

The Efficacy of SHIELD under Different Threat Models

Paper Type: Appraisal Paper of Existing Method

Cory Cornelius
cory.cornelius@intel.com
Intel Corporation
Hillsboro, OR, USA

Li Chen
li.chen@intel.com
Intel Corporation
Hillsboro, OR, USA

Nilaksh Das
nilakshdas@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Michael E. Kounavis
michael.e.kounavis@intel.com
Intel Corporation
Hillsboro, OR, USA

Shang-Tse Chen
schen351@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Duen Horng (Polo) Chau
polo@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

ABSTRACT

In this appraisal paper, we evaluate the efficacy of SHIELD [1], a compression-based defense framework for countering adversarial attacks on image classification models, which was published at KDD 2018. Here, we consider alternative threat models not studied in the original work, where we assume that an *adaptive* adversary is aware of the ensemble defense approach, the defensive pre-processing, and the architecture and weights of the models used in the ensemble. We define scenarios with varying levels of threat and empirically analyze the proposed defense by varying the degree of information available to the attacker, spanning from a full white-box attack to the gray-box threat model described in the original work. To evaluate the robustness of the defense against an adaptive attacker, we consider the *targeted-attack success rate* of the Projected Gradient Descent (PGD) attack, which is a strong gradient-based adversarial attack proposed in adversarial machine learning research. We also experiment with training the SHIELD ensemble from scratch, which is different from re-training using a pre-trained model as done in the original work. We find that the targeted PGD attack has a success rate of 64.3% against the original SHIELD ensemble in the full white box scenario, but this drops to 48.9% if the models used in the ensemble are trained from scratch instead of being retrained. Our experiments further reveal that an ensemble whose models are re-trained indeed have higher correlation in the cosine similarity space, and models that are trained from scratch are less vulnerable to targeted attacks in the white-box and gray-box scenarios.

ACM Reference format:

Cory Cornelius, Nilaksh Das, Shang-Tse Chen, Li Chen, Michael E. Kounavis, and Duen Horng (Polo) Chau. 2019. The Efficacy of SHIELD under Different Threat Models. In *Proceedings of KDD 2019 Workshop on Learning and Mining for Cybersecurity (LEMNCS'19)*, Anchorage, Alaska, United States, August 5th, 2019 (LEMNCS @ KDD'19), 7 pages.
<https://doi.org/>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

LEMNCS @ KDD'19, August 5th, 2019, Anchorage, Alaska, United States

© 2019 Copyright held by the owner/author(s).

ACM ISBN .
<https://doi.org/>

1 INTRODUCTION

Adversarial examples are inputs that appear innocuous to humans but fool a machine learning model into making incorrect predictions [2]. Even the most accurate image classification models that rely upon state-of-the-art deep neural networks are potentially vulnerable to these maliciously crafted inputs [3–5]. Thus, it is an important research problem to develop defenses against such adversarial examples, especially when these models are deployed in safety or security critical systems.

One such recently proposed defense is SHIELD, *Secure Heterogeneous Image Ensemble with Localized Denoising* [1], which uses a combination of techniques to defend against adversarial examples for image classification systems. SHIELD trains several image classification models using JPEG-compressed images at different compression levels. These JPEG-trained models are used in a majority vote ensemble to yield the final classification result for a given input image. At inference time, SHIELD also applies *Stochastic Local Quantization* (SLQ), which is a novel randomized form of JPEG compression, as a pre-processing step. In the face of a non-adaptive adversary that is only aware of the model architecture (but not its weights nor any defensive measures), SHIELD reports a 16.3% decrease in test accuracy against adversarial examples as compared to a 57.19% drop when there is no defense.

While experiment results in the original SHIELD paper emphasized its fast inference and practicality of being a readily available technique, it did not consider a full range of threat models or more recent stronger attacks. Thus, in this work, we aim to appraise SHIELD and make the following contributions:

Evaluating Threat Models with Adaptive Attacker. In this work, we further test the robustness of SHIELD against an *adaptive* adversary that has access to more parts of the defense than originally studied. We start with the white-box attack, in which the attacker has access to all parts of the inference pipeline, and peel back information given to the attacker by defining varying levels of threat models that are different from the gray-box threat model analyzed in [1]. We also find that SHIELD can be made more robust to targeted attacks even in the white-box case by using models in the ensemble whose weights are less correlated.

