
Stochastic Local Winner-Takes-All Networks Enable Profound Adversarial Robustness

Konstantinos P. Panousis
Cyprus University of Technology
Limassol, Cyprus
k.panousis@cut.ac.cy

Sotirios Chatzis
Cyprus University of Technology
Limassol, Cyprus
sotirios.chatzis@cut.ac.cy

Sergios Theodoridis
National and Kapodistrian University of Athens, Greece
Aalborg University, Aalborg, Denmark
stheodor@di.uoa.gr

Abstract

This work explores the potency of stochastic competition-based activations, namely Stochastic Local Winner-Takes-All (LWTA), against powerful (gradient-based) white-box and black-box adversarial attacks; we especially focus on Adversarial Training settings. In our work, we replace the conventional ReLU-based nonlinearities with blocks comprising locally and stochastically competing linear units. The output of each network layer now yields a sparse output, depending on the outcome of winner sampling in each block. We rely on the Variational Bayesian framework for training and inference; we incorporate conventional PGD-based adversarial training arguments to increase the overall adversarial robustness. As we experimentally show, the arising networks yield state-of-the-art robustness against powerful adversarial attacks while retaining very high classification rate in the benign case.

1 Introduction

In this paper, we revisit the novel stochastic formulation of deep networks with LWTA activations (Panousis et al., 2019, 2021), and delve deeper into its potency against adversarial attacks in the context of PGD-based Adversarial Training (AT) (Madry et al., 2017). We evaluate the robustness of the emerging networks against powerful gradient-based white-box as well as black-box adversarial attacks using the well-known AutoAttack (AA) framework (Croce and Hein, 2020). We provide the related source code at: <https://github.com/konpanousis/Adversarial-LWTA-AutoAttack>. We experimentally show that Stochastic LWTA-based networks not only yield state-of-the-art accuracy against all the considered attacks, but do so while retaining very high classification accuracy in the benign case.

2 Stochastic Local Winner-Takes-All (LWTA) Networks

Let us consider an input $\mathbf{x} \in \mathbb{R}^J$ presented to a conventional deep neural network layer comprising K hidden units and weight matrix $\mathbf{W} \in \mathbb{R}^{J \times K}$. Each hidden unit k within the layer performs an inner product computation $h_k = \mathbf{w}_k^T \mathbf{x} = \sum_{j=1}^J w_{jk} \cdot x_j \in \mathbb{R}$; this (usually) passes through a non-linear function $\sigma(\cdot)$. Thus, the final layer output vector $\mathbf{y} \in \mathbb{R}^K$ is formed via the concatenation of the non-linear activation of each individual hidden unit, such that $\mathbf{y} = [y_1, \dots, y_K]$, where $y_k = \sigma(h_k)$.

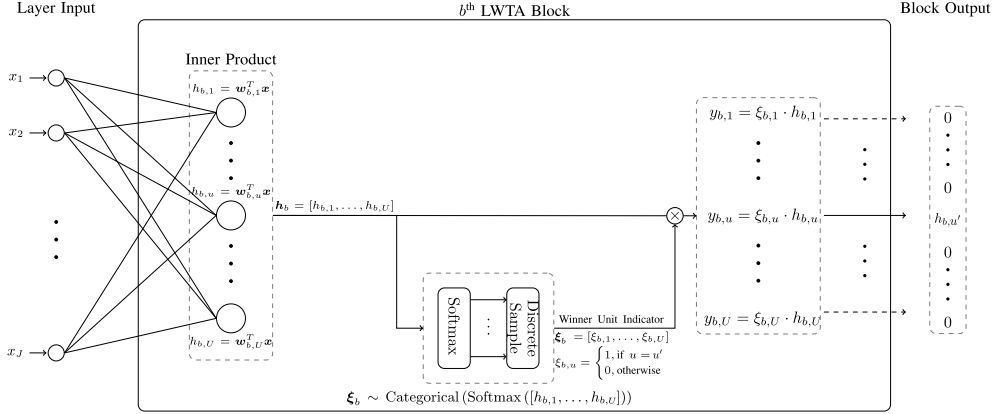


Figure 1: A detailed bisection of the b^{th} Stochastic LWTA block in an LWTA layer. Presented with an input $\mathbf{x} \in \mathbb{R}^J$, each unit $u = 1, \dots, U$ computes its activation $h_{b,u}$ via different weights $\mathbf{w}_{b,u} \in \mathbb{R}^J$, i.e., $h_{b,u} = \mathbf{w}_{b,u}^T \mathbf{x}$. The linear responses of the units are concatenated, such that $\mathbf{h}_b = [h_{b,1}, \dots, h_{b,U}]$, and transformed into probabilities via the softmax operation. Then, a Discrete sample $\boldsymbol{\xi}_b = [\xi_{b,1}, \dots, \xi_{b,U}]$ is drawn; this constitutes an one-hot vector with a single non-zero entry at position u' , denoting the winner unit in the block. This winner unit, u' , passes its linear response to the next layer; the rest pass zero values.

