

Unmasking the Vulnerabilities of Deep Learning Models: A Multi-Dimensional Analysis of Adversarial Attacks and Defenses

Firuz Juraev¹, Mohammed Abuhamad², Eric Chan-Tin², George K. Thiruvathukal², and Tamer Abuhmed¹

¹Department of Computer Science and Engineering, Sungkyunkwan University, Suwon, South Korea

²Department of Computer Science, Loyola University Chicago, IL, USA

Email: fjuraev@g.skku.edu, mabuhamad@luc.edu, chantin@cs.luc.edu, gkt@cs.luc.edu, tamer@skku.edu

Abstract—Deep Learning (DL) is rapidly maturing to the point that it can be used in safety- and security-critical applications, such as self-driving vehicles, surveillance, drones, and robots. However, adversarial samples, which are undetectable to the human eye, pose a serious threat that can cause the model to misbehave and compromise the performance of such applications. Addressing the robustness of DL models has become crucial to understanding and defending against adversarial attacks. In this study, we perform comprehensive experiments to examine the effect of adversarial attacks and defenses on various model architectures across well-known datasets. Our research focuses on black-box attacks such as *SimBA*, *HopSkipJump*, *MGAAttack*, and *boundary attacks*, as well as preprocessor-based defensive mechanisms, including *bits squeezing*, *median smoothing*, and *JPEG filter*. Experimenting with various models, our results demonstrate that the level of noise needed for the attack increases as the number of layers increases. Moreover, the attack success rate decreases as the number of layers increases. This indicates that model complexity and robustness have a significant relationship. Investigating the diversity and robustness relationship, our experiments with diverse models show that having a large number of parameters does not imply higher robustness. Our experiments extend to show the effects of the training dataset on model robustness. Using various datasets such as ImageNet-1000, CIFAR-100, and CIFAR-10 are used to evaluate the black-box attacks. Considering the multiple dimensions of our analysis, *e.g.*, model complexity and training dataset, we examined the behavior of black-box attacks when models apply defenses. Our results show that applying defense strategies can significantly reduce attack effectiveness. This research provides in-depth analysis and insight into the robustness of DL models against various attacks, and defenses.

Index Terms—Threat Analysis, Deep Learning, Black-box Attacks, Adversarial Perturbations, Defensive Techniques.

I. INTRODUCTION

Deep Learning (DL) has grown into a powerful and effective tool that contributes to solving a wide range of complicated learning tasks that were previously impossible to solve using conventional Machine Learning (ML) approaches. DL has been extensively used in a wide variety of modern day-to-day applications [4], [5], and it has made tremendous progress in terms of equal or beyond human-level performance. Consequently, DL-based solutions are widely used

in safety-, privacy-, and security-critical situations, such as malware detection [23], self-driving vehicles, drones, personal assistants, smart homes, and robots. Despite the success of DL, concerns have been raised regarding efficiency, energy consumption, data and model privacy, and susceptibility to adversarial examples. After the findings of Szegedy *et al.* [24], several intriguing studies appeared in the research community involving adversarial attacks on DL. For instance, vulnerabilities of automatic speech recognition [7], voice controllable systems [28], and autonomous vehicles [16] have been demonstrated by adversarial attacks. Existing adversarial attacks can be classified into three categories: **white-**, **gray-**, and **black-box attacks**. The adversaries' knowledge is what distinguishes the three attacks. Generally, DL models are of black-box nature, and black-box attacks are more realistic in many cases and are closer to real-world challenges.