Evaluating Stronger Attacks. Since around the time SHIELD was published, the Projected Gradient Descent (PGD) [4] attack had

started to gain reputation in becoming one of the strongest gradient-based iterative attacks proposed in adversarial machine learning research. As PGD was not evaluated in the original SHIELD paper, in this work, we contribute evaluation focusing on PGD. Specifically, we use the *targeted* version of PGD to study the robustness of the defense framework under different threat models. We focus on reporting the *attack success rate* for targeting the *least-likely* model prediction, as it is more compelling if an attack is able to affect the prediction of a model significantly, opposed to only slightly changing it. Since we target the least likely model prediction for the attack, it is only easier for the adversary to target other predictions as they already have a higher rank in the prediction likelihood.

Analyzing the Effect of Training Procedure on Robustness of SHIELD. We take a closer look into the training procedure of SHIELD, whose original approach of creating an ensemble of JPEG-trained models involves re-training the models from a pre-trained ResNet-50 v2 model. We find that this re-training approach yields an ensemble of models whose weights are highly correlated, making the defense less robust to attacks when the attacker has access to one or more of the models in the ensemble. To explore how to alleviate this issue, we experiment with newly trained models that have been trained on JPEG-compressed images from scratch, and perform the same analysis for comparison. We find that the training procedure indeed affects the robustness of the defense at reducing the success of strong targeted attacks. In the full white-box case, simply training the models of the ensemble from scratch can reduce the attack success rate by 24%.

2 BACKGROUND

2.1 SHIELD

SHIELD is a defense framework published in KDD 2018 that employs (1) Stochastic Local Quantization and central cropping as pre-processing; (2) model re-training; and (3) an ensemble of models to defend against adversarial attacks on image classification systems. Below, we summarize the various components of the SHIELD defense framework.

2.1.1 Stochastic Local Quantization (SLQ). SLQ is a randomized form of image compression introduced in [1]. This method leverages JPEG compression as a form of defense. Since JPEG compression performs quantization in the frequency domain, the authors hypothesize that this step also removes adversarial perturbations from the image. The authors also argue that since the quantization step is non-differentiable, the attacker cannot obtain any useful gradients to directly perform gradient-based attacks. Additionally, SLQ uses randomization to further mask the gradients by breaking up the image into smaller blocks, and applying a different JPEG compression level to each block.

More concretely, let us denote the input image as x , and the JPEG compression at quality level q applied to the input image as $JPEG(x, q)$. SLQ uses K qualities $\{q_1, \dots, q_K\}$ to compute $JPEG(x, q_1), \dots, JPEG(x, q_K)$, and use these images to randomly pick patches from the corresponding locations to stitch up the final image.

2.1.2 Training Models on JPEG-compressed Images. Since SHIELD employs SLQ pre-processing that uses different levels of JPEG compression, the authors of [1] also re-train the image classification

model to be more robust to the pre-processing operation. Specifically, SHIELD takes a pre-trained ResNet-50 v2 model, denoted as \mathcal{M} , and re-trains it on images compressed at JPEG quality q to obtain a new ResNet-50 v2 model \mathcal{M}_q . The weights of \mathcal{M} are used as the initial weights while training \mathcal{M}_q to get faster convergence. SHIELD trains multiple models in this manner using different JPEG quality levels $\{q_1, \dots, q_K\}$ to obtain $\mathcal{M}_{q_1}, \dots, \mathcal{M}_{q_K}$, which are then used as a majority vote ensemble.

2.2 Projected Gradient Descent

Given a benign input instance x , a targeted adversarial attack aims to find a small perturbation δ that changes the prediction of model \mathcal{M} to a target class t different from the true class y , i.e., $\mathcal{M}(x+\delta) = t$, where $t \neq y$, $\|\delta\| \leq \epsilon$. We call ϵ the perturbation strength with respect to a specific norm $\|\cdot\|$. In this work we only examine strength in terms of ℓ_∞ because it is easy to understand—the maximum pixel value deviation an attacker can apply across all image channels.