In contrast, in a fully connected LWTA-based network layer, singular nonlinear units are replaced by U linear competing units aggregated together in a (LWTA) block; in the following, we denote with B the number of LWTA blocks in an LWTA-based layer. The associated weights are now arranged as a three dimensional matrix $\mathbf{W} \in \mathbb{R}^{J \times B \times U}$ signifying that the input \mathbf{x} is presented to each block and each unit therein. Specifically, in the LWTA-based framework, the u^{th} competing unit within the b^{th} block computes its activation $h_{b,u}$ via the standard inner product computation $h_{b,u} = \mathbf{w}_{b,u}^T \mathbf{x} = \sum_{j=1}^J w_{j,b,u} \cdot x_j \in \mathbb{R}$; then, competition takes place among the units in the block.

The underlying principle is that out of the U units in an LWTA block, *only one* can be the *winner*; this unit gets to pass its *linear activation* to the next layer, while the rest output zero values. Thus, the output of an LWTA layer $\mathbf{y} \in \mathbb{R}^{B \cdot U}$ is composed of B subvectors $\mathbf{y}_b \in \mathbb{R}^U$, one for each LWTA block and each with a single non-zero entry. It is apparent that this competition process results in a *sparse representation*, since all units, *except one* in each block, produce a zero output. In related literature, the competition procedure is deterministic, i.e., the winner unit is the one with the highest activation. However, stochastic competition principles have been recently proposed in Panousis et al. (2019, 2021); Voskou et al. (2021).

In this context, to encode the winner unit in each of the B LWTA blocks that constitute a stochastic LWTA layer, we introduce an appropriate set of discrete latent indicator vectors $\boldsymbol{\xi} \in \text{one_hot}(U)^B$. This vector comprises B component subvectors, where each component entails *exactly one non-zero value* at the *index position* that corresponds to the *winner unit* in each respective LWTA block.

Thus, the output \mathbf{y} of a stochastic LWTA layer's $(b, u)^{\text{th}}$ component $y_{b,u}$ is defined as:

$$y_{b,u} = \xi_{b,u} \sum_{j=1}^J w_{j,b,u} \cdot x_j \in \mathbb{R} \quad (1)$$

where $\xi_{b,u}$ denotes the u^{th} component of $\boldsymbol{\xi}_b$, and $\boldsymbol{\xi}_b \in \text{one_hot}(U)$ holds the b^{th} subvector of $\boldsymbol{\xi}$.

We postulate that the latent winner indicator variables $\boldsymbol{\xi}_b, \forall b$ in Eq.(1) are drawn from a data-driven Categorical distribution with probabilities proportional to the intermediate linear computations that each unit performs. Therefore, the higher the linear response of a particular unit in a particular block,

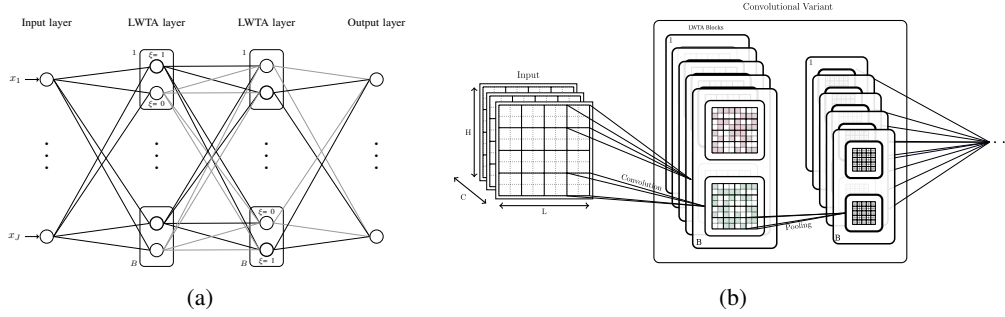


Figure 2: (a) A graphical representation of our competition-based modeling approach. Rectangles denote LWTA blocks, and circles the competing units therein. The winner units are denoted with bold contours ($\xi = 1$). (b) The convolutional LWTA variant. Competition takes place *position-wise* among the feature maps comprising a kernel. For each position, only the winner feature map contains a non-zero entry; for the rest feature maps in the kernel, the value at said position is zero.