Black-box attacks target the model with no knowledge of the systems' inner processes. Such attacks are relatively difficult to conduct and alarming-to-devastating when they succeed [14], [3]. We shed light on the characteristics, effects across multiple dimensions, and defenses against black-box attacks. In this work, we focus on the most frequently accessible methods among black-box attacks, such as Simple Black-box Attack (*SimBA*) [12], *HopSkipJump* [9], Microbial Genetic Algorithm-based Attack (*MGAAttack*) [25] and **boundary attacks** [6]. We study the behavior of these attacks in different settings and various models. Recent studies of defensive strategies against adversarial attacks suggest a plethora of trainer-, detector-, and preprocessor-type defenses. Unlike trainer- and detector-type defenses, preprocessor-type defenses neither optimize the model's weight nor change the model's architecture. Instead, the adversarial input is preprocessed to make it less-to-non-adversary. This work also investigates the effects of commonly used preprocessor-type defenses (*e.g.*, bit squeezing [27], median smoothing [18], and JPEG filter [10]) and Adversarial training [21] on black-box attacks.

In addition, this work examines the complexity, diversity, and robustness relationship among various DL model families. In the case of complexity-robustness relationship, we study **ResNet**, **VGG**, and **DenseNet** families. For the diversity-robustness relationship, we explore **InceptionV3**, **Xception**, **GoogleNet**, **MobileNetV2**, and **ShuffleNetV2**. Interestingly,

The code, trained models, and detailed settings are available at <https://github.com/InfoLab-SKKU/Adversarial-Attacks-Analysis>
Corresponding author: Tamer Abuhmed (tamer@skku.edu)

it has long been assumed that larger models with more trainable parameters outperform smaller models because they can capture more complex patterns in the data. However, our in-depth analysis and experiments on the relationship between the complexity and robustness of models show that the number of parameters in a model is not a concrete indicator of robustness. In fact, our results indicate that middle-size models may be more robust against adversarial attacks compared to heavy-weight models. For example, models like VGG-19 and ResNet-152, which have a large number of parameters (143,667,240 and 60,192,808, respectively), can be more vulnerable to attacks compared to middle-weight models such as Xception (22,855,952 parameters) and GoogLeNet (6,624,904 parameters). On the other hand, our results suggest that light-weighted models like MobileNet V2 and ShuffleNet V2 (3,504,872 and 2,278,604 parameters, respectively) can also be less robust compared to larger models. This study highlights the importance of taking steps to ensure that the model (of any size) has sufficient generalization and robustness capabilities.

This research also investigates the effects of using different training datasets on attacking various models. Although the sensitivity of models to adversarial examples in the image context and throughout the literature was typically assessed using specific datasets, Carlini *et al.* [8] showed that DL approaches developed using some datasets may not always extend to other datasets. For example, CIFAR [15] and MNIST [17] datasets, compared to ImageNet [20], include images with lower resolution and fewer categories. Therefore, most of the research efforts recently have favored ImageNet over CIFAR and MNIST [12]. Our analysis is done on three datasets, *i.e.*, ImageNet-1000, CIFAR-100, and CIFAR-10.

Contributions. In this work, we investigate the various aspects of black-box attacks and preprocessor-type defenses on the state-of-the-art models, including **ResNet**, **VGG**, **DenseNet** families using benchmark datasets such as ImageNet-1000, CIFAR-100, and CIFAR-10. Furthermore, different architectures are examined for the diversity-robustness relationship of models including, **InceptionV3**, **Xception**, **GoogleNet**, **MobileNetV2**, and **ShuffleNetV2**. In order to address the following research questions, large-scale experiments are performed and analyzed:

- ① The influence of model complexity on the effectiveness of various black-box attacks remains unclear. Specifically, *Does the model's complexity impact the attack's success/failure rate?* To shed light on this aspect, our study investigates the relationship between the complexity and robustness of models (from multiple DL families) across various black-box attacks and datasets.
- ② We investigate whether there exists a relationship between the adopted architecture/design and the robustness of models against adversarial attacks. Specifically, *do model design choices impact robustness?* We examine the behavior and vulnerability of diverse models under adversarial conditions.

- ③ *When adversarial examples are provided to the model, how do models address the attack problem while implementing defenses?* The most commonly used preprocessor defense strategies are analyzed to address this question.

Organization. The remainder of the paper is organized as follows: Section II highlights the methods and the essential principles for black-box attacks and defensive strategies. Section III discusses the experimental findings and analysis. The results are summarized in Section IV and the open challenges in attacks and defenses. The paper is concluded in Section V.

