

Training Robust Deep Neural Networks via Adversarial Noise Propagation

Aishan Liu, Xianglong Liu*, Chongzhi Zhang[†], Hang Yu[†], Qiang Liu, Dacheng Tao

Abstract—In practice, deep neural networks have been found to be vulnerable to various types of noise, such as adversarial examples and corruption. Various adversarial defense methods have accordingly been developed to improve adversarial robustness for deep models. However, simply training on data mixed with adversarial examples, most of these models still fail to defend against the generalized types of noise. Motivated by the fact that hidden layers play a highly important role in maintaining a robust model, this paper proposes a simple yet powerful training algorithm, named *Adversarial Noise Propagation* (ANP), which injects noise into the hidden layers in a layer-wise manner. ANP can be implemented efficiently by exploiting the nature of the backward-forward training style. Through thorough investigations, we determine that different hidden layers make different contributions to model robustness and clean accuracy, while shallow layers are comparatively more critical than deep layers. Moreover, our framework can be easily combined with other adversarial training methods to further improve model robustness by exploiting the potential of hidden layers. Extensive experiments on MNIST, CIFAR-10, CIFAR-10-C, CIFAR-10-P, and ImageNet demonstrate that ANP enables the strong robustness for deep models against both adversarial and corrupted ones, and also significantly outperforms various adversarial defense methods.

Index Terms—Adversarial Examples, Corruption, Model Robustness, Deep Neural Networks.

I. INTRODUCTION

RECENT advances in deep learning have achieved remarkable successes in various challenging tasks, including computer vision [1]–[3], natural language processing [4], [5] and speech [6], [7]. In practice, deep learning has been routinely applied on large-scale datasets containing data collected from daily life, which inevitably contain large amounts of noise including adversarial examples and corruption [8], [9]. Unfortunately, while such noise is imperceptible to human beings, it is highly misleading to deep neural networks, which presents potential security threats for practical machine learning applications in both the digital and physical world [10]–[14].

Over the past few years, the training of robust deep neural networks against noise has attracted significant attention. The

most successful strategy has tended to involve developing different adversarial defense strategies [15]–[21] against adversarial examples. A large proportion of these defensive methodologies attempt to supply adversaries with non-computable gradients to avoid common gradient-based adversarial attacks. While they can obtain a certain degree of stabilization for DNNs in adversarial setting, these methods can be easily circumvented by constructing a function to approximate the non-differentiable layer on the backward pass. [22]. By contrast, *adversarial training* [8] can still mount a appropriate defense by augmenting training data with adversarial examples. However, while adversarially trained deep models [23] are robust to some single-step attacks, they remain vulnerable to iterative attacks. More recently, [24] proposed to improve adversarial robustness by integrating an adversarial perturbation-based regularizer into the classification objective.

In addition to adversarial examples [15], [25], corruption such as snow and blur also frequently occur in the real world, which also presents critical challenges for the building of strong deep learning models. [26] found that deep learning models behave distinctly subhuman to input images with Gaussian noise. [27] proposed stability training to improve model robustness against noise, but this was confined only to JPEG compression. More recently, [28] was the first to establish a rigorous benchmark to evaluate the robustness of image classifier to 75 different types of corruption.

Despite the progress already achieved, few studies have been devoted to improving model robustness against corruption. Most existing adversarial defense methods remain vulnerable to the generalized noise, which is mainly due to the use of a simple training paradigm that adds adversarial noise to the input data. It is well understood that, in deep neural networks, the influence of invasive noise on prediction can be observed directly in the form of sharp variations in the middle feature maps during forward propagation [25], [29]. Prior studies [25] have proven that hidden layers play a very important role in maintaining a robust model. In [25], adversarially robust models can be created by constraining the Lipschitz constant between hidden layers (e.g., linear, conv layers) to be smaller than 1. [30] generates robust features with the help of the penultimate layer of a classifier, which are in turn used to help with training a robust model. More recently, [31] also noted the importance of robust feature representation and used this approach to deal with many computer vision tasks. This indicates that the noise resistance of hidden layers plays a highly important role in training robust models. Motivated by this fact, we aim to build strong deep models by obtaining robust hidden representations during training.

A. Liu, X. Liu, C. Zhang, and H. Yu are with the State Key Lab of Software Development Environment, Beihang University, Beijing 100191, China. X. Liu is also with Beijing Advanced Innovation Center for Big Data-Based Precision Medicine (*Corresponding author: Xianglong Liu, xliu@nlsde.buaa.edu.cn), ([†] indicates equal contributions, authors in alphabetical order)

Q. Liu is with the Department of Computer Science, University of Texas at Austin, Austin, TX 78712, USA

D. Tao is with the UBTECH Sydney Artificial Intelligence Centre and the School of Information Technologies, Faculty of Engineering and Information Technologies, The University of Sydney, Darlingtown, NSW 2008, Australia

Accordingly, in this paper, we propose a simple but very powerful training algorithm named *Adversarial Noise Propagation* (ANP), designed to enable strong robustness for deep models against generalized noise (including both adversarial examples and corruption). Rather than perturbing the inputs, as in traditional adversarial defense methods, our method injects adversarial noise into the *hidden layers* of neural networks during training. This can be accomplished efficiently by a simple modification of the standard backward-forward propagation, without introducing significant computational overhead. Since the adversarial perturbations for the hidden layers are considered and added during training, models trained using ANP are expected to be more robust against more types of noise. Moreover, ANP takes advantage of hidden layers and is orthogonal to most adversarial training methods; thus, they could be combined together to build stronger models. To facilitate further understanding of the contributions of hidden layers, we provide insights into their behaviors during training from the perspectives of hidden representation insensitivity and human visual perception alignment. Our code has also been released at <https://github.com/AnonymousCodeRepo/ANP>.

Extensive experiments in both black-box and white-box settings on MNIST, CIFAR-10 and ImageNet are conducted to demonstrate that ANP is able to achieve excellent results compared to the common adversarial defense algorithms, including the adversarial defense methods won at *NeurIPS 2017*. Meanwhile, experiments on CIFAR-10-C and CIFAR-10-P [28] prove that ANP can enhance strong corruption robustness for deep models. By investigating the contributions of hidden layers, we found that we only need to inject noise into shallow layers, which are more critical to model robustness. In addition, we can further improve model robustness by combining ANP with other adversarial training methods in different settings.

II. RELATED WORKS

A number of works have been proposed to improve adversarial robustness, including those utilizing network distillation [32], input reconstruction [33], [34], gradient masking [15], [35], etc. Among these, adversarial training [9], [36], [37] in particular has been widely studied in the adversarial learning literature and determined to be the most effective method for improving model robustness against adversarial examples. The concept of adversarial training, which was first proposed by [9], involves feeding model with adversarial examples to facilitate data augmentation during training:

$$\min_{\theta} \rho(\theta), \quad \rho(\theta) = \mathbb{E}_{(x,y) \sim D} \left[\max_{r \in S} L(y, F(x + r, \theta)) \right],$$

where r is a small ball that controls the magnitude of the noise.

These authors also proposed a gradient-based attack method, called FGSM, to generate adversarial examples for adversarial training. To further improve the effectiveness of adversarial training, PGD-based adversarial training [36] was subsequently introduced to adversarially train deep models via PGD

attack [36]. It can be readily seen that both FGSM- and PGD-based adversarial training only consider adversarial noise in the input data; thus they can be considered a special case of ANP in which only adversarial noise in the 0-th hidden layer is considered.