Projected Gradient Descent (PGD) [4] is one of the strongest first-order attacks in the adversarial machine learning literature. PGD iteratively minimizes a loss function $L(x, t)$. In this work, we use cross-entropy for L that computes the distance between the softmax of the logit layer and the one-hot representation of the target class t . We denote the logit layer of model \mathcal{M} as $f(x, \mathcal{M})$. In each iteration i , it computes the direction of perturbation by taking the sign of the gradient of loss function L with respect to the current perturbed instance x^i , and then perform a projection step back to the feasible set, i.e., within ϵ ball of the original instance while remaining a valid image:

$$x^{i+1} = Proj\left(x^i - \alpha \cdot \text{sign}(\nabla_{x^i} L(x^i, t))\right)$$

Instead of starting from the original instance, we use a random starting point within ϵ ball of the original instance [4]. Following [6], we use the least likely class y_{LL} predicted by the model as our target class.

3 APPRAISAL METHODOLOGY

When deploying a defense to mitigate adversarial examples, one must understand the limits of that defense in terms of the strength of the attacker and threat models it was tested for. We usually quantify a defense's capability as a well-defined threat model along with a security curve that plots the accuracy of the model in the face of an adversary that has the ability to manipulate the image at different strengths for that threat model [7]. We measure strength as the difference between a test image and its corresponding adversarial image generated via an attack method. Such differences are high-dimensional so we often summarize this difference by computing some metric like ℓ_∞ (largest absolute difference) or ℓ_2 (euclidean distance). At the limit of the strength (i.e., the ability to manipulate pixels with any value), the attacker can simply use those examples that are already incorrectly classified or just replace the image with another image of their desired target. It turns out, however, that many models exhibit weaknesses at much smaller strengths where it is difficult for humans to visually distinguish any adversarial perturbation. Note that the study of these security curves is of interest to defenders due to the asymmetry between adversaries

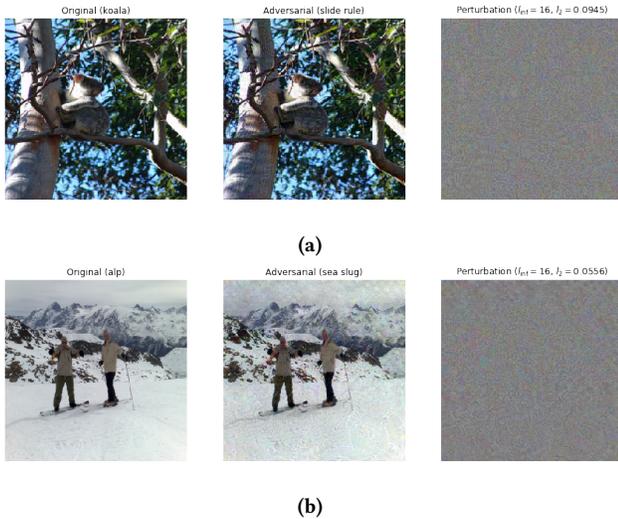


Figure 1: Example adversarial images generated against SHIELD. The first column shows the original images and their corresponding predictions. The second column shows the adversarial images their corresponding predictions. The final image is the difference between the original and adversarial image, which we call the perturbation. Note that both perturbations have ℓ_∞ distance of 16 but different ℓ_2 distances. The perturbation in the *alps* photo is visually noticeable, while the perturbation in the *koala* photo is not. Despite being visually noticeable, the *alps* perturbation has lower ℓ_2 than the *koala* perturbation. Ideally a chosen distance would be well ordered with respect to visual distinguishability. Images are best viewed in color.

and defenders. That is, adversaries need only find one particular weakness while defenders ought to mitigate all possible weaknesses.

In this work, we examine strength in terms of ℓ_∞ . Figure 1 shows two different adversarial images with their corresponding perturbations. Both perturbations have ℓ_∞ of 16 yet exhibit different ℓ_2 distances. The ℓ_2 distance of the *koala* perturbation is larger than the ℓ_2 distance of the *alps* perturbation, yet the *alps* perturbation is easily discernible in the adversarial image. Finding better metrics to summarize adversarial perturbations that take into account human perception remains an open problem. Because SHIELD was evaluated against ImageNet, we chose the same attacker strength that other state-of-the-art defenses on ImageNet use [8]: an ℓ_∞ of 16 out of 255.