II. EXPERIMENTAL SETTINGS

Dataset. ImageNet [20], CIFAR-10 [15], and CIFAR-100 [15] are three common datasets used in all experiments. The ImageNet dataset, with a 224×224 image size and 14 million samples, is organized into 1000 classes. The CIFAR-100 dataset, with a 32×32 image size and 60,000 samples, is divided into 100 classes. The CIFAR-10 dataset contains 60,000 images with a 32×32 image size for 10 classes. Due to the time complexity of the black-box attacks, all of the results are centered on test data of 1000 images. These images are selected based on two conditions: ① one example from each class of ImageNet, and ② each image must be classified correctly by all selected models. In both CIFAR-100 and CIFAR-10 datasets, a balanced sampling approach was used to select an equal number of images from each respective class, and all selected images are classified correctly by all selected models.

Models. To examine the complexity and robustness relationship (Section III-A), 12 models from the ResNet [13], VGG [22], and DenseNet families are employed. In the second set of experiments, seven distinct models are trained on each dataset, including GoogLeNet, InceptionV3, Xception, MobileNetV2, and ShuffleNetV2, to explore the model diversity-robustness relationship (Section III-B). All of these models are used across our sets of experiments to explore the behavior of attacks in different settings. Pre-trained models for the ImageNet-1000 dataset based on the Pytorch framework were employed. We trained all models for the CIFAR-100 and CIFAR-10 datasets from scratch. For both CIFAR-100 and CIFAR-10 datasets, the training continues for 200 epochs with a batch size of 128 and a learning rate of 0.01. The accuracy of all models across datasets is shown in Tables I and II.

Evaluation Metrics. The used metrics are: ① *Misclassification confidence (MC)* [29]: the probability that a model assigns incorrectly to classify an adversarial image. ② *Attack time*: the estimated time taken to carry out a successful attack. We presented an average attack time in seconds for each model. ③ *Noise rate*: the degree of noise is determined using the Structural Similarity Index Measure (SSIM) [26]. SSIM measures how two images are similar to each other. To calculate the noise rate, we subtracted the SSIM value from 1 (*i.e.*, $\text{noise rate} = 1 - \text{SSIM}$). ④ *Attack success rate* [29]:

TABLE I
THE ACCURACY OF MODELS USING DIFFERENT DATASETS FOR EXP 1.

Model	Top 5% Accuracy (ImageNet)	Accuracy (CIFAR 100)	Accuracy (CIFAR 10)
ResNet 18	89.08%	75.94%	94.79%
ResNet 34	91.42%	76.72%	94.77%
ResNet 50	92.86%	76.96%	95.30%
ResNet 101	93.55%	76.22%	94.99%
ResNet 152	94.05%	74.88%	94.60%
VGG 11	88.63%	71.08%	91.94%
VGG 13	89.25%	73.50%	93.70%
VGG 16	90.38%	72.86%	93.30%
VGG 19	90.88%	71.84%	93.18%
DenseNet 121	91.97%	77.54%	95.06%
DenseNet 169	92.81%	78.67%	95.10%
DenseNet 201	93.37%	78.63%	95.17%

TABLE II
THE ACCURACY OF MODELS USING DIFFERENT DATASETS FOR EXP 2.

Model	Top 5% Accuracy (ImageNet)	Accuracy (CIFAR 100)	Accuracy (CIFAR 10)
VGG 19	90.88%	71.84%	93.18%
ResNet 152	94.05%	74.88%	94.60%
DenseNet 169	93.56%	79.05%	95.43%
Inception V3	93.45%	77.12%	94.77%
Xception	94.29%	74.67%	93.49%
GoogLeNet	89.53%	75.93%	95.01%
MobileNet V2	90.28%	69.44%	94.20%
ShuffleNet V2	88.31%	64.19%	90.04%

TABLE III
THE PARAMETERS AND NUMBER OF LAYERS OF SELECTED MODELS FOR EXP 2.