the higher its probability of it being the winner in said block; this yields:

$$q(\xi_b) = \text{Categorical}\left(\xi_b \mid \text{softmax}\left(\sum_{j=1}^J [w_{j,b,u}]_{u=1}^U \cdot x_j\right)\right) \quad (2)$$

where $[w_{j,b,u}]_{u=1}^U$ denotes the vector concatenation of the set $\{w_{j,b,u}\}_{u=1}^U$. A graphical illustration of the proposed stochastic LWTA block is depicted in Fig. 1. Each stochastic LWTA layer comprises multiple such LWTA blocks, as illustrated in Fig. 2a. Note that stochastically selecting the winner of each LWTA block in each layer introduces stochasticity to the network activations. Presented with the same input, different *subnetworks* may be activated and a different *subpath* is followed to the output, as a result of winner sampling.

Further, to also accommodate in our study architectures based on the convolutional operation, we adopt the corresponding variant of the Stochastic LWTA activation in Panousis et al. (2019). We assume an input tensor $\mathbf{X} \in \mathbb{R}^{H \times L \times C}$ and define a set of kernels, each with weights $\mathbf{W}_b \in \mathbb{R}^{h \times l \times C \times U}$, where h, l, C, U are the kernel height, length, channels and competing feature maps, and $b = 1, \dots, B$. Analogously to the grouping of linear units in dense layers, in this case, local competition is performed among feature maps on a *position-wise* basis. Each kernel is treated as an LWTA block with competing feature maps; each layer comprises B kernels. Specifically, each feature map $u = 1, \dots, U$ in the b^{th} LWTA block of a convolutional LWTA layer computes:

$$\mathbf{H}_{b,u} = \mathbf{W}_{b,u} \star \mathbf{X} \in \mathbb{R}^{H \times L} \quad (3)$$

Then, competition takes place among competing feature maps on a *position-wise* basis. The competitive random sampling procedure reads:

$$q(\xi_{b,h',l'}) = \text{Categorical}\left(\xi_{b,h',l'} \mid \text{softmax}([\mathbf{H}_{b,1,h',l'}, \dots, \mathbf{H}_{b,U,h',l'}])\right), \forall h', l' \quad (4)$$

In each kernel, $b = 1, \dots, B$, for each position, only the winner feature map contains a non-zero entry; all the rest feature maps contain zero values at this position. This yields sparse feature maps with mutually exclusive active positions.

Thus, at a given layer of the proposed convolutional variant, the output $\mathbf{Y} \in \mathbb{R}^{H \times L \times B \cdot U}$ is obtained via concatenation of the subtensors $\mathbf{Y}_{b,u}$ that read:

$$\mathbf{Y}_{b,u} = \Xi_{b,u}(\mathbf{W}_{b,u} \star \mathbf{X}), \forall b, u \quad (5)$$

where $\Xi_{b,u} = [\xi_{b,u,h',l'}]_{h',l'=1}^{H,L}$. The corresponding illustration of the proposed stochastic convolutional LWTA block is depicted in Fig. 3. Convolutional Stochastic LWTA-based layers comprise multiple such blocks as shown in Fig. 2b.

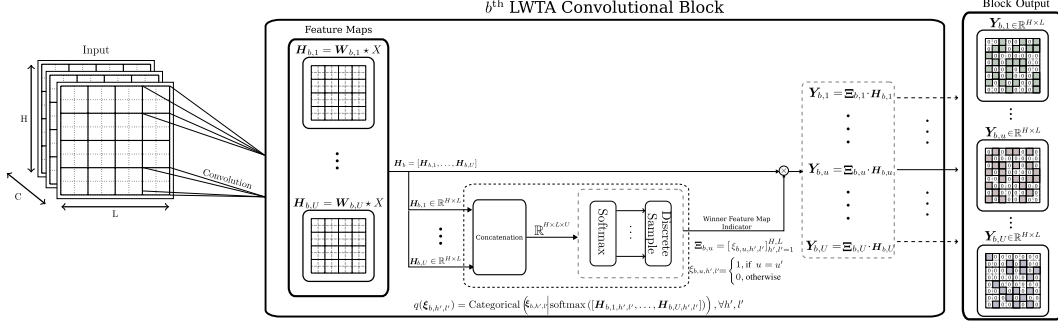


Figure 3: A detailed bisection of the b^{th} convolutional stochastic LWTA block. Presented with an input $\mathbf{X} \in \mathbb{R}^{H \times L \times C}$, competition now takes place among feature maps on a position-specific basis. Only the winner feature map contains a non-zero entry in a specific position. This leads to sparse feature maps, each comprising uniquely position-wise activated pixels.

