

Enhancing Resilience of Deep Learning Networks By Means of Transferable Adversaries

1st Moritz Seiler
Information Systems and Statistics
University of Münster
Münster, Germany
moritz.seiler@uni-muenster.de

2nd Heike Trautmann
Information Systems and Statistics
University of Münster
Münster, Germany
trautmann@wi.uni-muenster.de

3rd Pascal Kerschke
Information Systems and Statistics
University of Münster
Münster, Germany
kerschke@uni-muenster.de

Abstract—Artificial neural networks in general and deep learning networks in particular established themselves as popular and powerful machine learning algorithms. While the often tremendous sizes of these networks are beneficial when solving complex tasks, the tremendous number of parameters also causes such networks to be vulnerable to malicious behavior such as adversarial perturbations. These perturbations can change a model’s classification decision. Moreover, while single-step adversaries can easily be transferred from network to network, the transfer of more powerful multi-step adversaries has – usually – been rather difficult.

In this work, we introduce a method for generating strong adversaries that can easily (and frequently) be transferred between different models. This method is then used to generate a large set of adversaries, based on which the effects of selected defense methods are experimentally assessed. At last, we introduce a novel, simple, yet effective approach to enhance the resilience of neural networks against adversaries and benchmark it against established defense methods. In contrast to the already existing methods, our proposed defense approach is much more efficient as it only requires a single additional forward-pass to achieve comparable performance results.

Index Terms—Deep Learning, Adversarial Training, Multi-step Adversaries

I. INTRODUCTION

First notions of *artificial neural networks (ANNs)* – in the context of supervised learning – date back to the early 1940s [1]. Although considerable research has been conducted in the succeeding decades, people had to wait for the emergence of powerful computers, and the accompanying (significant) drop in costs for storing and processing data, until ANNs turned into highly competitive machine learning algorithms [2], [3]. Nowadays, ANNs represent the state of the art in a variety of supervised learning tasks, such as image, sound or object recognition [4]–[6].

However, while the strength of ANNs undoubtedly results from their huge amount of parameters, their enormous complexity also makes them rather incomprehensible. Although numerous works have introduced methods for interpreting the decisions made by ANNs [7]–[9], these methods lack explanations for malicious and intentional behaviour against ANNs.

One possibility for fooling ANNs is the addition of carefully crafted noise, so-called *adversarial perturbations* (cf. Figure 1), to the model’s input [10]–[17]. It should be noted

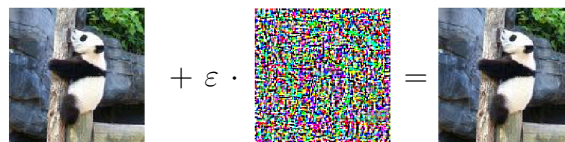


Fig. 1. Example for an adversary (from left to right): the original image (classified as *great panda* with a confidence of 99.99%), adversarial noise with $\epsilon = 3$ and the resulting adversarial image (classified as *saxophone* with a confidence of 83.8%).

that these perturbations do not need to be crafted individually per network; instead, ANNs that are trained for a similar task often can be fooled by the same adversaries [10], [11], [18].

While a plethora of works have shown the advantages of (deep learning) neural networks, their black-box characteristic makes them very vulnerable. Hence, in order to enable a more frequent and, even more importantly, more secure integration of neural networks in our daily lives, we need to ensure their robustness against adversarial attacks.

The issue of a network’s vulnerability and how to defend it against adversarial attacks will be investigated in this work. As such, our contributions can be summarized as follows:

- 1) We present a **novel method to craft transferable adversaries**, which in turn helps to improve the understanding of a network’s vulnerability against adversarial attacks.
- 2) We introduce a novel approach to regularize decision boundaries and, thus, **enhance the resilience of neural networks against adversarial attacks**. Compared to previous attempts, our approach does not require additional backward-passes, which can decrease training speed significantly (cf. Table I).

The remainder of this manuscript is structured as follows. In Section II we set our work into context. We then introduce a method for crafting transferable adversaries in Section III, and list possible defense strategies against adversarial attacks in Section II-B. Thereafter, we provide our experimental setup in Section V and discuss our findings in Section VI. At last, Section VII concludes our work.

II. RELATED WORK

Several methods to craft adversarial examples have been published in recent years [11], [16], [19], [20]. In principle,

one can categorize these methods into *single-* and *multi-step* attacks: the former only require a single gradient while the latter calculate several gradients. Kurakin et al. [21] found in their work that multi-step adversaries are less transferable than single-step ones and they concluded that models are indirectly robust against multi-step adversaries. Yet, in line with the findings of [22], we observed that the transfer of adversaries between models is rather easy when using a so-called *ensemble attack*, i.e., a multi-step attack, which is executed on an ensemble of models simultaneously [22].

As generating adversarial perturbations is rather simple, we will at first address the characteristics of these perturbations along with their accompanying risks of black-box attacks (see Section II-A). Note that in contrast to white-box methods, black-box approaches have no access to the model’s gradient. Thereafter, in Section II-B, we briefly outline selected strategies to enhance resilience.

A. Characteristics of Adversaries

Szegedy et al. [10] crafted artificial noise, which causes ANNs to misclassify images whose original versions they classified correctly. As a result, it could be demonstrated that already slight changes, which are hardly recognizable to the human eye, are sufficient to fool the network. As further demonstrated by [10], [11], and also in this work, adversarial perturbations are not model-specific, but instead generalize well over networks with different hyper-parameters. In addition, even networks trained on disjoint data sets, yet fulfilling a similar classification task, are likely vulnerable to the same adversarial image [11], [18]. Thus, two models with different architectures, hyper-parameters, and trained on different data sets, often misclassify the same images. This property can be used to execute black-box attacks as demonstrated by [18].