Model	Number of parameters	Number of Layers	Activation functions
VGG 19	143,667,240	19	ReLU
ResNet 152	60,192,808	152	ReLU
DenseNet 169	14,149,480	169	ReLU
Inception V3	27,161,264	48	ReLU
Xception	22,855,952	71	ReLU
GoogLeNet	6,624,904	22	ReLU
MobileNet V2	3,504,872	53	ReLU
ShuffleNet V2	2,278,604	65	ReLU




Fig. 1. Relation between the number of layers and the amount of noise and time needed (green line) for an attack to succeed. Noise rate and time increase as the number of layers increases in models of different families in three black-box attacks.

the ratio of successful trials to total trials. ⑤ *Confidence score*: the model probability of assigning an input to an output.

Experiment Workstation. The experiments are conducted on a machine equipped with an Intel Xeon(R) CPU E5-2620 v3 @ 2.40 GHz × 24 with Cuda-10.0 and three GEFORCE RTX 2080 Ti 12 GB GPUs, as well as Python 3.7.7 distributed in Anaconda 4.8.3 (64-bit).

Attack algorithms settings. The used adversarial attacks: ① *Simple Black-box Attack (SimBA)* with the following parameters: epsilon is 0.5 (overshoot parameter), maximum iterations are 6000. ② *HopSkipJump Attack* with the following parameters: batch size is 64, and maximum iterations are 20. ③ *Boundary Attack* with the following parameters: epsilon is 0.01, and maximum iterations are 1000. ④ *MGA Attack* with the following parameters: epsilon is 0.047, maximum queries are 1000, and mutation rate is 0.001. The parameters of attacks are tuned to get the maximum attack success rate.

Defense algorithms settings. The used defense algorithms: ① *Bits Squeezing* with the parameter of bit depth = 4. ② *Median Smoothing* with the parameter of kernel size as 2. ③ *JPEG Filter* with the parameter of quality as 75. The parameters of defense algorithms are selected with minimal settings to keep the image quality.

III. RESULTS AND ANALYSIS

Adversarial perturbations are intriguing because the adversaries are quite observable to DNNs over a broad range of training regimens. First, we must distinguish between the model’s complexity and diversity. The number of layers employed in the network design is used to describe *model’s complexity*. On the other hand, *model diversity* is characterized




Fig. 2. Noise, Misclassification confidence (MC), and time needed by the HopSkipJump attack for ResNet-18, ResNet-50, and ResNet-152 models.

as distinct network architectures. In this section, the findings of the experiments are discussed.

A. Exp 1— Model Complexity and Robustness

The ImageNet dataset is used in this set of experiments to explore the complexity-robustness relationship of models. The complexity is implied by the depth and breadth of models, *i.e.*, the number of layers and the number of parameters. The robustness of models is measured by their ability to maintain high accuracy on the test set when attacked by adversarial examples. These experiments aim to explore whether there is a trade-off between model complexity and robustness. For a fair comparison, the same family of models is utilized with a varying number of layers, *e.g.*, ResNet family (ResNet-18, ResNet-34, ResNet-50, ResNet-101, and ResNet-152), VGG (VGG-11, VGG-13, VGG-16, and VGG-19), and DenseNet (DenseNet-121, DenseNet-169, and DenseNet-201). The experiments are repeated 10 times on the same selected image set to obtain more stable results.

Complexity via Model’s Depth. Figure 1 shows that as the number of layers increases, the noise rate steadily increases, indicating that a model becomes more resistant to adversar-

ial attacks as the number of layers increases. Furthermore, successfully attacking deeper models requires more time and computation. The figure also shows that the time (green line) needed to generate adversarial samples increases for deeper models within a family. The varying noise and attack times are evident across models from ResNet and DenseNet with a varying number of layers. However, because there is less variation in the number of layers of VGG models, the variance is less obvious. This suggests that as models become more complex (depth-wise), generating effective adversarial attacks becomes increasingly difficult and resource-intensive. For example, Figure 2 the correct label for a benign image is an airship. Attacking ResNet-18 requires 8.4% adversarial noise using the HopSkipJump attack in 51 seconds, while attacking ResNet-50 and ResNet-152 requires 16.3% and 28.6% adversarial noise in 59 and 78 seconds, respectively.