3.1 Studying Effect of Training Procedure on Ensemble Robustness

Since the SHIELD framework deploys an ensemble of models re-trained from the same ResNet-50 v2 model \mathcal{M} , we posit that the model weights may be highly correlated, making it easier for the attacker to attack the ensemble even under weak gray-box restrictions in which the attacker has access to only some of the models from the ensemble. To study the effect of this potential issue, we train the

models from scratch on JPEG-compressed images of corresponding qualities used in the original work.

In [1], the model *re-trained* from a pre-trained ResNet-50 v2 model \mathcal{M} using images with JPEG quality q were denoted as \mathcal{M}_q . We call these models as “**Derivative**” models. Adopting a similar notation, we call the new ResNet-50 v2 model trained from scratch on images with JPEG quality q as \mathcal{M}_q^* , and we refer to them as “**Originative**” models. Hence, we perform our analysis on two ensembles, the newly trained $\{\mathcal{M}_{20}^*, \mathcal{M}_{40}^*, \mathcal{M}_{60}^*, \mathcal{M}_{80}^*\}$ (or SHIELD-Originative) and the original $\{\mathcal{M}_{20}, \mathcal{M}_{40}, \mathcal{M}_{60}, \mathcal{M}_{80}\}$ (or SHIELD-Derivative).

Also, to quantify the correlation between 2 models, we flatten, normalize and concatenate the weights from each layer of a model to form a single vector, and calculate the cosine similarity with the corresponding vector for the other model. Indeed, we find that SHIELD-Derivative (which has been obtained using initial weights from the same pre-trained ResNet-50 v2 model) has a higher average cosine similarity of 0.64 between any pair of models in the ensemble. Whereas, SHIELD-Originative (whose models have been trained from scratch) has a lower average cosine similarity of 0.42 between the pairs of models. One can further decrease this correlation by explicitly adding the similarity measure to the loss function as a regularizer, and train an ensemble jointly to make it more robust to transferability attacks.

3.2 Extending to More Threat Models

To fully understand the robustness of the proposed defense under different circumstances, it is important to set meaningful restrictions on the capabilities of the attacker according to the application domain [9]. Hence, we define varying degrees of threat models, and evaluate the proposed SHIELD defense under different scenarios. These threat models vary from a full white-box attack (no restrictions) to strong gray-box restrictions as originally studied in [1]. Table 1 summarizes these threat models, and we describe them in more detail below.

TM-White. This threat model corresponds to the full white-box attack that imposes no restrictions on the attacker, who is allowed to access to all parts of the proposed defense. For SHIELD, this includes the model architecture, weights, any JPEG pre-processing (i.e., SLQ), as well as all the defense parameters (e.g., SLQ in [1] uses JPEG qualities 20, 40, 60 and 80). In practice, attackers typically do not have full access to all information of the defense (if they do, that could mean an attacker is in complete control of the defender and could do anything they want). However, studying white-box attacks help researchers better understand a defense’s robustness under the most hostile situation.

TM-Gray1. This is a gray-box threat model that imposes weak restrictions on the attacker. Namely, the attacker now only has access to *some* of the models’ weights from the ensemble. The attacker is still aware of the defense strategy and the defense parameters. Since SHIELD uses 4 models in the ensemble, we study the robustness of the defense by varying the number of models that the attacker can access, and report the average attack success rate. For example, if the attacker can access 2 out of the 4 models from the ensemble, this yields 6 combinations of models that the attacker can attack,

Threat Model (TM)	Attacker has access to			
	Model Architecture	Weights of Models in Ensemble	Defense Strategy	Defense Parameters
TM-White	Yes	All	Yes	Yes
TM-Gray1	Yes	Some Models	Yes	Yes
TM-Gray2	Yes	No	Yes	Yes
TM-SHIELD	Yes	No	No	No

Table 1: Threat models studied in this work, ordered by decreasing severity. We define different scenarios with varying levels of threat. TM-White (top) allows the attacker to access the whole inference pipeline including the defense parameters, whereas TM-SHIELD (bottom), originally studied in [1], only provides the model architecture to the attacker.

i.e., $\binom{4}{2} = 6$, and we report the average success rate from all the 6 attack trials in this case.

TM-Gray2. This is a moderate gray-box threat model that imposes some more restrictions on the attacker. Specifically, now the attacker is aware of the the model architecture and the defense strategy and parameters, but cannot access any of the model weights and is unaware of the training strategy. In this scenario, the attacker is faced with training their own models on JPEG-compressed images. We evaluate this scenario by creating adversarial examples using the newly trained *Originative* models $\{M_q^*\}$ and testing them against with SHIELD’s default *Derivative* models $\{M_q\}$.