Several works have been conducted to examine the effects of adversaries. The neighborhoods of adversarial examples and their original counterparts were investigated by [13]. The authors found that while already small and random noise applied to adversaries often caused it to be shifted back to their original class, the classification of the original images did not change (even when rather large random noise was applied). Therefore, the authors concluded that adversarial examples likely inhabit small sub-spaces within the decision space.

Next, [23] investigated these subspaces by measuring the number of orthogonal attack directions and found that the number of these directions is rather large. They concluded that the more orthogonal attack directions exist at a given point, the larger these adversarial subspaces are and the more likely adversaries transfer between models (since larger subspaces are more likely to intersect between different models). This may be the explanation on *why* adversarial perturbations are often universal across different model designs and often universal across datasets.

B. Enhancing Resilience through Adversarial Training

Defense strategies against adversarial attacks can mainly be distinguished into two categories: approaches that either

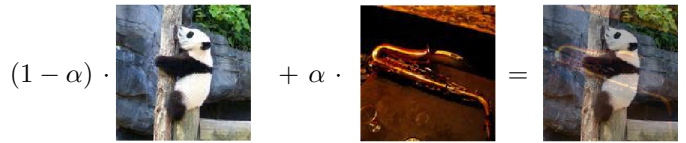


Fig. 2. Idea of our proposed method: linear combination of the original image (*great panda*) with a randomly selected image (*saxophone*) using weight $\alpha \in (0, 0.5)$ (here: $\alpha = 0.3$). The resulting training image (right) inherits the class of the original image. This allows us to generate, on the one hand, more training examples and, on the other one, training examples which lie in close proximity to the decision boundaries.

change (i) a model’s topology, or (ii) the way a model is trained. Gu and Rigazio [12] stated that a model’s robustness “is more related to intrinsic deficiencies in the training procedure and objective function than to [its] topology”.

Under the concept of *Adversarial Training*, several works have introduced regularizers which decrease a model’s vulnerability to adversarial perturbations. The first one was introduced by Goodfellow et al. [11]. Here, the objective function of the *Fast Gradient Sign Method* (FGSM) [11] is added to the classification loss. For every training image, a single-step adversary is calculated and the network’s sensitivity to this adversary is minimized. Thus, *Adversarial Training* acts as a regularizer to improve resilience [11], [21].

The training method of [11] only uses a single back-propagation step to calculate adversarial perturbations. In contrast to the aforementioned approach, other works introduced a min-max optimization method [24], [25]. It prepares the network for the worst-case scenario under a limited perturbation and minimizes its sensitivity to it. The algorithm works in two steps. First, it maximizes the loss of the input image by adding a small adversarial perturbation. However, the network trained with a min-max regularizer performed worse on the classification tasks [24] – even when trained with rather small perturbations. Due to the additional maximization problem, the algorithm is even more complex to execute, which further increases training time.

Similar approaches were introduced by [20], [26]. The approach of [26], denoted *Virtual Adversarial Training* (VAT), works in a similar way as the min-max optimization, yet it maximizes the Kullback-Leibler divergence [27] between the probability vectors of the training image and an adversary. The KL-divergence between the original and an adversary can be interpreted as the local smoothness of the decision space. Thus, the network not only learns to classify the training images correctly, but also smoothens the decision boundaries around each of the training points. Miyato et al. [26] used *Projected Gradient Descent* (PGD) to solve the inner maximization problem. The authors could demonstrate a guaranteed robustness even against multi-step adversaries under small l_∞ -norm – with $l_\infty = \max\{|x - x^{adv}|\}$ – where x is the unchanged input image while x^{adv} is the image with additional adversarial perturbation. Multi-step black- and white-box attacks have been conducted under a $l_\infty \leq 8$ -constraint. Even under multi-step white-box attacks, the accuracy of their model is



Fig. 3. Exemplary images from the modified *Large Scale Visual Recognition Challenge 2017* dataset: Seal, Bear, Dragonfly and Rabbit (left to right).

45%, whereas the accuracy of a non-robustified, and thus defenseless, model would be close to zero. However, as there are additional forward-backward passes required, training time takes accordingly longer.

III. CRAFTING STRONG AND TRANSFERABLE ADVERSARIES

Goodfellow et al. [11] introduced the *Fast Gradient Sign Method* (FGSM):

$$x^{adv} = x + \varepsilon \cdot \text{sign}(\nabla_x L(\Theta, x, y)) \quad (1)$$

The loss function* L of the model Θ is differentiated with respect to the input image x and its true class y . The sign of the resulting gradient $\nabla_x L$ is used to calculate adversarial noise which, when applied to an image, increases the model’s loss and, thereby, could shift its classification decision. The parameter ε is used to control the amount of perturbations. Note that a larger ε does not necessarily lead to a higher chance of fooling a model as demonstrated by Figure 6. Instead, tuning ε is highly important. The resulting adversary is a single-step one.

One can execute the FGSM multiple times with a small ε -value to converge towards the adversarial sub-space iteratively. The resulting adversaries are then referred to as multi-step ones. In order to find multi-step adversaries under limited amount of perturbations and close to the original image, several methods such as [16], [19]–[21] have been introduced. When the perturbation exceeds a certain limit – usually the l_∞ - or l_2 -norm of the adversarial perturbation is used – the perturbation is projected back into the search space, e.g., by using pixel-wise clipping.

As pointed out by [13], [23], adversaries lie in sub-spaces which are in close proximity to true training examples and which often intersect between models. One can argue that when crafting adversaries with multi-step methods, these adversaries often do not lie in intersecting sub-spaces because the resulting adversary may be overfitted to the particular model as its loss is maximized. To overcome this issue, [22] introduced the (i) *Momentum Iterative Gradient Based Method* (MI-FGSM) and (ii) the *ensemble in logits*. The former method uses a momentum to overcome local maxima and the latter one combined the logits of an ensemble of models to find a common and intrinsic direction. Attacking an ensemble of models increases the likelihood of transferable adversaries as these adversaries are not optimized for a single model but for

*In this work, we refer to the cross-entropy as loss function. Yet, any other function can most likely be used, too.

several models simultaneously. Hence, the adversaries lie more likely in intersecting adversarial sub-spaces.