Complexity via Model’s Parameters. On the other hand, the complexity of the models in terms of the number of parameters might suggest different results. For example, using HopSkipJump and Boundary Attack (Figure 1-(b) and -(c)), VGG models (e.g., VGG-19 with 143,667,240 parameters) are more susceptible to perturbations than ResNet (e.g., ResNet-152 with 60,192,808 parameters) and DenseNet (e.g., DenseNet-169 with 14,149,480 parameters). This also holds true when comparing results from the ResNet and DenseNet families, as DenseNet models are more resilient to perturbations.

These results indicate that the model complexity (*i.e.*, number of layers and parameters)—alone does not necessarily determine the susceptibility to adversarial attacks and that the architectural design itself impacts the robustness. This encourages the following set of experiments to explore other factors, *e.g.*, model design.

Complexity against Sophisticated Attacks. MGAAttack adopts a distinct approach from previous attacks that use heuristics to generate adversarial samples with minimal noise through an evolutionary approach. This approach, among many emerging sophisticated methods, allows a more efficient and effective way of creating adversarial samples. Therefore, we investigated the relationship between model complexity and robustness against MGAAttack through two key metrics, namely attack success rate and attack time. Figure 3 illustrates that as the number of layers in the models increases, the success rate of the attack decreases and the number of queries for the attack increases. In terms of the number of parameters, similar observations as from previous experiments were made, where models with a larger number of parameters tended to be more susceptible to attacks.

B. Exp 2— Model Diversity and Robustness

To explore the effects of architectural design on the robustness of models, we use various models, including MobileNet-V2, ShuffleNet-V2, Inception-V3, Xception, GoogLeNet, VGG-19, ResNet-152, and DenseNet-169. In this set of experiments, diversity refers to the underlying designs of the models and their associated impacts on the number of parameters.




Fig. 3. The amount of noise and time needed (green line) for MGAAttack to succeed against various models. As the number of layers increases, the attack success rate decreases (above), and the number of queries increases (below).




Fig. 4. Attack needed noise across light-weight, middle-weight, and heavy-weight models when using SimBA, HopSkipJump, and Boundary attacks.

MobileNet-V2 and ShuffleNet-V2 are designed to be lightweight models, suitable for deployment on mobile and embedded devices with limited computational resources. They use depth-wise separable convolutions to reduce the number of parameters. MobileNet-V2 has a higher accuracy than ShuffleNet-V2, but ShuffleNet-V2 is faster and has fewer parameters (see Tables II and III). The Inception module was introduced by GoogLeNet (also known as Inception-V1), which allows for various filter sizes and reduces the number of parameters. It also contains auxiliary classifiers to aid in training. Based on the Inception architecture, Inception-V3, and Xception employ depth-wise separable convolutions to achieve high accuracy with fewer parameters. The VGG-19 design is straightforward, consisting of stacked convolutional layers with small filter sizes followed by max-pooling layers. It

TABLE IV
ATTACK ADDED NOISE ACROSS 8 DIVERSE MODELS WHEN USING SIMBA, HOPSKIPJUMP, AND BOUNDARY ATTACKS.