2.1 Training

Since a network composed of such stochastic LWTA layers entails latent variables ξ , we resort to a Bayesian treatment to perform effective parameter estimation. To this end, we turn to a stochastic gradient variational Bayes treatment (Kingma and Welling, 2014) for scalability. The resulting objective takes the form of an evidence lower-bound (ELBO) as described next.

Considering data $\mathcal{D} = \{X_i, Y_i\}_{i=1}^N$, we define the categorical cross-entropy between the data labels Y_i and the class probabilities $f(X_i; \hat{\xi})$, generated by the penultimate Softmax layer of a Stochastic LWTA-based network, as $\text{CE}(Y_i, f(X_i; \hat{\xi}))$. Here, $\hat{\xi}$ denotes *sample instances* of all the latent variables, ξ , in all layers. We stress that the output class probabilities depend on the winner selection process in each layer, which is stochastic.

This way, the Evidence Lower Bound (ELBO) reads:

$$\mathcal{L} = - \sum_{X_i, Y_i \in \mathcal{D}} \text{CE}(Y_i, f(X_i; \hat{\xi})) - \text{KL}[q(\xi) || p(\xi)] \quad (6)$$

For simplicity, and without loss of generality, we consider a symmetric Categorical distribution for the latent variable indicators ξ ; hence, $p(\xi_b) = \text{Categorical}(1/U) \forall b$ for dense layers, and $p(\xi_{b,h',l'}) = \text{Categorical}(1/U) \forall b, h', l'$ for convolutional ones. In our work, we perform Monte-Carlo sampling using a single reparameterized sample for each of the corresponding latent variables. These are obtained via the reparameterization trick of the continuous relaxation of the Categorical distribution (Maddison et al., 2017; Jang et al., 2017) as described next. We focus on the reparameterization trick for the dense case; the convolutional case is analogous.

Let $\tilde{\xi}$ denote the probabilities of $q(\xi)$ (Eqs. (2) and (4)). Then, the samples $\hat{\xi}$ are expressed as:

$$\hat{\xi}_{b,u} = \text{Softmax}((\log \tilde{\xi}_{b,u} + g_{b,u})/\tau), \quad \forall b = 1, \dots, B, u = 1, \dots, U \quad (7)$$

where $g_{b,u} = -\log(-\log V_{b,u})$, $V_{b,u} \sim \text{Uniform}(0, 1)$, and $\tau \in (0, \infty)$ is a temperature factor, controlling how “closely” the continuous relaxation approximates the Categorical distribution.

On this basis, we can write the KL divergence term present in Eq. (6) as:

$$\begin{aligned} \text{KL}[q(\xi_b) || p(\xi_b)] &= \mathbb{E}_{q(\xi_b)}[\log q(\xi_b) - \log p(\xi_b)] \\ &\approx \log q(\hat{\xi}_b) - \log p(\hat{\xi}_b), \quad \forall b \end{aligned} \quad (8)$$

Hence, the final ELBO expression yields:

$$\mathcal{L} = - \sum_{X_i, Y_i \in \mathcal{D}} \text{CE}(Y_i, f(X_i; \hat{\xi})) - \sum_b (\log q(\hat{\xi}_b) - \log p(\hat{\xi}_b)) \quad (9)$$

2.2 Prediction

At prediction time, we directly draw L samples from the trained posteriors $q(\xi)$ in order to determine the winning units in each block of the network. As mentioned before, each time we sample for

Table 1: Natural and Robust accuracy under a conventional PGD attack with 20 steps and 0.007 step-size using WideResNet-34 models with different widen factors. We use the same PGD-based Adversarial Training scheme for all models (Madry et al., 2017).

Adversarial Training-PGD				
Widen Factor	Natural Accuracy (%)		Robust Accuracy (%)	
	Baseline	Stochastic LWTA	Baseline	Stochastic LWTA
1	74.04	87.0	49.24	81.87
5	83.95	91.88	54.36	83.4
10	85.41	92.26	55.78	84.3

the same input, a different *subpath* is followed. This is a key aspect that *stochastically* alters the information flow in the network and *obstructs* an adversary from attacking the model.