TM-SHIELD. This corresponds to the original gray-box threat model studied in [1] which imposes strong restrictions on the attacker. In this scenario, the attacker is only aware of the model architecture used in the ensemble and is oblivious of the defense pre-processing and does not have access to the model weights.

3.3 Constructing and Evaluating Adaptive Attacks

SHIELD integrates several techniques that mask the gradients of a model so that an adversary cannot perform an exact gradient-based attack that uses backpropagation. This backpropagation computation is exactly the same method one would use to learn such a model, except that rather than changing the weights of the model, the attacker seeks to change only the input image. For this work, we rely upon the PGD attack as implemented in Cleverhans [10] (which contains technical details).

Backpropagation requires that all operations in the model be differentiable. By using JPEG compression, which is non-differentiable, SHIELD forces attackers to create JPEG compression approximations that are differentiable. A recent work [11] describes how to do this, and we apply the authors’ differentiable JPEG approximation to attack SHIELD. Similarly, because the majority voting of the ensemble of models in SHIELD is also non-differentiable, we approximate the majority vote by averaging the logits output of each model before applying the softmax function. One can also apply alternative ensemble approximations [12], but we found averaging to be effective.

A further difficulty in naively applying PGD to SHIELD is the stochasticity of the SLQ pre-processing step. Because SLQ introduces randomness in the input presented to the model, the model effectively sees a different, albeit similar, input image every time for

the same image. SHIELD relies upon its ensemble of JPEG trained models to ensure that the same input image even after SLQ pre-processing still maps to the same prediction. The attacker, however, typically has no control of this randomness. To circumvent this issue, we compress the image at multiple compression levels — 20, 40, 60 and 80 (the JPEG qualities used by SLQ) — and average over these compression levels to find an adversarial perturbation. More concretely, for the purpose of attacking, we compute the logits of a given ensemble \mathcal{E} for the input x as,

$$\frac{1}{|Q| \cdot |\mathcal{E}|} \sum_{M \in \mathcal{E}} \sum_{q \in Q} f(\text{JPEG}(x, q), M)$$

where Q is the set of JPEG compression levels used by SLQ and $f(x, M)$ is the logits output of the model M for input x .

4 RESULTS

For our experiments, we sampled 1000 of the 50000 ImageNet [13] validation images and used PGD with 20 iterations to generate adversarial images. We performed the targeted version of the PGD attack with the *least-likely* prediction of the model under attack as the targets. We report the *attack success rate*, which is the fraction of images that were incorrectly classified by the model with the target label specifically chosen by the attacker. The attack is not successful if the model misclassifies the image but the output label is not the same as the target label. In the following section, we now present our results on the threat models being studied.

4.1 White-box Attack (TM-White)

The results from a full white-box targeted attack are shown in Table 2. We see that only 17 of the 1000 adversarially crafted images are correctly classified by the original SHIELD model (SHIELD-Derivative). The remaining 983 were incorrectly classified as something other than the ground truth label.

Furthermore, an attacker is able to force SHIELD-Derivative to make a prediction of their choosing. Here we chose the prediction target to be the least-likely class that the differentiable version of ensemble predicts on the original image. Such a target should be difficult for an attacker to hit, yet our results show that 643 of those 1000 adversarial image were successfully predicted as the least-likely label. However, we see in the case of SHIELD-Originative that an ensemble having uncorrelated models can alleviate this problem and increase the defense robustness to a targeted attack.

Ensemble	White-box Attack Success	Acc.
SHIELD-Derivative	0.643	0.017
SHIELD-Originative	0.489	0.022

Table 2: White-box robustness to targeted attacks. SHIELD-Originative thwarts more targeted attacks (i.e., attacker is less successful), even when adversary has access to the entire inference pipeline, since models trained from scratch have lower correlation than models re-trained from the another model as in SHIELD-Derivative. Accuracy of both models are adversely affected.