Model Type	Target Model	SimBA	HopSkipJump	Boundary Attack
Heavy-weight	VGG 19	0.035±0.041	0.024±0.036	0.027±0.039
	ResNet 152	0.039±0.043	0.055±0.061	0.062±0.069
	DenseNet 161	0.036±0.040	0.081±0.077	0.091±0.086
Middle-weight	Inception V3	0.038±0.046	0.058±0.071	0.071±0.091
	Xception	0.050±0.048	0.083±0.080	0.088±0.091
	GoogLeNet	0.045±0.044	0.081±0.088	0.088±0.093
Light-weight	MobileNet V2	0.029±0.036	0.028±0.038	0.031±0.044
	ShuffleNet V2	0.021±0.031	0.015±0.029	0.017±0.033

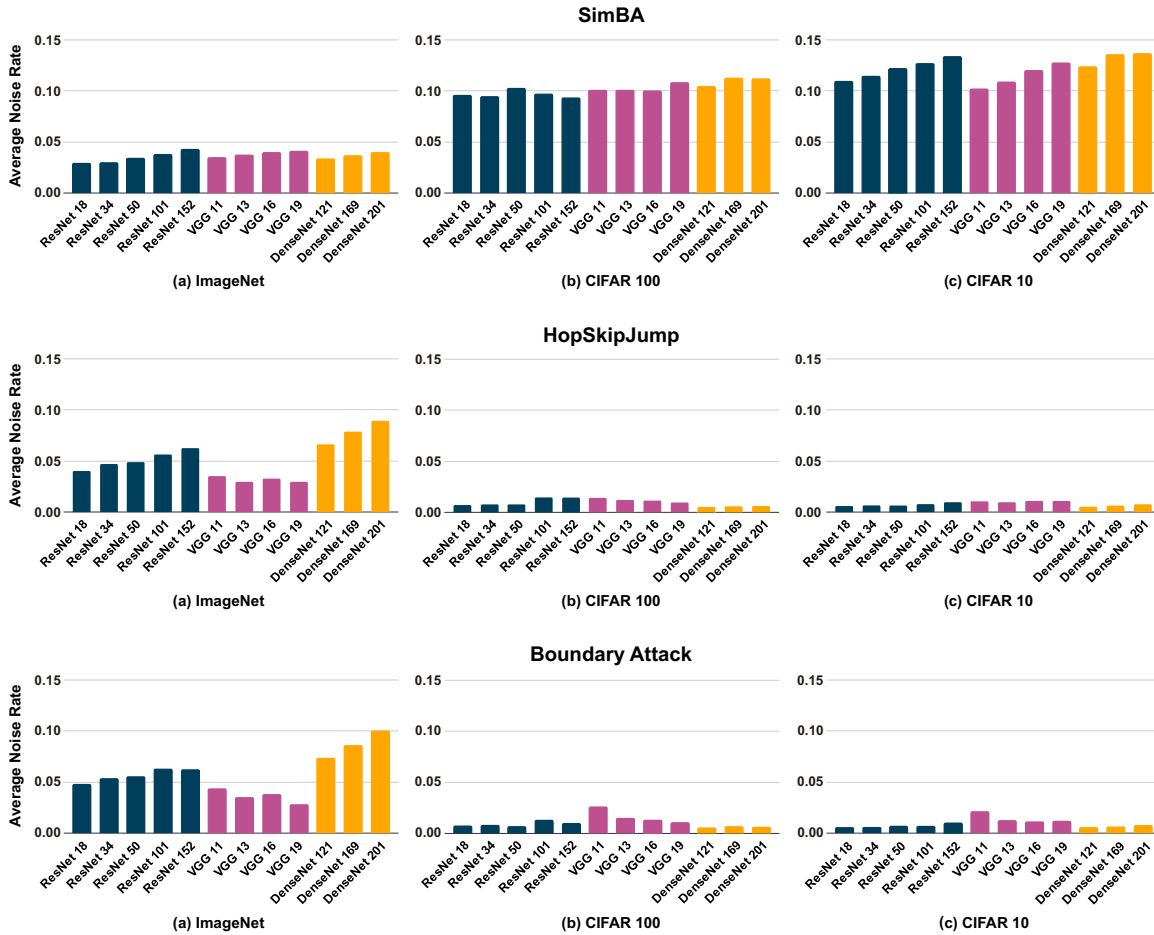


Fig. 5. The impact of datasets on needed noise for SimBA, HopSkipJump, and Boundary attacks to succeed against 12 models from VGG, ResNet, and DensNet families. HopSkipJump and Boundary Attack seem to be more successful when the number of classes and input size are smaller. SimBA requires more noise against models trained on CIFAR-10 and CIFAR-100.

has a large number of parameters and necessitates a substantial amount of computational power. ResNet-152 is a very deep model that employs residual connections to create very deep networks without vanishing gradients. The concept of dense connections, where each layer is linked to every other layer in a feed-forward manner, is the foundation of DenseNet-169. This results in feature reuse and a reduction in parameter count while increasing accuracy. Table III shows the selected models to explore the effects of architectural designs on robustness.