To improve the diversity of adversarial examples, ensemble adversarial training [38] was devised; this approach employs a set of models F to generate different adversarial examples for data augmentation. In ANP, various types of adversarial noise (with different noise sizes, iteration steps, and noise magnitudes) are generated in each layer, which could increase the diversity and complexity of adversarial noise injected into the model.

Layer-wise adversarial training [39] was proposed as a regularization mechanism to prevent overfitting. During training, adversarial gradients in the current mini-batch are computed with reference to the previous mini-batch. Their method is primarily designed to improve the model's generalization ability, which is a different goal from ours. For ANP, the adversarial noise for one specified mini-batch is computed only in the same mini-batch. We believe that the high correlation of adversarial gradients within a single mini-batch will lead to better performance. We therefore introduce the progressive backward-forward propagation in order to fully utilize the information contained in every group of mini-batch data. We further prove that only shallow layers should be considered during training to obtain a robust model.

In summary, our ANP can be considered as a more general framework that considers the contribution of hidden layers and observes adversarial noise in more flexible ways.

III. PRELIMINARIES

A. Terminology and Notation

Given a dataset with feature vector $x \in \mathcal{X}$ and label $y \in \mathcal{Y}$, the deep supervised learning model aims to learn a mapping or classification function $f: \mathcal{X} \rightarrow \mathcal{Y}$. More specifically, in this paper, we consider the visual recognition problem.

B. Adversarial Example

Given a network f_{θ} and an input x with ground truth label y , an adversarial example x^{adv} is an input such that

$$f_{\theta}(x^{adv}) \neq y \quad \text{s.t.} \quad \|x - x^{adv}\| < \epsilon,$$

where $\|\cdot\|$ is a distance metric used to quantify that the semantic distance between the two inputs x and x^{adv} is small enough. By contrast, the adversarial example makes the model predict the wrong label: namely, $f_{\theta}(x^{adv}) \neq y$.

C. Corruption

Image corruption refers to random variations of the brightness or color information in images, such as Gaussian noise, defocus blur, brightness, etc. Supposing, we have a set of corruption functions C in which each $c(x)$ performs a different kind of corruption function. Thus, average-case model

performance on small, general, classifier-agnostic corruption can be used to define model corruption robustness as follows

$$\mathbb{E}_{c \sim C}[P_{(x,y) \sim D}(f(c(x)) = y)].$$

In summary, corruption robustness measures the classifier's average-case performance on corruption C , while adversarial robustness measures the worst-case performance on small, additive, classifier-tailored perturbations.

IV. PROPOSED APPROACH

In this section, we introduce our proposed approach, *Adversarial Noise Propagation* (ANP).

A. Adversarial Formulation

In a deep neural network, the sharp variations in the hidden representation will propagate through hidden layers, leading to undesired predictions. Therefore, model robustness can be greatly improved by the noise insensitivity and guaranteeing the stable behavior in hidden layers. To obtain a model that is robust to small degree of noise, we try to improve the layer-wise noise resistance ability in the deep learning models.

Instead of manipulating only the input layer, as in traditional adversarial defense methods, we instead attempt to add adversarial noise to each hidden layer of a deep learning model by propagating backward from the adversarial loss during training. This strategy forces the model to minimize the model loss for a specific task, exploiting the opposite adversarial noise in each hidden layer that expects to maximize the loss. Subsequently, the learned parameters in each layer enable the model to maintain consistent and stable predictions for the clean instance and its noisy surrogates distributed in the neighborhood, thus building strong robustness for deep models.

From a formal perspective, let us first recall that a deep neural network $y = F(x; \theta)$ is a composition of a number of nonlinear maps, each of which corresponds to a layer:

$$z_{m+1} = f(z_m; \theta), \quad m = 0, \dots, M,$$

where $z_0 = x$ denotes the input, $z_M = y$ the output, and z_m the output of the m -th hidden layer. Moreover, θ collects the weights of the network.

In our framework, we introduce an adversarial noise r_m on the hidden state z_m at each layer, as follows:

$$z_{m+1} = f(z_m + r_m, \theta), \quad m = 0, \dots, M.$$

We use $\tilde{y} = F(x; \theta, r)$ to denote the final network output with the injected noise $r = \{r_m\}_{m=0}^M$ at all layers. We then learn the network parameter θ by minimizing the following adversarial loss:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim D} \left[\max_r (L(y, F(x; \theta, r)) - \eta \cdot \|r\|_p) \right], \quad (1)$$

where, for each data point (x, y) , we search for an adversarial noise r , subject to an ℓ_p norm constraint. The coefficient η controls the magnitude of the adversarial noise.

B. Noise Propagation

Evidently, the key challenge of this framework is solving the inner maximization for individual input data points. This is efficiently addressed in our *Adversarial Noise Propagation* (ANP) method by performing gradient descent on r , utilizing a natural backward-forward style training that adds minimum computational cost over the standard back-propagation. This induces noise propagation across layers and noise injection into hidden layers during the backward-forward training.

More specifically, in each iteration, we first select a mini-batch of training data. For each data point (x, y) in the mini-batch, we approximate the inner optimization by running k steps of gradient descent to utilize the most information possible in each mini-batch. After initializing at $r^{m,0} = 0$ for all $m = 0, \dots, M$, we have the adversarial gradient for the m -th hidden layer z^m :

$$g^{m,t} = \nabla_{r^m} L(y, F(x; \theta, r^t)) = \nabla_{z^m} L(y, F(x; \theta, r^t)),$$

from which we can make the key observation that the gradient of the adversarial noise r^m is equal to the gradient of the hidden states z^m . Evidently, this has already been calculated in the standard backward propagation, and thus introduces no additionally computational cost.

More specifically, the noise gradient is calculated recursively during the standard backward propagation, as follows:

$$\begin{aligned} g^{m-1,t} &= \frac{\partial L(y, F(x; \theta, r))}{\partial z^{m-1,t}} \\ &= \frac{\partial L(y, F(x; \theta, r))}{\partial z^{m,t}} \cdot \frac{\partial z^{m,t}}{\partial z^{m-1,t}} \\ &= g^{m,t} \cdot \frac{\partial z^{m,t}}{\partial z^{m-1,t}}. \end{aligned}$$

Performing gradient descent on r^m yields a further update, namely:

$$r^{m,t+1} \leftarrow (1 - \eta) r^{m,t} + \frac{\varepsilon}{k} \frac{g^{m,t}}{\|g^{m,t}\|_p}, \quad (2)$$

where ε is the step size; moreover, it is normalized by the number of steps k , so that the overall magnitude of update for k steps equals ε . Practically speaking, η and ε control the contribution made by the previous noise value $r^{m,t}$ and the gradient $g^{m,t}$ to the new noise $r^{m,t+1}$, respectively.