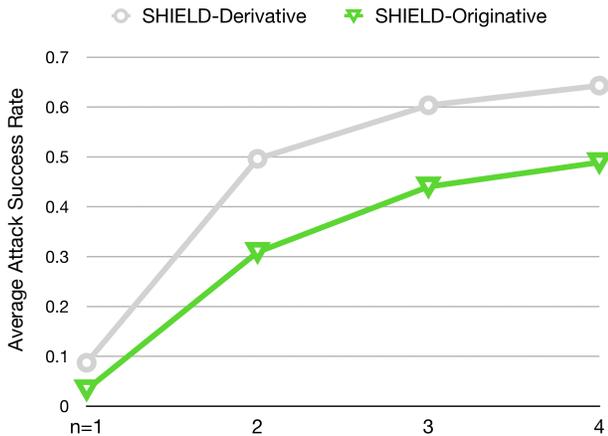


Figure 2: Comparison of the average attack success for TM-Gray1, where the attacker has access to n out 4 models from the ensemble. The attack is least successful in both cases when the attacker has access to only one model, and most successful when the attacker has access to all 4 models (TM-White). Models trained from scratch (SHIELD-Originative, in green) are significantly more robust to an adaptive targeted attack as compared to the original SHIELD ensemble, in gray.

The attacker is only able to target 489 images out of 1000 in this case.

4.2 Weak Gray-box attack (TM-Gray1)

In this threat model, the attacker can only access some of the models from the ensemble, while having full knowledge of the defense strategy. We perform this ablation study by averaging over the attack success rate from attacking all possible model combinations corresponding to the attacker’s budget constraint (i.e., the number of models the attacker can access). The most restrictive case is when the attacker can only access one of the models. In this scenario, the attacker has 4 choices for performing the attack (on any one of the 4 models from the ensemble). Hence the average attack success is reported from the 4 trials that the attacker can make. Similarly, the attacker has 6 choices when allowed to attack 2 models together, and so on.

The results of this experiment are shown in Figure 2, which plots the average attack success from attacking n out of 4 models, by varying n on the x-axis. As is evident, this study spotlights the effect of training from scratch and the risk of having models that are correlated. Note here that TM-White is an extreme case of this threat model when the attacker can access the whole ensemble ($n = 4$).

Even with just one model available to the attacker, the accuracy of the original ensemble (SHIELD-Derivative) falls to 20.33%, and further to 1.7% when all models are available to the attacker (TM-White). The SHIELD-Originative ensemble also shows similar trends with respect to accuracy. However, there is a clear contrast when it comes to the attack success rates for the two ensembles. The targeted attack has a lower success rate with the SHIELD-Originative ensemble, since the weights of the models in this ensemble are less correlated, and hence the attack does not transfer so well to the other models.

4.3 Moderate Gray-box attack (TM-Gray2)

This is a gray-box threat model that places more restrictions on the attacker. The attack can no longer access the weights of the ensemble and is unaware of the training procedure, but is aware of the model architecture and defense parameters. We believe such a threat model can be realistic since many applications employ existing architectures but with weights learned for their specific task. Furthermore, weights of a model can remain confidential by using a trusted execution environment or deployment as a cloud API. Finally, we believe our assumption that the attacker is aware of a JPEG-based defense is reasonable. Like model architectures, the set of known defenses is small enough that an attacker can exhaustively try them until they are successful.

In this threat model, the attacker is forced to train their own proxy models for performing the attack. To mimic this, we first generated adversarial examples against newly trained models from the SHIELD-Originative ensemble using the adaptive attack procedure described above. We then fed these adversarial examples into the original SHIELD-Derivative model to determine whether they fool SHIELD. We also evaluate these attacked images using the original JPEG-trained models (\mathcal{M}_{20} , \mathcal{M}_{40} , \mathcal{M}_{60} , \mathcal{M}_{80}) for comparison.

The results from this experiment are shown in Table 3. In all of these tests, we found that the adversary was unable to control the prediction target. That is, the adversarial image we generated did not classify to the ground truth label nor the targeted label, but to some other label. It is application dependent whether this ability of the attacker to control the predicted output is of concern. In authentication usages, for example, the attacker might impersonate a specific user with targeted attacks or simply any authorized user with untargeted attacks.

We also tested the Fast Gradient Method (FGM) [14] which is known to be transferable even across model architectures. Our results showed that FGM targeted attack in this scenario is still not successful in forcing a prediction on the ensemble, resulting in a 0% attack success rate, although it affects the model accuracy more adversely than PGD.