In contrast to [22], we used the combined gradient of an ensemble of models to craft adversaries which lie in intersecting sub-spaces. We call our method *gradient ensemble attack*. We found that these adversaries are likely to transfer between models of the same architecture – and frequently transfer to other architectures as well (see Section VI). To ensure that every model’s gradient has the same impact on the resulting adversarial perturbation, we normalize each gradient to length one:

$$\hat{\nabla}_x L(\Theta, x, y) = \frac{\nabla_x L(\Theta, x, y)}{\|\nabla_x L(\Theta, x, y)\|_2}. \quad (2)$$

Afterwards, the different gradients are summed up to find a common direction of intersecting adversarial sub-spaces. Similar to [21], we approximated the true direction by using the sign-method of the summed and normalized gradients to the image in an iterative process:

$$x_i^{adv} = x_{i-1} + \lambda \cdot \text{sign} \left(\sum_{n=0}^N \hat{\nabla}_x L(\Theta_n, x_{i-1}, y) \right) \quad (3)$$

s.t. $\|x_i^{adv} - x\|_\infty \leq \varepsilon$ and $x_0 = x$.

To ensure that the magnitude of the perturbations stays within a given limit, we used pixel-wise clipping. Further, we used a learning rate of $\lambda = 1$ (to slowly convert towards the adversarial spot) and set the number of steps to $I = \min(\varepsilon + 4, 1.25 \cdot \varepsilon)$ as proposed by [21] to craft adversarial images. In addition, we chose $\varepsilon \in \{4, 8, 16\}$ to limit the amount of perturbations in order to test and compare different magnitudes of adversarial perturbations.

IV. PROPOSED DEFENSE STRATEGIES

As demonstrated by [28] and [29], adversarial training may lead to *gradient masking*. This term describes a phenomenon in which the decision boundaries are roughly the same, yet the corresponding gradient is damaged or obfuscated in a way that white-box attacks fail. However, the model does not become more resilient against adversarial examples in black-box attack scenarios or against transferable adversaries as these attacks are based on surrogate models.

In order to avoid the risk of gradient masking, we propose a method that does not require gradient calculations and still flattens the decision space. In addition, we avoid the expensive optimization of an inner maximization problem as done in VAT or PGD. We found that while random noise is mostly ignored by the model, superimposing two pictures does distract a model. Therefore, as illustrated in Figure 2, we designed training images \tilde{x}_i by placing a randomly selected image x_r on top of the original training image x_i :

$$\tilde{x}_i = (1 - \alpha) \cdot x_i + \alpha \cdot x_r. \quad (4)$$

The parameter $\alpha \in (0, 0.5)$ controls the impact of image x_r on the generated image \tilde{x}_i . As the majority of the image originates from the original image x_i , the generated image \tilde{x}_i will inherit

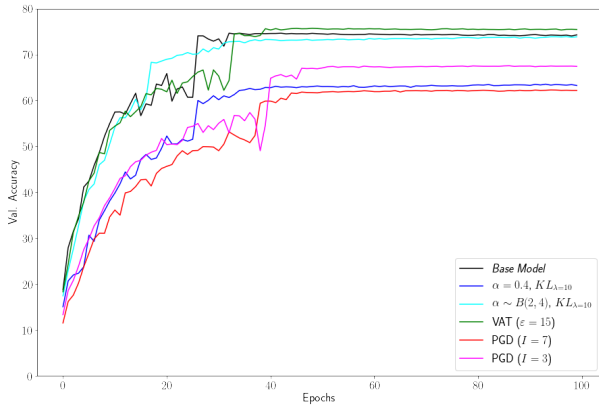


Fig. 4. Visualization of the validation accuracy during training. The *Projected Gradient Descent* [20] substantially reduces the overall accuracy compared to the *Base Model*, our proposed method, or the *Virtual Adversarial Training* [26]. Furthermore, the model $\alpha \sim B(2, 4)$, $KL_{\lambda=10}$ achieves an accuracy of more than 70% with fewer methods than the *Base Model*.

the class label y_i of the original image. Our proposed approach thus allows to (i) generate *more training examples*, and (ii) create images that are closer to the decision boundaries of at least two classes – and thus *harder to distinguish* – as \tilde{x}_i contains properties of two image classes. Thereby, the space between the different classes is flattened and the boundaries are regularized. Thereafter, the networks will be trained by minimizing the loss $L(\Theta, \tilde{x}_i, y_i)$. As the results depend on the choice of α , we considered three different approaches:

- 1) Using a fixed α -value.
- 2) Following the idea of [26], i.e., first predicting the classes y_i and \tilde{y}_i based on x_i and \tilde{x}_i (and using a fixed α as in 1.), and then minimizing the loss of the unmodified training examples *and* the KL-Divergence between the predicted classes of the unmodified and modified training examples:

$$\underset{\Theta}{\text{minimize}} \quad L(\Theta, x_i, y_i) + \lambda \cdot KL(\hat{y}_i \parallel \tilde{y}_i).$$

- 3) Similarly to 2., but this time drawing α from a beta-distribution instead of a fixed α -value; $\alpha \sim B(p, q)$ with $p = 2$ and $q \in \{4, 6, 10\}$.

Note that while the first method does not require any additional passes, the latter two methods require a second forward-pass to calculate the KL-Divergence – but no additional backward-pass. As stated by [26], the KL-divergence can be interpreted as the local smoothness of the decision space. If the divergence is high, the predictions of both input images are dissimilar, implying that the decision space between the two activations is not flattened. By minimizing the divergence, the models learn to find similar activations for both images and, thereby, flatten the decision space between the two activations.