1) *Learning with a Noise Register*: In practice, ANP is implemented efficiently through a backward-forward training process. In more detail, for each mini-batch training, we store the corresponding adversarial noise r for each hidden layer during the backward propagation process; during the forward propagation process, moreover, we simply fetch r as the noise and add it to the input in the corresponding hidden layer following the affine transformation and before the activation function. This training procedure introduces no substantial increase in either computation and memory consumption, except that we need to integrate a register S into each neuron in order to store the adversarial noise r (as illustrated in Figure 1). At the time of inference, S can be discarded, meaning that it has no influence on the model complexity. For example, for the m -th hidden layer, during the forward propagation process at the t -th iteration, we first compute the affine transformation

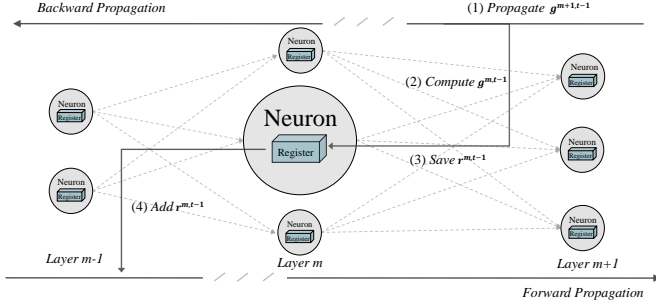


Fig. 1. Adversarial noise propagation with the noise register during backward-forward training.

$z^{m,t} = a^{m-1,t} w^{m-1} + b^{m-1}$. More specifically, w^{m-1} and b^{m-1} denote the weight and bias for the affine transformation, and $a^{m-1,t}$ represents the activation at the previous layer. We next fetch and add the adversarial noise $r^{m,t}$ to $z^{m,t}$, and compute the activation $a^{m,t} = \text{relu}(z^{m,t})$. The adversarial noise is subsequently propagated to the next hidden layer. In contrast to the traditional adversarial training methods, we feed only clean examples to the models during training. Adversarial noise for each layer is computed and propagated to train the robust models. Algorithm 1 outlines more details of the training process with ANP in each mini-batch.

Algorithm 1 Adversarial Noise Propagation (ANP)

Input: mini-batch data (x, y)

Output: robust model parameters θ

Hyper-parameter: η, ε and k

- 1: **for** t in k steps **do**
 - 2: // Backward propagation
 - 3: Update model parameters θ using standard back-propagation.
 - 4: Compute and propagate adversarial gradient:
 $g^{m,t-1} = g^{m+1,t-1} \cdot \frac{\partial z^{m+1,t-1}}{\partial z^{m,t-1}}$
 - 5: Compute, propagate and save adversarial noise:
 $r^{m,t} = (1 - \eta)r^{m,t-1} + \frac{\varepsilon}{k} \frac{g^{m,t-1}}{\|g^{m,t-1}\|_p}$
 - 6: // Forward propagation
 - 7: Compute the affine transformation:
 $z^{m,t} = a^{m-1,t} w^{m-1} + b^{m-1}$
 - 8: Fetch and add adversarial noise:
 $z^{m,t} += r^{m,t}$
 - 9: Compute the activation:
 $a^{m,t} = \text{relu}(z^{m,t})$
 - 10: **end for**
-

V. EXPERIMENTS AND EVALUATION

In this section, we will evaluate our proposed ANP on the popular image classification task. Following the guidelines from [40], we compare ANP with several state-of-the-art adversarial defense methods against both adversarial noise and corruption as well.

A. Experimental Setup

Datasets and models. To assess the adversarial robustness, we adopt the widely used MNIST, CIFAR-10 and ImageNet datasets. MNIST is a dataset of 10 classes of handwritten digits of size 28×28 , containing 60K training examples and 10K test instances [41]. We use LeNet for MNIST. CIFAR-10 consists of 60K natural scene color images, with 10 classes, of size $32 \times 32 \times 3$ [42]. We further use VGG-16, ResNet-18, DenseNet and InceptionV2 for CIFAR-10. ImageNet contains 14M images with more than 20k classes [43]. In the interests of simplicity, we only choose 200 classes from the 1000 available in ILSVRC-2012, with 100K and 10k images used for training set and test set, respectively. The models we use for ImageNet are ResNet-18 and AlexNet.

Adversarial attacks. We apply a diverse set of adversarial attack algorithms including FGSM [9], BIM [37], Step-LL [37], MI-FGSM [44], PGD [36], and BPDA [22] in terms of ℓ_∞ -norm. We also use C&W [45] in terms of ℓ_2 -norm.

Corruption attacks. To assess the corruption robustness, we test our proposed method on CIFAR-10-C and CIFAR-10-P [28]. These two datasets are the first choice for benchmarking model static and dynamic model robustness against different common corruption and noise sequences at different levels of severity [28]. They are created from the test set of CIFAR-10 using 75 different corruption techniques (e.g., Gaussian noise, Poisson noise, pixelation, etc.).

Defense methods. We select several state-of-the-art adversarial defense methods, including NAT [23] (adversarially training a model with FGSM using different training strategies), PAT [36] (adversarial training with PGD), EAT [38] (adversarial training with FGSM generated by multiple models), LAT [39] (injecting noise into hidden layers) and Rand [15] (randomly resizing input images). Among these methods, EAT and Rand were ranked No.1 and No.2 in *NeurIPS 2017* adversarial defense competition.

To ensure that our experiments are fair, we use FoolBox [46] and select the hyper-parameters for attack and defense methods that are suggested in the relevant papers and works, e.g., [23], [25], [36], [40], etc., which are comparable to other defense strategies. Further details can be found in the Supplementary Material.

B. Evaluation Criteria

In this part, we will explicate the metrics used in our paper to evaluate model robustness against adversarial perturbations, corruption, and more generalized noise.

1) *Adversarial robustness evaluation:* We use top-1 *worst case classification accuracy* for our black-box attack defense. For a specific test set, corresponding adversarial example sets are generated using attack methods from different hold-out models. Subsequently, the worst results are selected among them as the final result. However, top-1 *classification accuracy* is utilized for the white-box attack. In this situation, adversaries know every detail of the target model and generate adversarial examples for direct attack. For these evaluation metrics, the higher the better.

2) *Corruption robustness evaluation*: We adopt mCE, Relative mCE and mFR, following [28], to comprehensively evaluate a classifier's robustness to corruption. More specifically, mCE denotes the mean corruption error of the model compared to the baseline model, while Relative mCE represents the gap between mCE and the clean data error. Moreover, mFR stands for the classification differences between two adjacent frames in the noise sequence for a specific image. For these evaluation metrics, the lower the better.

mCE. The first evaluation step involves taking a classifier f , which has not been trained on CIFAR10-C, and computing the clean dataset top-1 error rate as E_{clean}^f . The second step involves testing the classifier on each corruption type c at each level of severity s , denoted as $E_{s,c}^f$. Finally, mCE is computed by dividing the errors of a baseline model as follows:

$$CE_c^f = \frac{\sum_{s=1}^5 E_{s,c}^f}{\sum_{s=1}^5 E_{s,c}^{base}}.$$

Thus, mCE is the average of 15 different Corruption Errors (CEs).

Relative mCE. A more nuanced corruption robustness measure is that of Relative mCE. If a classifier withstands most corruption, the gap between mCE and the clean data error is minuscule. Thus, Relative mCE is calculated as follows:

$$\text{Relative } mCE_c^f = \frac{\sum_{s=1}^5 E_{s,c}^f - E_{clean}^f}{\sum_{s=1}^5 E_{s,c}^{base} - E_{clean}^{base}}.$$

mFR. Let us denote m noise sequences with $S = \{(x_1^{(i)}, x_2^{(i)}, \dots, x_n^{(i)})\}_{i=1}^m$, where each sequence is created with noise p . The ‘‘Flip Probability’’ of network f on noise sequences S is:

$$FP_p^f = \frac{1}{m(n-1)} \sum_{i=1}^m \sum_{j=2}^n 1(f(x_j^{(i)}) \neq f(x_{j-1}^{(i)})) \\ = \mathbb{P}_{x \sim S}((f(x_j) \neq f(x_{j-1}))).$$

The Flip Rate can thus be obtained by $FR_p^f = FP_p^f / FP_p^{base}$, where mFR is the average value of FR.

