al., “Language models are few-shot learners,” *arXiv preprint arXiv:2005.14165*, 2020.
- [3] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al., “A general reinforcement learning algorithm that masters chess, shogi, and go through self-play,” *Science*, vol. 362, no. 6419, pp. 1140–1144, 2018.
- [4] Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al., “Solving rubik’s cube with a robot hand,” *arXiv preprint arXiv:1910.07113*, 2019.
- [5] Andrew W Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Žídek, Alexander WR Nelson, Alex Bridgland, et al., “Improved protein structure prediction using potentials from deep learning,” *Nature*, vol. 577, no. 7792, pp. 706–710, 2020.
- [6] Bruno A Olshausen and David J Field, “Sparse coding with an overcomplete basis set: A strategy employed by V1?,” *Vision research*, vol. 37, no. 23, 1997.
- [7] Kunihiko Fukushima, Sei Miyake, and Takayuki Ito, “Neocognitron: A neural network model for a mechanism of visual pattern recognition,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, no. 5, pp. 826–834, 1983.
- [8] Giuseppe Amato, Fabio Carrara, Fabrizio Falchi, Claudio Gennaro, and Gabriele Lagani, “Hebbian learning meets deep convolutional neural networks,” in *Image Analysis and Processing – ICIAP 2019*, Elisa Ricci, Samuel Rota Bulò, Cees Snoek, Oswald Lanz, Stefano Messelodi, and Nicu Sebe, Eds., Cham, 2019, pp. 324–334, Springer International Publishing.
- [9] Matteo Carandini and David J Heeger, “Normalization as a canonical neural computation,” *Nature Reviews Neuroscience*, vol. 13, no. 1, pp. 51–62, 2012.
- [10] Max F Burg, Santiago A Cadena, George H Denfield, Edgar Y Walker, Andreas S Tolias, Matthias Bethge, and Alexander S Ecker, “Learning divisive normalization in primary visual cortex,” *PLOS Computational Biology*, vol. 17, no. 6, pp. e1009028, 2021.
- [11] Mengye Ren, Renjie Liao, Raquel Urtasun, Fabian H Sinz, and Richard S Zemel, “Normalizing the normalizers: Comparing and extending network normalization schemes,” *arXiv preprint arXiv:1611.04520*, 2016.
- [12] Joel Dapello, Tiago Marques, Martin Schrimpf, Franziska Geiger, David Cox, and James DiCarlo, “Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations,” *bioRxiv*, 2020.
- [13] Zhe Li, Wieland Brendel, Edgar Walker, Erick Cobos, Taliah Muhammad, Jacob Reimer, Matthias Bethge, Fabian Sinz, Zachary Pitkow, and Andreas Tolias, “Learning from brains how to regularize machines,” *Advances in neural information processing systems*, vol. 32, 2019.
- [14] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [16] Patrik O Hoyer, “Non-negative matrix factorization with sparseness constraints.,” *Journal of machine learning research*, vol. 5, no. 9, 2004.
- [17] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu, “Towards deep learning models resistant to adversarial attacks,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [19] Maura Pintor, Fabio Roli, Wieland Brendel, and Battista Biggio, “Fast minimum-norm adversarial attacks through adaptive norm constraints,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [20] Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein, “Robustbench: a standardized adversarial robustness benchmark,” *arXiv preprint arXiv:2010.09670*, 2020.