As a side effect, our method allows to generate a multitude of training images (using Equation (4)), which in turn allows to train a larger area.

V. METHODOLOGY AND EXPERIMENTS

In the following, we briefly outline the models’ topologies and used training data.

TABLE I

MEASURED TRAINING TIME FOR DIFFERENT MODELS (ALL MODELS WERE TRAINED ON A SINGLE NVIDIA QUADRO RTX 6000). TRAINING A SINGLE EPOCH WITH OUR PROPOSED METHOD TAKES ONLY SLIGHTLY LONGER THAN THE *Base Model* WITHOUT ANY DEFENSE. IN ADDITION, $\alpha \sim B(2, 4)$, $KL_{\lambda=10}$ REACHES 60% AND 70% VALIDATION ACCURACY ABOUT 30 MINUTES EARLIER THAN THE *Base Model* AS IT CONVERGES FASTER (CF. FIGURE 4). TRAINING MODELS USING THE *Projected Gradient Descent* [20] OR *Virtual Adversarial Training* [26] TAKES NOTICEABLY LONGER PER EPOCH, AND THEY TAKE EVEN LONGER TO CONVERGE.

Model	Median Time p. Epoch	Time until 50% Acc.	Time until 60% Acc.	Time until 70% Acc.
<i>Base</i>	6:43m	00:59h	02:15h	03:17h
$\alpha = 0.4$	6:55m	02:18h	03:14h	-/-
$\alpha \sim B(2, 4)$	6:51m	01:01h	01:42h	02:44h
VAT ($\epsilon = 15$)	14:45m	02:12h	03:55h	08:05h
PGD ($I = 7$)	22:40m	11:46h	15:56h	-/-
PGD ($I = 3$)	11:37m	03:53h	07:45h	-/-

A. Dataset

Within our experiments, we used the data from the second challenge of the *Large Scale Visual Recognition Challenge 2017*[†] [30] for training and validating our models. Yet, as down-scaling the whole images to the required size was not reasonable, we first used bounding boxes to cut out the objects within each of the images. Then, as the bounding boxes were of different sizes, the cropped images were down-scaled to the desired resolution of 128×128 pixels (cf. Figure 3).

As a result, our dataset contains roughly 400,000 images. Unfortunately, the images are not uniformly distributed across all 200 classes. As the mismatch between the classes is too large to use oversampling, we selected 100 classes with an identical number of images per class. For training and validation purposes, we performed holdout with a 80-20-split. Thereby, we created a dataset which is larger and has a higher resolution than CIFAR10/100 [31] while (at the same time) being significantly smaller and less complex than ImageNet [30]. In addition, by using balanced classes, we eliminate any side-effects which may occur from unbalanced datasets and may influence our results.

B. Models

To investigate the effects of adversarial perturbations, we trained multiple ANNs. First, an “unprotected” model – denoted *base model* in the following – is trained. For this purpose, we took a VGGNet [5], as it provides a straightforward and uncluttered design of convolutions, pooling and dense layers. In addition, we modified the design by using a *Global-Average-Pooling* layer [32] and extended each CNN layer with a *Batch Normalization* layer [33]. Afterwards, we compared our proposed methods to different ResNets [4] to verify our findings.

In order to compare different methods to defend against adversarial attacks, we trained several models with and without different defense methods (see Section VI for more details). All models were trained on a single Nvidia RTX 6000 or

[†]<http://image-net.org/challenges/LSVRC/2017/>

a single Nvidia V100 GPU. We found, that by using VAT [26] or PGD [20] the training speed is reduced, significantly. Each training epoch not only takes longer, but the models also converge more slowly towards the optimum. In contrast, our proposed method is more time efficient as it only requires an additional forward pass and it even converges faster (cf. Figure 4 and Table I).

C. Assessing a Network’s Resilience against Adversarial Perturbations

To assess the robustness of a network against adversaries, we trained six models per considered network architecture. We then used an ensemble of one, three or five of the six ANNs[‡] (see Section VI for details) to extract a set of adversarial images using Equation (3). More precisely, an image is considered being adversarial, if it is misclassified by *all* of the ensemble’s networks – note that for simplicity the networks do not have to agree on the same wrong class. The set of the extracted adversarial images is then classified by the sixth model and its accuracy is taken as quality indicator of the respective network architecture.

VI. RESULTS AND DISCUSSION

At first, adversaries have been generated as described in Section V-C. For the gradient alignment, we considered an ensemble of one, three and five models, respectively. In a first analysis, we investigated the impact of the perturbation parameter for the values $\varepsilon \in \{4, 8, 16\}$. Interestingly, we observed that the success rate for crafting adversaries is not sensitive to the tested values for ε . Therefore, we are only referring to the adversaries with $\varepsilon = 16$ as they are the most powerful ones. If a single unprotected model – which is fully identical to the *Base Model* in terms of topology and training execution – is used to calculate multi-step adversaries, the *Base Model’s* classification accuracy is still 23.8% as shown in Table II. However, aligning the gradients of an ensemble of three or five models, the *Base Model’s* accuracy on these adversaries decreases to 4% and 0.9%, respectively.

Next, we trained multiple models with *Virtual Adversarial Training (VAT)*, *Projected Gradient Descent Training (PGD)* and the three different defense methods proposed in this work (see Section II-B). For VAT we used $I = 3$, $\lambda = 1$ and $\varepsilon \in \{5, 10, 15, 20, 25\}$. Miyato et al. [26] recommended using $I = 1$ and $\lambda = 1$ as they found it to be sufficient. We increased the *power of iterations* to $I = 3$ to ensure a better conversion (cf. Miyato et al. [26] for more details). As tuning ε is most important, we tried several different values and compared them to each other. Next, we adapted the default parameter for PGD as proposed by [20]: $\varepsilon = 8$, $\lambda = 2$ and $I \in \{3, 7\}$ as the number of iterations. $\varepsilon = 8$, $\lambda = 2$ and $I = \{7\}$ are the default settings used for on CIFAR10 by [20]. In addition, we used $I = 3$ to speed up training.