3) *Model Structure Robustness*: Standard methods for both attack and defense [17], [22], [23] typically evaluate model robustness using different variants of classification accuracy (e.g., worst-case, average-case, etc.). However this type of measurement only focuses on the final output predictions of the models (i.e., their final behaviors) and reveals only limited information regarding how and why robustness is achieved. To further understand how model the structure of a model can affect its robustness, we evaluate model robustness from the model structure perspective. We consequently proposed two metrics, namely *Empirical Boundary Distance* and ε -*Empirical Noise Insensitivity*, which are based on decision boundary distance and the Lipschitz constant, respectively.

Empirical Boundary Distance. The minimum distance to the decision boundary among the data points reflects the model robustness to small noise [47], [48]. Due to the computation difficulty for decision boundary distance, we propose *Empirical Boundary Distance* (denoted as W_f) in a heuristic way. Intuitively, a larger W_f means a stronger model. Given a learnt

model f and point x_i with class label y_i ($i = 1, \dots, N$), we first generate a set V of m random orthogonal directions [49]. For each direction in V , we then estimate the root mean square (RMS) distances $\phi_i(V)$ to the decision boundary of f until the model's prediction changes: i.e., $f(x_i) \neq y_i$. Among $\phi_i(V)$, d_i denotes the minimum distance moved to change the prediction for instance x_i . Our *Empirical Boundary Distance* is thus defined as follows:

$$W_f = \frac{1}{N} \sum_{i=1}^N d_i, \quad d_i = \min \phi_i(V). \quad (3)$$

ε -Empirical Noise Insensitivity. [50] first introduced the concept of learning algorithms robustness, which is based on the idea that if two samples are ‘‘similar’’, then their test errors will be very close. Inspired by this, we propose ε -*Empirical Noise Insensitivity* to measure the model robustness against generalized noise from the view of the Lipschitz constant. Evidently, a lower value indicates a stronger model. We first randomly select N clean examples, after which M examples are generated from each clean example via various methods (e.g., adversarial attack, Gaussian noise, blur, etc.). The differences between model loss functions are computed when a clean example and corresponding polluted examples are fed into the model. The different severities in the loss function are used to measure the model's insensitivity and stability to generalized small noise within constraint ε :

$$I_f(\varepsilon) = \frac{1}{N \times M} \sum_{i=1}^N \sum_{j=1}^M \frac{|l_f(x_i|y_i) - l_f(\mu_{ij}|y_i)|}{\|x_i - \mu_{ij}\|_\infty} \quad (4) \\ s.t. \quad \|x_i - \mu_{ij}\|_\infty \leq \varepsilon,$$

where x_i , μ_{ij} and y_i denote the clean example, corresponding polluted example and class label, respectively. $l_f(\cdot)$ represents the loss function of model f .

C. Is It Necessary to Inject Noise into All Layers?




Fig. 2. The results of VGG-16 trained with adversarial noise added to different top- m and bottom- m layer groups. The horizontal axis of (b) is arranged in an order opposite to that of the layer number (i.e., by adding noise to the bottom-1 layers, we mean perturbing the 13th layer). The solid lines denote models trained with different layer groups via ANP, while the dashed lines represent the baseline method NAT. The model achieves the best performance when noise is injected only into the top-4 layers.

A significant body of research exists regarding the representation power of the neural network [51], [52]; however, studies for hidden layers are rarely involved. Since not all layers are created equal [53], an intuitive question for ANP training emerges, namely: *do we need to add noise to all layers?* We




Fig. 3. The results of VGG-16 trained with adversarial noise added to distant and neighboring layers. In subfigure (a) and (b), We select the 1st and 7th layer respectively as the base layer. We perturb a pair of layers with different intervals ranging from 0 to 6.




Fig. 4. Subfigure (a) denotes the layer statistics re-assignment experiment; here dashed lines represent the baseline model with no layer weight re-assignment, i.e. ANP. Subfigure (b) illustrates single hidden layer perturbing experiment, i.e., in which we only perturb the m -th layer. The model robustness decreases as the layer goes deeper.

therefore try to investigate the contribution of hidden layers when adversarial noise is injected and propagated. In this section, all adversarial examples are generated via white-box FGSM with $\epsilon=8$ in terms of ℓ_∞ norm.

We first study the contribution made by hidden layers from the perspective of layer groups, i.e., multiple layers. First, we train a number of VGG-16 models on CIFAR-10 with adversarial noise injected into different layer groups: namely, the top- m and the bottom- m layers. From the results in Figure 2 (a), we can observe that the model robustness improves (i.e., both the classification accuracy for adversarial examples and the negative mCE for corruption increase) when injecting noise into the top- m ($m \leq 4$) hidden layers. We can further observe a similar phenomenon from Figure 2 (b): as we increase the number of bottom layers to be perturbed, the model robustness increases, especially when the shallow layers are involved. Surprisingly, however, it is not true that the model becomes increasingly robust the more layers are perturbed; this indicates that different hidden layers contribute to different extents to model robustness in deep architecture. Shallow layers are more critical to model robustness; by contrast, the importance of deep layers is somewhat more limited.

Meanwhile, we trained ANP with the progressive number $k=1$; this is roughly similar to NAT with noise confined only to the top-1 layer. Its performance curves are indicated by the orange and green solid lines without a marker (ANP-k1) in Figure 2 (a), which are almost close to that of NAT, while much lower than the model trained with $k=3$ (ANP-k3). This, in turn, demonstrates the importance of our progressive

noise injection to model robustness.

We further investigate the contributions of different layers by perturbing both neighboring layers and distant layers. For a VGG-16 model, we choose to add noise to a pair of layers by selecting a base layer and second layer, with different intervals between these two layers ranging from 0 to 12. In other words, we add noise to neighboring layers when the interval is 1, and perturb more distant layers when interval exceeds 1. As shown in Figure 3, adding distant layers at specific positions is superior to adding noise to two neighboring layers; for example, the model is most robust when we perturb the 1st and 4th layer in Figure 3 (a), or the 7th and 13th layer in Figure 3 (b). We attribute these results to the strong contributions made by the shallow layers towards model robustness, since we also obtain the strongest model when we choose the top-4 layers in Fig 2 (a).

We further investigate the contributions made by hidden layers from the perspective of single layer via layer statistics re-assignment and single hidden layer perturbation.

We first re-assign the statistics of each Batch Normalization (BN) [54] layer (i.e., running mean and running variance) of an ANP-trained VGG16 model using a corresponding vanilla model at the time of inference. Intuitively speaking, a stronger model performance gap following layer statistics re-assignment results in larger layer-wise differences, which in turn indicate a more non-trivial hidden layer on model robustness. As shown in Figure 4 (a), the influence of individual layer statistics re-assignment to model robustness reduces (i.e., the performance gap is smaller) as the layer depth increases.

We then further perturb single hidden layers individually; i.e., noise is injected only into the m -th layer. As shown in Figure 4 (b), the model robustness reduces (i.e., classification accuracy for adversarial examples and negative mCE for corruptions reduces) as the layer depth increases.