The sampling process results in a set of L output logits of the network, which we can average to obtain the final prediction:

$$f(X_i; \hat{\xi}) = \frac{1}{L} \sum_{l=1}^L f(X_i; \hat{\xi}^l) \quad (10)$$

where $\hat{\xi}^l$ denotes a sample drawn from $q(\xi)$.

3 Experimental Results

We investigate the potency of LWTA-based networks against adversarial attacks under an Adversarial Training regime; we employ a PGD adversary (Madry et al., 2017). To this end, we use the well-known WideResNet-34 (Zagoruyko and Komodakis, 2016) architecture, considering three different widen factors: 1, 5, and 10; we focus on the CIFAR-10 dataset and adopt experimental settings similar to Wu et al. (2021). We use a batch-size of 128 and an initial learning rate of 0.1; we halve the learning rate at every epoch after the 75th epoch. We use a single sample $L = 1$ for prediction. All experiments were performed using a single NVIDIA Quadro P6000.

For evaluating the robustness of the proposed structure, we initially consider the conventional PGD attack with 20 steps, step size 0.007 and $\epsilon = 8/255$. In Table 1, we compare the robustness of LWTA-based WideResNet networks against the baseline results of Wu et al. (2021). As we observe, our Stochastic LWTA-based networks yield significant improvements in robustness under a traditional PGD attack; they retain extremely high natural accuracy (up to $\approx 13\%$ better), while exhibiting a staggering, up to $\approx 32.6\%$, difference in robust accuracy compared to the *exact same architectures* employing the conventional ReLU-based nonlinearities and trained in the *exact same fashion*.

Further, and to ensure that our approach does not cause the well-known obfuscated gradient problem (Athalye et al., 2018), we resort to stronger parameter-free attacks using the newly introduced AutoAttack (AA) framework (Croce and Hein, 2020). AA comprises an ensemble of four powerful white-box and black-box attacks, e.g., the commonly employed A-PGD attack; this is a step-free variant of the standard PGD attack (Madry et al., 2017), which avoids the complexity and ambiguity of step-size selection. In addition, for the entailed L_∞ attack, we use the common $\epsilon = 8/255$ value. Thus, in Table 2, we compare the LWTA-based networks to several recent state-of-the-art approaches evaluated on AA¹. The reported accuracies correspond to the final reported robust accuracy of the methods after sequentially performing all the considered AA attacks. Once again, we observe that the proposed networks yield state-of-the-art robustness against all SOTA methods, with an improvement of $\approx 16.72\%$, even when compared with methods that employ substantial data augmentation to increase robustness, e.g. Goyal et al. (2021). These results vouch for the potency of Stochastic LWTA networks in adversarial settings.

Finally, since our considered networks consist of stochastic components, i.e. the competitive random sampling procedure to determine the winner in each LWTA block, the output of the classifier might change at each iteration; this obstructs the attacker from successfully altering the final decision. To counter such randomness in the involved computations, Croce and Hein (2020) combine the

¹<https://github.com/fra31/auto-attack>

Table 2: Robust Accuracy (%) comparison under the AutoAttack framework. † denotes models that are trained with additional unlabeled data. The AutoAttack performance corresponds to the final robust accuracy after employing all the attacks in AA. Results directly from the AA leaderboard.

Method	AutoAttack
TRADES(Zhang et al., 2019)	53.08
Early-Stop (Rice et al., 2020)	53.42
FAT (Zhang et al., 2020)	53.51
HE (Pang et al., 2020)	53.74
WAR (Wu et al., 2021)	54.73
<hr/>	
Pre-training (Hendrycks et al., 2019)†	54.92
MART (Wang et al., 2020)†	56.29
HYDRA (Schwag et al., 2020)†	57.14
RST (Carmon et al., 2019)†	59.53
Gowal et al. (2021)†	65.88
WAR (Wu et al., 2021)†	61.84
<hr/>	
Ours (Stochastic-LWTA/PGD/WideResNet-34-1)	74.71
Ours (Stochastic-LWTA/PGD/WideResNet-34-5)	81.22
Ours (Stochastic-LWTA/PGD/WideResNet-34-10)	82.60

APGD attack with an averaging procedure of 20 computations of the gradient at the same point. This technique is known as Expectation over Transformation (EoT) (Athalye et al., 2018). Thus, we use AA combined with EoT for further performance evaluation of the proposed LWTA-based networks. The corresponding results are presented in Table 3. As we observe, all of the considered networks retain state-of-the-art robustness against the powerful AA & EoT attacks. This conspicuously supports the usefulness of Stochastic LWTA activations towards adversarial robustness.