Table II shows the results of different methods based on their accuracy on our crafted adversarial images. As indicated

[‡]Note that our method applied to an ensemble of one model is identical to the i-FGSM.

by the base model’s accuracy values on the adversarial data, the more models are used for our proposed *gradient ensemble attack*, the higher is the success rate of transferring the adversaries to other models. This demonstrates that our adversaries, crafted from an ensemble of models, are likely transferable to other networks. Moreover, the VAT models seem to perform best on adversarial images.

To test the generalizability of our approach, we additionally assessed our adversaries on ResNet32 and ResNet50 [4]. As shown in Table III, when applying a *gradient ensemble attack* on VGGNet13 and the ResNet models together, the resulting adversaries likely transfer between both topologies. The accuracy of unprotected (base) models on our combined adversarial dataset is 0.01% for the VGGNet13 network and about 26 to 27% for both ResNet models – although all three models have an accuracy of over 90% on the original images. Even adversaries that were originally crafted for a different topology reduce a model’s accuracy noticeably.

We further tested our method on different ResNets. As shown in the bottom half of Table III, we found that not only the originally considered VGGNet13 models, but also ResNet32 and ResNet50 became more resilient against the transferable adversaries.

However, comparing the performance of adversarial defense methods merely based on the model’s accuracy or on the success rate for crafting adversaries is problematic. Adversarial sub-spaces may occur a little aside of the original ones or gradient masking could prevent gradient-based attack methods from being successful. Therefore, we do not only refer to a model’s accuracy on strong and transferable adversaries, but also investigate the surrounding decision space, as well as its gradient.

Figure 5 illustrates the loss of six different models based on the decision space of a single input image in two adversarial directions – one taken from an unprotected model and the other one in direction of the related model. As depicted by Figure 5 (a), i.e., the image of the *Base Model*, one can see an adversarial sub-space in close proximity of the input image as indicated by the high loss value (shown in red). Thus, it only takes a small change in ε ($\varepsilon = 6$ is optimal in this particular case) to push the input image X deep into this adversarial sub-space. A similar adversarial sub-space occurs in case of the VAT model as shown in Figure 5 (b). Here, the distance is greater, but the sub-space still exists. So, it takes a larger ε to fool the VAT model.

In case of our proposed methods, using a fixed α -value is not sufficient either as the blind-spot still exists, as illustrated in Figures 5 (d) and (e). Although our proposed methods with a fixed α -value weaken the adversarial spot, they do not eliminate it. Probably, the α -value is chosen too high in this case as there is another adversarial spot behind the first one. We found that sampling $\alpha \sim B(p, q)$ is essential to flatten the decision space. As demonstrated by Figure 5 (f) the decision space of our method with varying α -values is significantly more flat than all the others. In fact, it is even more flat than the one of the PGD model – especially for large ε -values as