Thus, the experiments above constitute double confirmation that shallow layers are more critical to model robustness, while on the contrary, the importance of deeper layers is somewhat lesser.

3) Theoretical analysis: In this section, we try to provide theoretical analysis for the above conclusions (model performance with or without ANP).

We first consider the feed-forward neural networks with ReLU activation. Here, the function F is parameterized by a sequence of matrices $W = (W_1, W_2, \dots, W_L)$, i.e., $F = F_W$. For the sake of convenience, we assume that $W_h \in \mathbb{R}^{d_h \times d_{h-1}}$, where d_h and d_{h-1} are the dimensions of two adjacent layers, while $\rho(\cdot)$ is the ReLU function. For the vectors, $\rho(x)$ is the vector generated by applying $\rho(\cdot)$ on each coordinate of x , i.e., $\rho(x)_i = \rho(x_i)$. Thus, we have

$$F_W(x) = W_L \rho(W_{L-1} \rho(\dots \rho(W_1 x) \dots)).$$

For a K -classification problem, we have the dimension $d_L = K$, $F_W(x) : \mathbb{R}^d \rightarrow \mathbb{R}^K$, while $[F_W(x)]_k$ is the score for the k -th class.

Let $\epsilon_l \in \{\epsilon_1, \epsilon_2, \dots, \epsilon_L\}$ denote the adversarial perturbation added in the hidden layers of ANP and corresponding

perturbation vector $\epsilon_l = \epsilon_l \cdot \mathbf{1}_{d_l}$ with dimension d_l in the l -th layer. We denote the function f with layer-wise adversarial perturbations of ANP as follows:

$$F_{W,\epsilon}(x) = W_L(\rho(W_{L-1}(\rho(\cdots \rho(W_1(x+\epsilon_1)) \cdots) + \epsilon_{L-1})) + \epsilon_L)$$

Theorem 1: Consider the function $F_W(\cdot)$ and the layer-wise adversarial perturbation vector ϵ_l in the l -th layer. Given an input sample x , the difference between the $F_{W,\epsilon}(x)$ (model trained with ANP) and $F_W(x)$ (model trained without ANP, i.e., vanilla) could be derived as

$$\|F_{W,\epsilon}(x) - F_W(x)\| \leq \|\Pi_{i=1}^L W_i \epsilon_1\| + \|\Pi_{j=2}^{L-1} W_j \epsilon_{L-1}\| + \cdots + \|W_L \epsilon_L\|$$

Consequently, the difference in output scores (expressed in the accuracy on clean sample x) between classifiers $F_{W,\epsilon}(x)$ and $F_W(x)$ could be influenced more strongly by the adversarial perturbations in shallow layers. As can be seen from Figure 2 (b), as we perturb more bottom- m layers (from deep layers to shallow layers), the accuracy on clean samples is decreasing. Thus, we can draw the following conclusions: (1) shallow layers have stronger negative influences than deep layers on model accuracy for clean examples; (2) shallow layers are more critical to model robustness than deep layers; (3) the contribution of each layer with respect to layer depth is not strictly linear.

In summary, it is sufficient to *inject noise just the shallow layers only to achieve better model robustness*; this finding will be used to guide our experiments throughout the rest of the paper.

D. Adversarial Robustness Evaluation

We first evaluate model adversarial robustness by considering black-box and white-box attack defense.

1) *Black-box setting:* In black-box attack defense, adversaries have no knowledge of the target models (e.g., architectures, parameter weights, etc.). We first generate adversarial examples (10k images) using various hold-out models which are different from the target models; we then use these adversarial examples to attack the target model.

MNIST. With LeNet as the target model, the experimental results for black-box and white-box defense are listed in Table I.

CIFAR-10. As shown in Table II, on CIFAR-10, we employ VGG-16 as the target model and compare ANP with various other defense methods. The hold-out models include ResNet-50, DenseNet and Inception-v2. Among these defense methods, EAT is trained ensemble with VGG-16 and Inception-v2; meanwhile, Rand resizes input images from 32 to 36 and follows a NAT-trained VGG-16.

ImageNet. AlexNet is applied as the target model, and adversarial examples are generated from ResNet-18 and AlexNet. Meanwhile, EAT is trained with AlexNet and ResNet-18, while Rand resizes input images from 224 to 254 and follows a NAT-trained AlexNet. The results on ImageNet are presented in Table III.

From the above black-box experiments, we can make the following observations: (1) In the black-box setting, in almost

all cases, ANP achieves the best defense performance among all methods; (2) ANP makes the widely used deep models strongly robust against both single-step and iterative black-box attacks; (3) Normally, the classification performance for clean examples degrades significantly (e.g., by 5-10% in terms of accuracy) when more noise is introduced into the model; however, ANP supplies the models with good generalization ability, resulting in stable classification accuracy that is close to that of the vanilla models on CIFAR-10 as well as MNIST.

2) *White-box setting:* In the white-box scenario, we apply PGD, C&W and the attacking framework BPDA for adversarial attacks. Since we compare with Rand, which is an adversarial defense strategy that obfuscates gradients by resizing and padding the input images, we also use BPDA to circumvent the undifferentiable component and ensure a thorough analysis. Moreover, BIM is utilized in BPDA to demonstrate the white-box attack after circumventing the gradient mask with iteration number 5. The results on CIFAR-10 and ImageNet are listed in Table IV.

We can conclude that in the white-box setting, ANP exhibits a significant advantage in terms of defending against adversarial examples over other methods on these datasets, although ANP is slightly weak against PGD attack on ImageNet compared to PAT. Overall, the results indicate that training with ANP enables the model to be robust against various attack methods. However, we can also see that for large datasets like ImageNet, clean accuracy drops for all methods. The reasons for this might be two-folded. Firstly, AlexNet, a relatively small model, may not have sufficient capacity to fit adversarial noise while maintaining high accuracy on clean examples. Secondly, the trade-off between robustness and accuracy does exist [55], [56] especially for high-dimensional data distributions. We will study this in more depth in future work.

E. Corruption Robustness Evaluation

To assess corruption robustness, we conduct experiments using 10K images from CIFAR-10-C with 15 different corruption levels and 5 severity levels. To test the model's dynamic robustness, we use CIFAR-10-P, which differs from CIFAR-10-C in that noise sequences are generated for each image with more than 30 frames. As shown in Section 5.2.2, mCE indicates the average corruption error, while mFR is the average flip rate of noise sequence (for both of these, lower is better). According to the results in Figure 5 (a) and (b), ANP achieves the lowest mCE and mFR value among all methods, indicating strong corruption robustness. More precisely, as can be seen from Table V, ANP surpasses the compared strategies by large margins (i.e., almost 6 and 30 for mCE and mFR, respectively). The results demonstrate that ANP can reliably provide both static and dynamic robustness against corruption.

Although they perform well on adversarial examples, compared methods (especially PAT) show weak robustness to both static and dynamic corruptions. Another interesting phenomenon that can be observed is that all compared methods even perform worse than the vanilla model for dynamic corruption (as shown by higher mFR values). Most adversarial training methods attempt to inject noise into the inputs

TABLE I
BLACK-BOX AND WHITE-BOX ATTACK DEFENSE RESULTS ON MNIST WITH LeNET.