Table 3: Robustness against AA combined with 20 iterations of EoT. APGD-DLR corresponds to the APGD attack, using a different loss, i.e., the Difference of Logits Ratio (Croce and Hein, 2020).

Widen Factor	Nat. Acc.	APGD	APGD-DLR
1	87.00	79.67	76.15
5	91.88	81.67	77.65
10	92.26	82.55	79.00

4 Conclusions

In this work, we explored the potency of Stochastic LWTA-based networks against powerful white-box and black-box attacks. The experimental results vouch for the efficacy of the arising networks, yielding state-of-the-art robustness in all the experimental settings. We obtained an immense improvement in robustness compared to the second best performing alternative, which notably relies on substantial data augmentation. A potentially key principle towards adversarial robustness may be the stochastic alteration of the information flow in Stochastic LWTA-based networks; this, arises from the considered data-driven winner selection mechanism in each LWTA block. Different subpaths, stochastically emerging even for the same input, essentially obstruct the adversary from successfully attacking the model. Further evaluation against stochasticity countermeasures, i.e., Expectation over Transformation (EoT), further validate our findings, as the induced decrease in the final robustness was negligible.

Acknowledgements

This work has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 872139, project aiD.

References

- Athalye, A., Carlini, N., and Wagner, D. A. (2018). Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proc. ICML*.
- Carmon, Y., Raghunathan, A., Schmidt, L., Duchi, J. C., and Liang, P. S. (2019). Unlabeled data improves adversarial robustness. In *Proc. NIPS*.
- Croce, F. and Hein, M. (2020). Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proc. ICML*.
- Gowal, S., Qin, C., Uesato, J., Mann, T., and Kohli, P. (2021). Uncovering the limits of adversarial training against norm-bounded adversarial examples. *arXiv preprint arXiv:2010.03593*.
- Hendrycks, D., Lee, K., and Mazeika, M. (2019). Using pre-training can improve model robustness and uncertainty. In *Proc. ICML*.
- Jang, E., Gu, S., and Poole, B. (2017). Categorical reparametrization with gumbel-softmax. In *Proc. ICLR*.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. In *Proc. ICLR*.
- Maddison, C. J., Mnih, A., and Teh, Y. W. (2017). The concrete distribution: A continuous relaxation of discrete random variables. In *Proc. ICLR*.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2017). Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Pang, T., Yang, X., Dong, Y., Xu, K., Zhu, J., and Su, H. (2020). Boosting adversarial training with hypersphere embedding. In *Proc. NIPS*.
- Panousis, K., Chatzis, S., Alexos, A., and Theodoridis, S. (2021). Local competition and stochasticity for adversarial robustness in deep learning. In *Proc. AISTATS*.
- Panousis, K., Chatzis, S., and Theodoridis, S. (2019). Nonparametric Bayesian deep networks with local competition. In *Proc. ICML*.
- Rice, L., Wong, E., and Kolter, J. Z. (2020). Overfitting in adversarially robust deep learning. *arXiv preprint arXiv:2002.11569*.
- Sehwag, V., Wang, S., Mittal, P., and Jana, S. (2020). Hydra: Pruning adversarially robust neural networks. In *Proc. NIPS*.
- Voskou, A., Panousis, K. P., Kosmopoulos, D., Metaxas, D. N., and Chatzis, S. (2021). Stochastic transformer networks with linear competing units: Application to end-to-end sl translation. In *In Proc. ICCV*.
- Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., and Gu, Q. (2020). Improving adversarial robustness requires revisiting misclassified examples. In *Proc. ICLR*.
- Wu, B., Chen, J., Cai, D., He, X., and Gu, Q. (2021). Do wider neural networks really help adversarial robustness? In *In Proc. NIPS*.
- Zagoruyko, S. and Komodakis, N. (2016). Wide residual networks. In *Proc. BMVC*.
- Zhang, H., Chen, H., Song, Z., Boning, D., Dhillon, I. S., and Hsieh, C.-J. (2019). The limitations of adversarial training and the blind-spot attack. *arXiv preprint arXiv:1901.04684*.
- Zhang, J., Xu, X., Han, B., Niu, G., Cui, L., Sugiyama, M., and Kankanhalli, M. (2020). Attacks which do not kill training make adversarial learning stronger. In *Proc. ICML*.