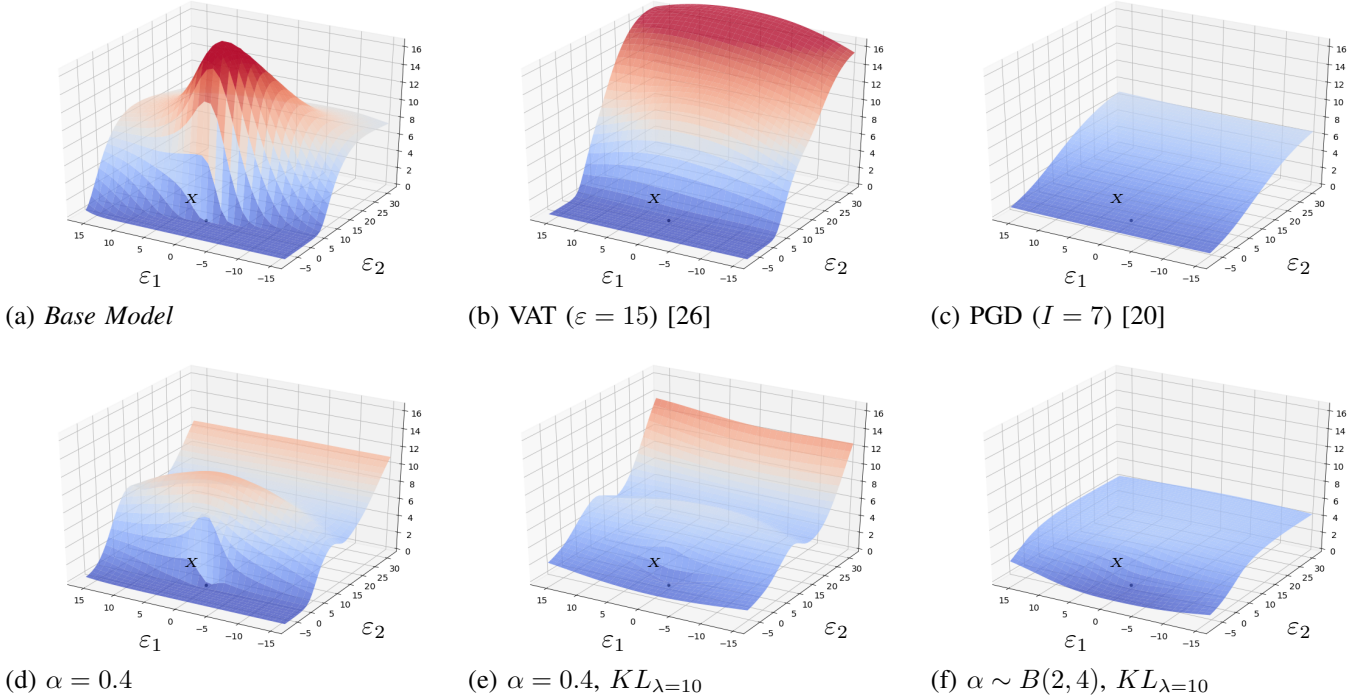


Fig. 5. Visualization of the decision space of six different VGGNet13 models in two adversarial directions of the same input image X . The loss is plotted using $x^* = x + \varepsilon_1 \cdot \text{sign}(\nabla_x L_1) + \varepsilon_2 \cdot \text{sign}(\nabla_x L_2)$ (cf. [28]). The red peak is an adversarial spot, which corresponds to a very high loss. The magnitude of the related model’s gradients $\text{sign}(\nabla_x L_2)$ is controlled by $\varepsilon_2 \in [-8, 32]$, whereas $\varepsilon_1 \in [-16, 16]$ controls the magnitude of an unprotected model’s gradient $\text{sign}(\nabla_x L_1)$. The six images visualize how the loss values change when the input x is moved in one of these two directions. The images in (d) to (f) correspond to our proposed models.

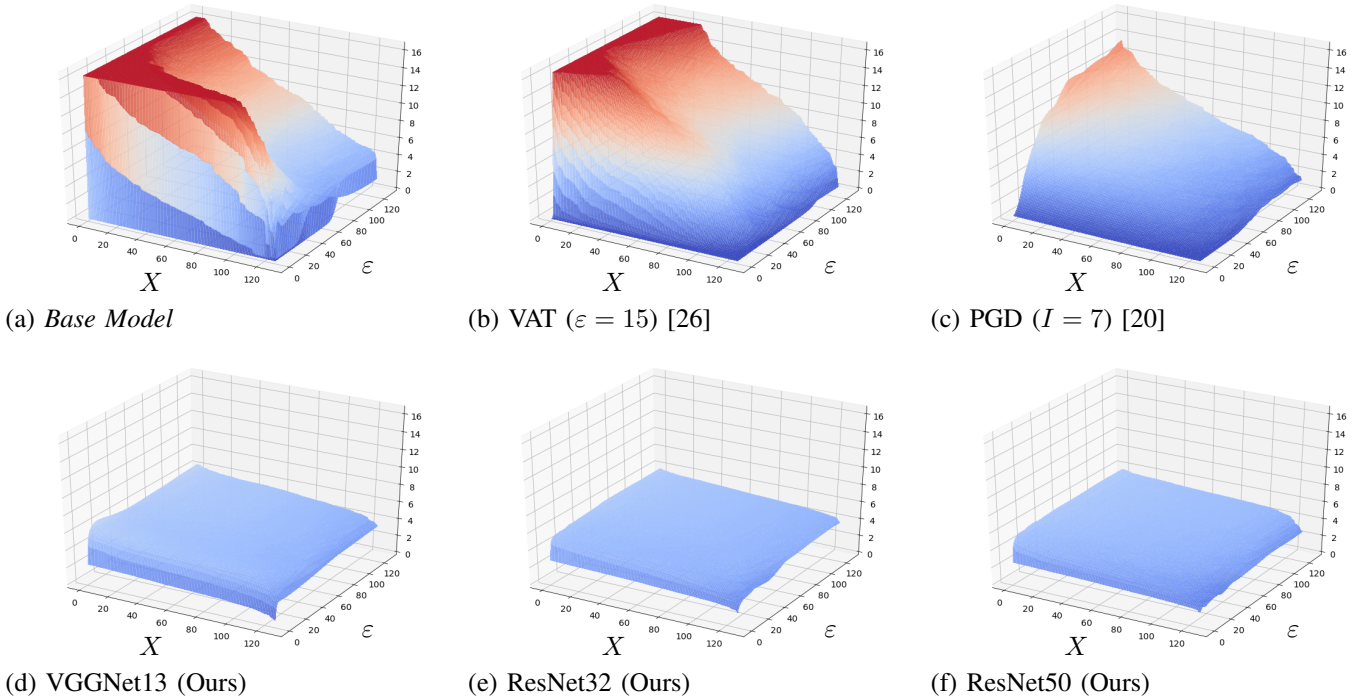


Fig. 6. Impact of different images X with additional perturbation on the loss value. The magnitude of the loss is controlled by varying ε . In comparison to the models from the literature (top row), our approach flattens the decision space significantly as demonstrated by the low loss values – even for very large values of ε . Note that our proposed models (bottom row) were trained using $\alpha \sim B(2, 4)$ and $KL_{\lambda=10}$.

TABLE II

ACCURACY OF DIFFERENT NETWORKS ON THE SETS OF ORIGINAL AND ADVERSARIAL IMAGES. ALL ADVERSARIAL DATASETS ARE CRAFTED USING OUR PROPOSED *gradient ensemble attack* APPROACH WITH ONE, THREE OR FIVE (VULNERABLE) VGGNET13 MODELS. OUR METHOD IS EQUIVALENT TO [20], [21] IF ONLY ONE MODEL IS USED. THE SHOWN RESULTS ARE FOR $\varepsilon = 16$. WHEN TRAINING A MODEL WITH PGD, THE PERFORMANCE ON THE ORIGINAL IMAGES IS SIGNIFICANTLY LOWER IN COMPARISON TO THE *Base Model*. IN CONTRAST, VAT AND OUR PROPOSED METHOD SLIGHTLY INCREASE THE VALIDATION ACCURACY. ACROSS ALL ADVERSARIAL PERTURBATIONS, VAT MODELS ACHIEVED THE HIGHEST ACCURACY. NOTEWORTHY VALUES ARE HIGHLIGHTED.