LeNET	CLEAN	BLACK-BOX			WHITE-BOX		
		FGSM			BIM	PGD	C&W
		$\epsilon=0.1$	$\epsilon=0.2$	$\epsilon=0.3$	$\epsilon=0.2$	$\epsilon=0.2$	
VANILLA	99.6%	72.0%	28.0%	4.0%	61.7%	22.3%	27.1%
PAT	99.0%	96.8%	90.7%	78.0%	90.6%	59.1%	63.7%
NAT	98.4%	94.0%	88.0%	72.0%	85.9%	45.1%	51.4%
ANP	99.3%	97.1%	93.2%	78.0%	91.5%	59.1%	69.1%

TABLE II
BLACK-BOX ATTACK DEFENSE RESULTS ON CIFAR-10 WITH DIFFERENT MODELS.

(a) VGG-16

VGG-16	CLEAN	FGSM		PGD	STEP-LL	MI-FGSM
		$\epsilon = 8$	$\epsilon = 16$	$\epsilon = 8, \alpha = 2$	$\epsilon = 8$	$\epsilon = 8$
VANILLA	92.1%	38.4%	19.3%	0.0%	7.5%	2.3%
PAT	83.1%	82.5%	76.3%	85.5%	80.1%	79.9%
NAT	86.1%	73.5%	70.2%	80.3%	79.1%	77.6%
LAT	84.4%	75.8%	63.7%	79.2%	78.3%	77.6%
RAND	85.2%	77.6%	70.8%	80.2%	79.3%	78.2%
EAT	87.5%	81.2%	76.2%	83.5%	82.7%	80.8%
ANP	91.7%	82.8%	76.4%	84.4%	83.3%	81.1%

(b) ResNet-18

RESNET-18	CLEAN	FGSM		PGD	STEP-LL	MI-FGSM
		$\epsilon = 8$	$\epsilon = 16$	$\epsilon = 8, \alpha = 2$	$\epsilon = 8$	$\epsilon = 8$
VANILLA	93.1%	12.8%	10.2%	6.0%	21.4%	1.0%
PAT	85.1%	82.4%	72.4%	87.4%	81.0%	79.8%
NAT	89.1%	78.1%	68.8%	83.6%	80.4%	77.7%
LAT	88.9%	45.8%	23.1%	58.3%	49.7%	33.0%
RAND	86.4%	57.6%	35.0%	70.4%	61.0%	33.0%
EAT	86.9%	80.8%	72.5%	84.7%	82.8%	80.4%
ANP	92.1%	83.6%	73.5%	86.5%	84.0%	80.4%

TABLE III
BLACK-BOX ATTACK DEFENSE RESULTS ON IMAGENET WITH ALEXNET.

ALEXNET	CLEAN	FGSM		PGD	STEP-LL	MI-FGSM
		$\epsilon = 8$	$\epsilon = 16$	$\epsilon = 8, \alpha = 2$	$\epsilon = 8$	$\epsilon = 8$
VANILLA	61.7%	12.6%	9.2%	4.3%	13.7%	3.5%
PAT	56.2%	41.5%	40.2%	42.1%	41.5%	41.2%
NAT	53.5%	39.1%	34.2%	39.6%	41.6%	37.7%
RAND	50.2%	39.2%	27.0%	39.2%	40.9%	39.7%
EAT	51.1%	39.6%	35.2%	39.9%	42.3%	37.9%
ANP	51.5%	41.7%	39.3%	42.1%	42.7%	41.2%

by searching for the worst-case perturbations, which indeed improve adversarial model robustness. However, these tactics seem to be worthless to average-case or general perturbations, and maybe also somehow counteract corruption robustness. Since robustness requires high data complexity [57], our proposed ANP introduces adversarial noise with high complexity and diversity via progressive iteration, which contributes both adversarial and corruption robustness. We can therefore conclude that ANP supplies models with stronger corruption robustness compared to other defense methods.

We also test a VGG-16 model trained with Gaussian noise $N(0, 0.1)$ added to input; however the model is weak to corruption with an mCE value 117.2.

F. Model Structure Robustness Evaluation

In this section, we evaluate the structural robustness of deep models using our proposed metrics: *Empirical Boundary Distance* and ϵ -*Empirical Noise Insensitivity*.

TABLE IV
WHITE-BOX ATTACK DEFENSE ON CIFAR-10 AND IMAGENET.
(a) CIFAR-10 with VGG-16

VGG-16	CLEAN	BPDA	PGD	C&W
		$\epsilon=8$	$\epsilon=8$	
VANILLA	92.1%	0.2%	0.0%	8.2%
PAT	83.1%	41.9%	41.4%	42.4%
NAT	86.1%	24.5%	8.1%	31.6%
RAND	85.2%	0.2%	9.1%	34.2%
EAT	87.5%	37.6%	9.9%	35.2%
FREE-4	88.9%	41.2%	38.2%	42.3%
ANP	91.7%	43.5%	28.9%	48.1%

(b) CIFAR-10 with ResNet-18

RESNET-18	CLEAN	BPDA	PGD	C&W
		$\epsilon=8$	$\epsilon=8$	
VANILLA	93.1%	0.0%	2.7%	8.4%
PAT	85.1%	45.0%	41.9%	43.9%
NAT	89.1%	33.5%	10.1%	32.6%
RAND	86.4%	1.9%	9.4%	31.2%
EAT	86.9%	40.1%	14.5%	39.1%
FREE-4	89.0%	42.9%	37.8%	44.1%
ANP	92.1%	54.0%	29.9%	47.3%

(c) ImageNet with AlexNet

ALEXNET	CLEAN	BPDA	PGD
		$\epsilon=8$	$\epsilon=8$
VANILLA	61.7%	7.9%	2.4%
PAT	56.2%	27.6%	28.7%
NAT	53.3%	26.8%	15.2%
RAND	50.2%	21.6%	16.7%
EAT	56.0%	25.0%	12.5%
ANP	53.5%	28.2%	27.4%




Fig. 5. Model corruption robustness evaluation. Our ANP outperforms other compared methods by large margins in terms of corruption robustness, with lower mCE and mFR values.

1) *Empirical Boundary Distance*: We choose 1,000 randomly selected orthogonal directions, and compute the minimum distance along each direction for a specific image to change the predicted label. As can be seen from Figure 6 (a), models trained by ANP have the largest distance. Moreover, Figure 6 (b) provides the minimum distances moved for each of 100 randomly picked images in order to change their labels. It is easy to see that the distance curve of ANP is almost the highest, with a large leading gap at the beginning. Table VI further reports W_f figures for different methods. The results

TABLE V
CORRUPTION ROBUSTNESS EVALUATION WITH MCE AND MFR.

VGG-16	ERROR	MCE	MFR
VANILLA	7.9	100.0	100.0
PAT	16.9	98.1	121.8
NAT	13.9	85.7	131.2
EAT	12.5	81.8	108.3
FREE-4	11.1	92.5	115.6
ANP	8.3	75.7	79.2

consistently prove that ANP supplies deep models with strong discriminating power by the largest margins, with the result that these models are the most robust.

2) *ϵ -Empirical Noise Insensitivity*: In this section, several different methods (including FGSM, PGD, Gaussian noise, etc.) are employed to generate polluted examples from 100 clean images. For each clean image, 10 corresponding polluted examples are generated using every method within noise constraint ϵ . As shown in Figure 7 (a) and (b), ANP obtains the smallest noise insensitivity in most cases, exhibiting strong robustness to both adversarial examples and corruption. More experimental results are shown in Figure 7 (c) to (f). Specifically, Figure 7 (c) and (d) illustrate the results on adversarial examples generated using FGSM and Step-LL; meanwhile, subfigures (e) and (f) show the results on Gaussian noise and Shot noise.