Attacked Ensemble	#model known	Defending Model				
		SHIELD-Derivative	\mathcal{M}_{20}	\mathcal{M}_{40}	\mathcal{M}_{60}	\mathcal{M}_{80}
SHIELD-Originative	1	0.000 (0.303)	0.000 (0.215)	0.000 (0.273)	0.000 (0.272)	0.000 (0.269)
	2	0.000 (0.265)	0.000 (0.184)	0.000 (0.240)	0.000 (0.237)	0.000 (0.228)
	3	0.000 (0.231)	0.000 (0.165)	0.000 (0.206)	0.000 (0.211)	0.000 (0.202)
	4	0.000 (0.214)	0.000 (0.139)	0.000 (0.185)	0.000 (0.198)	0.000 (0.185)
<i>No attack</i>		<i>0.633</i>	<i>0.538</i>	<i>0.597</i>	<i>0.574</i>	<i>0.603</i>

Table 3: Average attack success rate, and average model accuracy (in parenthesis) for the TM-Gray2 threat model. The targeted attack does not have any success in this scenario since the attacker does not have access to the model weights and the training procedure, and is thus forced to train their own models which are not guaranteed to transfer even with an adaptive targeted attack.

In this gray-box threat model, we see a degradation in the model accuracy of the ensemble as well as the JPEG-trained models; however, an attacker cannot control the prediction like they could in the white-box threat model.

4.4 Strong Gray-box attack (TM-SHIELD)

This is the threat model originally studied in [1] that places stronger gray-box restrictions on the attacker. The attacker in this scenario is only aware of the model architecture and cannot access any other parts of the inference pipeline. This threat model assumes that the attacker is non-adaptive. To mimic this, we turn off the differentiable JPEG component in our adaptive attack computation and attack a ResNet-50 v2 model that is not part of the original SHIELD ensemble. Although we observe a drop in model accuracy from 0.633 to 0.381 for the SHIELD-Derivative ensemble and from 0.77 to 0.423 for the SHIELD-Originative ensemble due to the added perturbation, we find that the non-adaptive targeted PGD attack is not successful with a 0% attack success.

5 CONCLUSION

As with any empirical security analysis, our results represent upper-bounds on the robustness of SHIELD. Attacks only get stronger. We believe that an adversary can better target gray-box attacks using a technique that recently won the both the non-targeted adversarial attack and targeted adversarial attack competitions at the 2017 Competition on Adversarial Attacks and Defenses [15], or adversaries will use query-based black-box methods to increase their targeted attack success [16]. More so, while understanding the limits of a model with respect to some ℓ_∞ limit is instructive, ultimately adversaries will find more clever ways to fool a model [17, 18].

Our main goal of this work is to understand the efficacy of SHIELD against adaptive attack techniques and under different threat models. We evaluate the robustness of the ensemble-based defense at resisting against targeted attacks. Our results show that SHIELD can be made more robust to such attacks in the white-box scenario by modifying the training procedure so as to deploy models whose weights are less correlated. This is further supported by our gray-box experiments, wherein the models trained from scratch are more robust to making maliciously targeted predictions. In this

work we simply train the models from scratch on different data distributions (by varying JPEG compression levels) to obtain models that have less correlated weights, but one can imagine a training paradigm that explicitly models this as a regularizer in the training optimization criteria.

We hope that future defenses will employ similar adaptive attack techniques to demonstrate their robustness at a variety of attack strengths and in the face of different, perhaps more realistic, threat models.

ACKNOWLEDGMENTS

We thank Richard Shin for sharing his differentiable JPEG implementation.