Parametrization	Accuracy on Original Data		Accuracy on Adversarial Dataset based on					
	Train	Val.	One Model		Three Models		Five Models	
			True	Adv.	True	Adv.	True	Adv.
<i>Base Model</i>								
	91.2%	73.4%	89.9%	23.8%	95.3%	4.1%	96.7%	0.9%
<i>Projected Gradient Descent (PGD) with $\varepsilon = 8$ and $\lambda = 2$</i>								
$I = 3$	86.7%	66.2%	79.6%	79.0%	84.6%	83.6%	86.6%	86.5%
$I = 7$	77.8%	62.4%	75.6%	75.1%	80.7%	80.0%	82.8%	81.8%
<i>Virtual Adversarial Training (VAT) with $I = 3$ and $\lambda = 1$</i>								
$\varepsilon = 5$	97.7%	76.2%	90.6%	85.1%	95.2%	87.2%	96.6%	86.9%
$\varepsilon = 15$	99.0%	74.6%	88.5%	86.0%	93.0%	86.0%	94.5%	90.3%
$\varepsilon = 25$	96.9%	76.5%	91.1%	83.1%	95.5%	82.8%	96.8%	80.2%
<i>Our Proposed Method with $KL_{\lambda=10}$</i>								
$\alpha \sim B(2, 4)$	90.6%	74.3%	88.8%	71.9%	93.5%	62.1%	95.1%	53.7%
$\alpha \sim B(2, 6)$	84.7%	72.4%	87.1%	74.2%	91.8%	70.0%	93.7%	66.0%
$\alpha \sim B(2, 10)$	86.9%	71.4%	85.8%	76.9%	90.4%	75.0%	92.4%	72.4%

TABLE III

ACCURACY OF DIFFERENT NETWORKS ON THE SETS OF ORIGINAL AND ADVERSARIAL IMAGES. ALL ADVERSARIAL DATASETS ARE CRAFTED USING OUR *gradient ensemble attack* APPROACH WITH FIVE UNPROTECTED VGGNET13, RESNET32 AND RESNET50 MODELS. FOR THE LAST DATASET (LAST COLUMN), WE COMBINED THREE VGG13, RESNET32 AND RESNET50 MODELS. THE TABLE SHOWS THE ACCURACY OF THREE UNPROTECTED MODELS ON DIFFERENT ADVERSARIAL DATASETS. AS ONE CAN SEE, MULTI-STEP ADVERSARIES DO TRANSFER BETWEEN TOPOLOGIES. FURTHER, OUR PROPOSED METHOD PROVIDES RESILIENCE AGAINST THESE TRANSFERABLE ADVERSARIES.

	Accuracy on Original Data		Accuracy on Adversarial Dataset							
	Train	Val.	VGGNet13		ResNet32		ResNet50		Combined	
			True	Adv.	True	Adv.	True	Adv.	True	Adv.
<i>Base Models</i>										
VGG13	91.2%	73.4%	96.7%	0.9%	96.2%	28.3%	96.9%	50.2%	98.4%	0.01%
ResNet32	89.4%	66.7%	87.9%	67.6%	95.5%	16.0%	95.3%	44.0%	96.1%	26.1%
ResNet50	87.5%	62.8%	84.1%	64.9%	92.2%	24.6%	94.7%	30.9%	94.5%	26.7%
<i>Our Proposed Method with $\alpha \sim B(2, 4)$ and $KL_{\lambda=10}$</i>										
VGG13	90.6%	74.3%	95.1%	53.7%	96.6%	69.6%	97.1%	81.5%	98.1%	32.4%
ResNet32	83.1%	59.8%	80.3%	67.8%	88.5%	38.5%	89.5%	56.9%	90.2%	45.1%
ResNet50	80.4%	62.4%	83.8%	70.3%	91.6%	44.1%	92.9%	56.4%	93.6%	46.9%

illustrated in Figures 5 (c) and (f), respectively.

To verify these findings, we generated 128 adversarial images X (using the FGSM) for each defense method and compared the loss values based on varying $\varepsilon \in [0, 128]$ (c.f. Figure 6). As for the VAT model, adversarial attacks can be highly successful. In fact, $\varepsilon = 26$ provides the greatest success-rate for single-step adversaries while for the *Base Model* $\varepsilon = 8$ works best (see Figures 6 (a) and (b)). This proves once again, that the adversarial spots are just a little further away compared to the *Base Model*. This may be an explanation, on why the VAT model performs rather good on our generated adversaries. However, the VAT model is not substantially more robust against adversarial attacks than the unprotected base model – it only requires a larger ε to fool it.

In contrast, the PGD model and our proposed method clearly flatten the decision space and thereby strongly reduce

the risk of adversarial spots (cf. Figure 6 (c) and (d)). As one can see, adversarial examples occur in greater distance compared to the *Base Model* and occur significantly less frequent as the lower loss values demonstrate. However, there are adversarial sub-spaces which cause a high distraction of the PGD models. In contrast to the PGD model, our model (Virtual Alpha with $p = 2$, $q = 4$ and $\lambda = 10$) provides a significantly more flattened space, i.e., high loss values rarely occur at all.

Noticeably, for $\varepsilon < 16$ the loss values of our model are significantly higher than the ones of the PGD model. In addition, loss values grow rapidly for small ε -values. The high average loss values for small ε -values may explain *why* our model has a lower accuracy on our generated adversarial datasets than the PGD model as the loss values do not need to be maximal in order to indicate misclassification.

VII. CONCLUSION

This work investigates the effects of adversarial attacks on deep learning networks. It analyzes different strategies for increasing a model’s resilience and, thus, countervailing malicious attacks. The performance of the different defense strategies is compared across large sets of transferable, carefully generated adversaries. Next, three new approaches to improve resilience against such perturbations were first introduced and then compared against the state-of-the-art techniques *Virtual Adversarial Training* and *Projected Gradient Descent*. In addition, a novel adversarial attack method called *gradient ensemble attack* has been introduced. Further, this work has demonstrated the transferability of adversaries, which have been crafted using our proposed method.