TABLE VI
EMPIRICAL BOUNDARY DISTANCE AMONG FIVE VGG-16 MODELS
TRAINED WITH DIFFERENT TRAINING METHODS.

METHOD	VANILLA	PAT	NAT	EAT	ANP
W_f	29.46	42.23	34.22	35.21	47.36




Fig. 6. Empirical Boundary Distance values are computed among five different VGG-16 models to change the predicted label: (a) the average distance moved in each orthogonal direction, and (b) the minimum distance (i.e., Empirical Boundary Distance) moved for 100 different images.

G. Combination with Other Adversarial Defense Strategies

Currently, most effective adversarial defense strategies involve adversarial training, in which adversarial examples are added during training. [9] first proposed adversarial training with an adversarial objective function based on the use of fast




Fig. 7. ϵ -Empirical Noise Insensitivity is calculated under adversarial examples and corruptions.

gradient sign method as an effective regularizer. Based on that, with the aim of universally robust networks, [36] trained models by feeding adversarially perturbed inputs with PGD into the loss term. [24] proposed to improve adversarial robustness by integrating an adversarial perturbation-based regularizer into the classification objective. Recently, [56] decomposed the prediction error for adversarial examples (robust error) into the sum of the natural (classification) error and boundary error, and consequently proposed the TRADES defense method that combines an accuracy loss and a regularization term for robustness.

Clearly, our proposed ANP framework takes advantage of hidden layers within a network and is orthogonal to the above mentioned adversarial training methods. It is therefore intuitive for us to combine our ANP with these elaborately designed objective functions, as this could further improve model robustness by fully exploiting the potential of hidden layers. Accordingly, we further train models by combining ANP with TRADES (No.1 in *NeurIPS 2018* adversarial defense competition) and PAT (the most commonly used adversarial training method).

As shown in Table VII, when ANP is combined with TRADES (ANP+TRADES) and PAT (ANP+PAT), this approach outperforms its counterparts (i.e., TRADES and PAT). More specifically, ANP+TRADES achieves the most robust performance and outperforms the compared methods. During

training, we simply combine PAT or TRADES into our ANP framework in terms of the ℓ_∞ norm. These rates enable us to draw an insightful conclusion that hidden layers could be considered and introduced along with other objective functions and regularizer terms to encourage stronger adversarial robustness in the future.

H. What Did Hidden Layers Do during Noise Propagation?

In this section, we aim to uncover the behavior and effect of hidden layers during noise propagation from the perspectives of hidden representation insensitivity and human vision alignment.

Considered from a high-level perspective, robustness to noise can be viewed as a global insensitivity property that a model satisfies [55]. A model that achieves small loss for noise in a dataset is necessarily one that learns representations that are insensitive to such noise. By injecting adversarial noise into hidden layers, ANP can be viewed as a method for embedding certain insensitivity into each hidden representation for models. We therefore try to explain the hidden layer behaviors from the perspective of hidden representation insensitivity. We measure the hidden representation insensitivity based on the degree of neuron activation value change within pairs of samples (x, x') , in which the distance between each pair is constrained with ϵ . Intuitively, when fed with benign and polluted examples, the more insensitively the neurons behave, the more robust the models will be. As shown in Figure 8, neurons in each layer behave more insensitively to ϵ -noise (PGD attack adversarial examples and corruption) after being trained with ANP.

TABLE VII
WHITE-BOX ATTACK DEFENSE ON CIFAR-10 WITH VGG-16.

VGG-16	CLEAN	FGSM	PGD	BBATTACK [58]
		$\epsilon=8$	$\epsilon=8$	$\epsilon=4$
VANILLA	92.1%	1.4%	0.0%	2.3%
TRADES	83.2%	47.9%	47.6%	64.7%
PAT	83.1%	45.3%	41.4%	61.2%
ANP+TRADES	82.1%	48.5%	49.8%	65.8%
ANP+PAT	82.0%	46.9%	44.3%	62.9%

Finally, we explain the model robustness from the perspective of alignment with human visual perception, such that a stronger model gains a more semantically meaningful gradient. As shown in Figure 9 (a) to (d), gradients for ANP-trained networks on top-4 layers (d) align well with the perceptually relevant features (such as edges) of the input images. By contrast, these gradients have no coherent patterns and appear very noisy to humans for vanilla networks (b), and are less semantically meaningful for the ANP-trained model on all layers (c).

VI. CONCLUSION

In order to improve model robustness against noise, this paper proposes a novel training strategy, named *Adversarial*




Fig. 8. Hidden representation insensitivity on different layers. Subfigure (a) to (h) represent the mean for all layers and hidden representation insensitivity in layer *conv2_relu*, *conv3_relu*, *conv4_relu*, *conv6_relu*, *conv8_relu*, *conv10_relu*, *conv11_relu*, *conv12_relu* and *conv13_relu*, respectively.

Noise Propagation (ANP), which injects diversified noise into the hidden layers in a layer-wise manner. ANP can be efficiently implemented through the standard backward-forward process, meaning that it introduces no additional computations. Our study of the behaviors of hidden layers yielded two significant conclusions: (1) empirical studies reveal that we



Fig. 9. Visualization of the loss gradient w.r.t. input pixels on CIFAR-10. Subfigure (a) denotes the input image, (b) to (d) represent gradients gained from Vanilla model, ANP on top-all layers and ANP on top-4 layers, respectively. No preprocessing was applied to the gradients (other than scaling and clipping for visualization).

only need to perturb shallow layers to train robust models; (2) theoretical proof demonstrates that the shallow layers have stronger negative influences than deep layers in terms of clean accuracy. Extensive experiments on the visual classification task demonstrate that ANP can enable strong robustness for deep networks, and can therefore aid in obtaining very promising performance against various types of noise.

Currently, most successful adversarial defense strategies considered only the input layers and trained models with elaborately designed loss functions and regularization terms. Fortunately, our proposed ANP framework takes advantage of hidden layers and is orthogonal to most of these adversarial training methods, which can be combined together to build stronger models. According to our experimental results, we can obtain a stronger model when ANP is provided with other adversarial training methods. Thus, further researchers could propose training methods that fully exploit the potential of hidden layers and can therefore promise stronger adversarial robustness.

Furthermore, our current strategy injects the same amount of adversarial noise into each layer and largely requires manual efforts (e.g., the magnitude of noise or specific layers). However, since every layer contributes to the model robustness to a different extent, it is preferable for us to devise a more adaptive algorithm that considers the heterogeneous behaviors

of different layers. It would therefore fully exploit the efficacy of every hidden layer and facilitate the building of stronger models. Meanwhile, it is intriguing to see that adding noise to shallow layers improves model robustness, while perturbing the deep layers has the opposite effect. The reasons for this remain unclear; we will investigate these further in future work.