REFERENCES

- [1] N. Das, M. Shanbhogue, S.-T. Chen, F. Hohman, S. Li, L. Chen, M. E. Kounavis, and D. H. Chau, “Shield: Fast, practical defense and vaccination for deep learning using jpeg compression,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, Aug. 2018. <http://doi.acm.org/10.1145/3219819.3219910>.
- [2] I. J. Goodfellow, “Defense against the dark arts: An overview of adversarial example security research and future research directions,” *CoRR*, vol. abs/1806.04169, 2018. <https://arxiv.org/abs/1806.04169>.
- [3] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” in *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014. https://openreview.net/forum?id=kklr_MTHMRQJG.
- [4] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” in *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, Apr. 2018. <https://openreview.net/forum?id=rjzIBfZAb>.
- [5] N. Carlini and D. Wagner, “Adversarial examples are not easily detected: Bypassing ten detection methods,” in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security (AISec)*, Nov. 2017. <http://doi.acm.org/10.1145/3128572.3140444>.
- [6] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *arXiv preprint arXiv:1607.02533*, 2016.
- [7] B. Biggio and F. Roli, “Wild patterns: Ten years after the rise of adversarial machine learning,” *CoRR*, vol. abs/1712.03141, 2017. <https://arxiv.org/abs/1712.03141>.
- [8] A. Kurakin, I. J. Goodfellow, and S. Bengio, “Adversarial machine learning at scale,” in *International Conference on Learning Representations*, 2017.
- [9] J. Gilmer, R. P. Adams, I. Goodfellow, D. Andersen, and G. E. Dahl, “Motivating the rules of the game for adversarial example research,” *arXiv preprint arXiv:1807.06732*, 2018.
- [10] N. Papernot, F. Faghri, N. Carlini, I. Goodfellow, R. Feinman, A. Kurakin, C. Xie, Y. Sharma, T. Brown, A. Roy, A. Matyasko, V. Behzadan, K. Hambardzumyan, Z. Zhang, Y.-L. Juang, Z. Li, R. Sheatsley, A. Garg, J. Uesato, W. Gierke, Y. Dong, D. Berthelot, P. Hendricks, J. Rauber, and R. Long, “Technical report on the cleverhans v2.1.0 adversarial examples library,” *arXiv preprint arXiv:1610.00768*, 2018. <http://arxiv.org/abs/1610.00768>.

- [11] R. Shin and D. Song, "Jpeg-resistant adversarial images," in *Proceedings of the Machine Learning and Computer Security Workshop*, Dec. 2017. https://machine-learning-and-security.github.io/papers/mlsec17_paper_54.pdf.
- [12] W. He, J. Wei, X. Chen, N. Carlini, and D. Song, "Adversarial example defense: Ensembles of weak defenses are not strong," in *Proceedings of the 11th USENIX Workshop on Offensive Technologies (WOOT)*, (Vancouver, BC), USENIX Association, 2017. <https://www.usenix.org/conference/woot17/workshop-program/presentation/he>.
- [13] 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*, pp. 248–255, 2009.
- [14] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [15] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," in *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. <https://doi.org/10.1109/CVPR.2018.00957>.
- [16] A. Ilyas, L. Engstrom, A. Athalye, and J. Lin, "Black-box adversarial attacks with limited queries and information," in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, July 2018. <http://proceedings.mlr.press/v80/ilyas18a.html>.
- [17] Y. Sharma and P.-Y. Chen, "Attacking the madry defense model with l_1 -based adversarial examples," in *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, Apr. 2018. <https://openreview.net/forum?id=Sy8WeUJPF>.
- [18] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer, "Adversarial patch," in *Proceedings of the Machine Learning and Computer Security Workshop*, Dec. 2017. https://machine-learning-and-security.github.io/papers/mlsec17_paper_27.pdf.

A IMPLEMENTATION ISSUES

The first author of this work discovered an issue in the original SHIELD evaluation code implementation that affects the results reported in the published paper [1]. During training, central crops are enabled to better learn discriminative features for the object of interest. However, during evaluation it is customary to turn off this central cropping. The implementation of SHIELD does not turn off this cropping during evaluation nor did the attacks in SHIELD take into account this cropping when generating perturbations.

The first author of this work discovered this issue after noticing adversarial images were not working against the public implementation of SHIELD. He found that the images were adversarial against another implementation of SHIELD using the same SLQ and model parameters, but not against the originally released SHIELD implementation¹. After some investigation, he found the aforementioned central-cropping-at-evaluation-time issue, disabled this central cropping, and found that the adversarial images remained adversarial.

As such, central cropping is now a feature of SHIELD in the sense that the evaluation portion of the published paper [1] is SLQ *with central cropping*. Further experiments reveal that central cropping contributes significantly to the reported robustness of SHIELD. In this work, we do not include this cropping as part of the pre-processing step, and only consider SLQ as is.

¹<https://github.com/poloclub/jpeg-defense>
at commit 1576429cf199c38065b941a48b0fcd7747901457