Within our investigations we have observed that VAT does not provide substantial resilience against adversarial perturbations as the adversarial sub-spaces are just pushed a little further away. However, the incidence of these spaces is similar to an unprotected model. Further, PGD trained models reduce the frequency of adversarial sub-spaces and strongly increase the distance to them. Yet, these sub-spaces still occur. Our proposed method, superimposing two images and minimizing the KL-Divergence between the two activations, reduces the risk of adversarial sub-spaces with high loss. In fact, our results demonstrate that these spaces rarely occur. However, the average loss value is significantly higher which explains why our models performed worse on our adversarial test sets. Nevertheless, our proposed method is very promising as it (i) is easily executable (it only requires an additional forward pass), and (ii) still provides a noticeable regularized decision space.

Our ideas for future work are two-fold: (i) we will compare additional methods to further decrease the overall loss of our proposed method and thereby improve its performance on adversaries; (ii) we will investigate the effects of our *gradient ensemble attack* for crafting strong and transferable adversaries in a wider context – especially applying it to different white- and black-box attack scenarios.

ACKNOWLEDGMENT

All three authors acknowledge support by the *European Research Center for Information Systems (ERCIS)*.

REFERENCES

- [1] W. S. McCulloch and W. Pitts, “A Logical Calculus of the Ideas Immanent in Nervous Activity,” *The Bulletin of Math. Biophysics*, vol. 5, no. 4, pp. 115 – 133, 1943.
- [2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [3] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer, 2001.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” *IEEE CVPR*, pp. 770 – 778, 2016.
- [5] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” in *ICLR*, 2015.
- [6] O. Abdel-Hamid, A. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu, “Convolutional Neural Networks for Speech Recognition,” *IEEE/ACM Trans. on Audio, Speech & Language Processing*, vol. 22, no. 10, pp. 1533 – 1545, 2014.
- [7] A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune, “Synthesizing the Preferred Inputs for Neurons in Neural Networks via Deep Generator Networks,” in *NIPS*, 2016, pp. 3387 – 3395.
- [8] D. Erhan, Y. Bengio, A. Courville, and P. Vincent, “Visualizing Higher-Layer Features of a Deep Network,” University of Montreal, Tech. Rep. 1341, 2009.
- [9] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, “Learning Deep Features for Discriminative Localization,” in *CVPR*, 2016, pp. 2921 – 2929.
- [10] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing Properties of Neural Networks,” in *ICLR*, 2014.
- [11] I. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and Harnessing Adversarial Examples,” in *ICLR*, 2015.
- [12] S. Gu and L. Rigazio, “Towards Deep Neural Network Architectures Robust to Adversarial Examples,” *CoRR*, vol. abs/1412.5068, 2014.
- [13] P. Tabacof and E. Valle, “Exploring the Space of Adversarial Images,” in *IJCNN*. IEEE, 2016, pp. 426 – 433.
- [14] A. Fawzi, S. Moosavi-Dezfooli, and P. Frossard, “Robustness of Classifiers: From Adversarial to Random Noise,” in *NIPS*, 2016, pp. 1624 – 1632.
- [15] N. Carlini and D. A. Wagner, “Defensive Distillation is Not Robust to Adversarial Examples,” *CoRR*, vol. abs/1607.04311, 2016.
- [16] N. Carlini and D. Wagner, “Towards Evaluating the Robustness of Neural Networks,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 39 – 57.
- [17] A. Kurakin, I. J. Goodfellow, and S. Bengio, “Adversarial Examples in the Physical World,” in *ICLR*, 2017.
- [18] N. Papernot, P. D. McDaniel, I. J. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical Black-Box Attacks against Machine Learning,” in *ASIA CCS*. ACM, 2017, pp. 506 – 519.
- [19] S. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “DeepFool: A Simple and Accurate Method to Fool Deep Neural Networks,” in *CVPR*. IEEE, 2016, pp. 2574 – 2582.
- [20] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards Deep Learning Models Resistant to Adversarial Attacks,” in *ICLR*, 2018.
- [21] A. Kurakin, I. J. Goodfellow, and S. Bengio, “Adversarial Machine Learning at Scale,” *CoRR*, vol. abs/1611.01236, 2016.
- [22] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, “Boosting Adversarial Attacks with Momentum,” in *CVPR*, 2018, pp. 9185 – 9193.
- [23] F. Tramèr, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “The Space of Transferable Adversarial Examples,” *CoRR*, vol. abs/1704.03453, 2017.
- [24] R. Huang, B. Xu, D. Schuurmans, and C. Szepesvári, “Learning with a Strong Adversary,” *CoRR*, vol. abs/1511.03034, 2015.
- [25] U. Shaham, Y. Yamada, and S. Negahban, “Understanding Adversarial Training: Increasing Local Stability of Neural Nets through Robust Optimization,” *Neurocomputing*, vol. 307, pp. 195 – 204, 2015.
- [26] T. Miyato, S. Maeda, M. Koyama, K. Nakae, and S. Ishii, “Distributional Smoothing with Virtual Adversarial Training,” in *ICLR*, 2016.
- [27] S. Kullback and R. A. Leibler, “On Information and Sufficiency,” *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79 – 86, 1951.
- [28] F. Tramèr, A. Kurakin, N. Papernot, I. J. Goodfellow, D. Boneh, and P. D. McDaniel, “Ensemble Adversarial Training: Attacks and Defenses,” in *ICLR*, 2018.
- [29] A. Athalye, N. Carlini, and D. Wagner, “Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples,” in *ICML*, vol. 80. PMLR, 2018, pp. 274 – 283.
- [30] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *IJCV*, vol. 115, no. 3, pp. 211 – 252, 2015.
- [31] A. Krizhevsky, “Learning Multiple Layers of Features from Tiny Images,” Master’s thesis, University of Toronto, ON, Canada, 2009.
- [32] M. Lin, Q. Chen, and S. Yan, “Network In Network,” in *ICLR*, 2014.
- [33] S. Ioffe and C. Szegedy, “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,” in *PMLR*, vol. 37, 2015, pp. 448 – 456.