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *International Conference on Neural Information Processing Systems*, 2012.
- [2] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [3] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, 2017.
- [4] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [5] A. Rajadesingan, R. Zafarani, and H. Liu, "Sarcasm detection on twitter: A behavioral modeling approach," in *8th ACM International Conference on Web Search and Data Mining*, 2015.
- [6] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, and T. N. Sainath, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Processing Magazine*, 2012.
- [7] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *International Conference on Machine Learning*, 2016.
- [8] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *International Conference on Learning Representations*, 2014.
- [9] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *International Conference on Learning Representations*, 2015.
- [10] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against deep learning systems using adversarial examples," *arXiv preprint*, 2016.
- [11] A. Liu, X. Liu, J. Fan, Y. Ma, A. Zhang, H. Xie, and D. Tao, "Perceptual-sensitive gan for generating adversarial patches," in *33rd AAAI Conference on Artificial Intelligence*, 2019.
- [12] K. R. Mopuri, A. Ganeshan, and R. V. Babu, "Generalizable data-free objective for crafting universal adversarial perturbations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [13] A. Liu, T. Huang, X. Liu, Y. Xu, Y. Ma, X. Chen, S. Maybank, and D. Tao, "Spatiotemporal attacks for embodied agents," in *European Conference on Computer Vision*, 2020.
- [14] A. Liu, J. Wang, X. Liu, b. Cao, C. Zhang, and H. Yu, "Bias-based universal adversarial patch attack for automatic check-out," in *European Conference on Computer Vision*, 2020.
- [15] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, "Mitigating adversarial effects through randomization," in *International Conference on Learning Representations*, 2018.
- [16] G. S. Dhillon, K. Azizzadenesheli, Z. C. Lipton, J. Bernstein, J. Kossaifi, A. Khanna, and A. Anandkumar, "Stochastic activation pruning for robust adversarial defense," in *International Conference on Learning Representations*, 2018.
- [17] J. Buckman, A. Roy, C. Raffel, and I. Goodfellow, "Thermometer encoding: One hot way to resist adversarial examples," in *International Conference on Learning Representations*, 2018.
- [18] C. Guo, M. Rana, M. Cisse, and L. Van Der Maaten, "Countering adversarial images using input transformations," in *International Conference on Learning Representations*, 2018.
- [19] Y. Song, T. Kim, S. Nowozin, S. Ermon, and N. Kushman, "Pixeldefend: Leveraging generative models to understand and defend against adversarial examples," in *International Conference on Learning Representations*, 2018.
- [20] T. Miyato, S. Maeda, M. Koyama, and S. Ishii, "Virtual adversarial training: A regularization method for supervised and semi-supervised learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [21] C. Zhang, A. Liu, X. Liu, Y. Xu, H. Yu, Y. Ma, and T. Li, "Interpreting and improving adversarial robustness of deep neural networks with neuron sensitivity," *IEEE Transactions on Image Processing*, 2020.
- [22] A. Athalye, N. Carlini, and D. A. Wagner, "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples," in *International Conference on Machine Learning*, 2018.
- [23] K. Alexey, G. Ian, and B. Samy, "Adversarial machine learning at scale," in *International Conference on Learning Representations*, 2017.
- [24] Z. Yan, Y. Guo, and C. Zhang, "Deep defense: Training dnns with improved adversarial robustness," in *Advances in Neural Information Processing Systems 31*, 2018.
- [25] M. Cisse, P. Bojanowski, E. Grave, Y. Dauphin, and N. Usunier, "Parseval networks: Improving robustness to adversarial examples," in *International Conference on Machine Learning*, 2017.
- [26] S. Dodge and L. Karam, "A study and comparison of human and deep learning recognition performance under visual distortions," in *IEEE International Conference on Computer Communication and Networks*, 2017.
- [27] S. Zheng, Y. Song, T. Leung, and I. Goodfellow, "Improving the robustness of deep neural networks via stability training," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [28] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," in *International Conference on Learning Representations*, 2019.
- [29] F. Liao, M. Liang, Y. Dong, T. Pang, X. Hu, and J. Zhu, "Defense against adversarial attacks using high-level representation guided denoiser," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [30] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, "Adversarial examples are not bugs, they are features," in *Advances in Neural Information Processing Systems*, 2019.
- [31] S. Santurkar, A. Ilyas, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, "Image synthesis with a single (robust) classifier," in *Advances in Neural Information Processing Systems*, 2019.
- [32] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," *IEEE Symposium on Security and Privacy*, 2015.
- [33] S. Gu and L. Rigazio, "Towards deep neural network architectures robust to adversarial examples," in *International Conference on Learning Representations*, 2015.
- [34] Y. Song, T. Kim, S. Nowozin, S. Ermon, and N. Kushman, "Pixeldefend: Leveraging generative models to understand and defend against adversarial examples," in *International Conference on Learning Representations*, 2018.
- [35] Q. Wang, W. Guo, K. Zhang, A. G. Ororbia II, X. Xing, C. L. Giles, and X. Liu, "Learning adversary-resistant deep neural networks," in *International Conference on Learning Representations*, 2017.
- [36] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*, 2017.
- [37] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," in *International Conference on Learning Representations, Workshop*, 2017.
- [38] F. Tramèr, A. Kurakin, N. Papernot, I. J. Goodfellow, D. Boneh, and P. D. McDaniel, "Ensemble adversarial training: Attacks and defenses," in *International Conference on Learning Representations*, 2018.
- [39] S. Sankaranarayanan, A. Jain, R. Chellappa, and S. N. Lim, "Regularizing deep networks using efficient layerwise adversarial training," in *32nd AAAI Conference on Artificial Intelligence*, 2018.
- [40] N. Carlini, A. Athalye, N. Papernot, W. Brendel, J. Rauber, D. Tsipras, I. Goodfellow, and A. Madry, "On evaluating adversarial robustness," *arXiv preprint arXiv:1902.06705*, 2019.
- [41] Y. LeCun, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [42] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.
- [43] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [44] Y. Dong, F. Liao, T. Pang, and H. Su, "Boosting adversarial attacks with momentum," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [45] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *IEEE Symposium on Security and Privacy*, 2017.
- [46] J. Rauber, W. Brendel, and M. Bethge, "Foolbox: A python toolbox to benchmark the robustness of machine learning models," 2017.
- [47] C. Cortes and V. Vapnik, "Support-vector networks," *Machine learning*, 1995.

- [48] G. Elsayed, D. Krishnan, H. Mobahi, K. Regan, and S. Bengio, "Large margin deep networks for classification," in *Advances in Neural Information Processing Systems*, 2018.
- [49] W. He, B. Li, and D. Song, "Decision boundary analysis of adversarial examples," in *International Conference on Learning Representations*, 2018.
- [50] H. Xu and S. Mannor, "Robustness and generalization," *Machine learning*, 2012.
- [51] K. Hornik, "Approximation capabilities of multilayer feedforward networks," *Neural networks*, 1991.
- [52] O. Delalleau and Y. Bengio, "Shallow vs. deep sum-product networks," in *Advances in Neural Information Processing Systems*, 2011.
- [53] C. Zhang, S. Bengio, and Y. Singer, "Are all layers created equal?" *arXiv preprint arXiv:1902.01996*, 2019.
- [54] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.
- [55] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, "Robustness may be at odds with accuracy," *stat*, 2018.
- [56] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, "Theoretically principled trade-off between robustness and accuracy," *arXiv preprint arXiv:1901.08573*, 2019.
- [57] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry, "Adversarially robust generalization requires more data," in *Advances in Neural Information Processing Systems*, 2018.
- [58] W. Brendel, J. Rauber, M. Kümmerer, I. Ustyuzhaninov, and M. Bethge, "Accurate, reliable and fast robustness evaluation," in *Advances in Neural Information Processing Systems*, 2019.