Depending on their parameter count, we group these models into three classes. ① *Light-weight* models with an average of 2,891,738 parameters, including MobileNet-V2 and ShuffleNet-V2. ② *Middle-weight* models such as Inception-V3, Xception, and GoogLeNet that have an average of 18,880,706 parameters. ③ The rest of the models are *Heavy-weight* with an average of 72,669,842 parameters.

Initially, and based on previous experiments (Exp 1 - Section III-A), the results show that the model complexity (*i.e.*, via the number of layers and parameters) would affect the robustness. Deeper networks tend to have higher resilience to perturbations in the input, whereas a larger number of parameters tends to make the model more susceptible to such perturbations. This set of experiments examines this assumption by employing a variety of models with varying

numbers of parameters and designs to explore whether the design impacts robustness.

Figure 4 shows that *heavy-weight* models are actually less robust compared to their *middle-weight* counterparts, as indicated by the average added noise. Considering only the number of parameters, the *middle-weight* models are still more robust than the *light-weight* despite the fewer parameters in *light-weight* models. This suggests that the number of parameters is not necessarily indicative of robustness and that model designs matter in terms of robustness. *Middle-weight* models, which feature various kernel sizes, auxiliary blocks, inception blocks, and stem designs, have been shown to be more robust. More robust models can be achieved through careful design choices rather than simply increasing the complexity of the model (more detailed results are in Table IV).

C. Exp 3— Training Dataset and Robustness

In our earlier experiments, we explored the behavior of attacks on the ImageNet dataset. In this set of experiments, we explore the attack behavior with different models when using different datasets with varying input sizes and classes. To this end, we use ImageNet, CIFAR-100, and CIFAR-10, to examine the behavior of SimBA, HopSkipJump, and Boundary attacks against 12 models from the VGG, ResNet,




Fig. 6. The impact of datasets on attacks targeting models with diverse architectural designs and sizes using different datasets.

and DenseNet families. Our results across different datasets are consistent with the results of previous experiments; *i.e.*, deeper networks within a model family tend to be more robust. Interestingly, the experiments with CIFAR-100 do not follow the same pattern. However, this might be due to the poor performance of models on CIFAR-100 as compared to their performance on other datasets (see Table I), as adversarial misclassification can be partially caused by poor generalization.

Number of Classes and Robustness. Based on our experiments, the success and stealthiness of adversarial attacks can depend on the specific attack strategy and the characteristics of the training dataset. For example, HopSkipJump and Boundary Attack seem to be more successful when the number of classes is smaller. SimBA, on the other hand, requires more noise to succeed against all models. The primary idea behind a SimBA is to move the prediction from the benign class to the neighboring class. As a result, for the ImageNet dataset, SimBA perturbs considerably less noise than HopSkipJump as can be seen in Figure 5. The dense distribution of ImageNet’s 1000 classes allows even a small amount of noise to change the prediction. However, due to the lower number of classes in the CIFAR-100 and CIFAR-10, SimBA introduces greater noise for input images to succeed.

Input Size and Robustness. Considering the results of HopSkipJump and Boundary Attack, the needed perturbations to succeed are significantly lower when attacking models trained in the CIFAR-10 and CIFAR-100 datasets than when attacking models trained on ImageNet (see Figures 5 and 6). The results suggest that both the number of classes and the size of the input may influence the robustness of the models. Additionally, evaluating the robustness of models across different datasets and input sizes is crucial to ensuring their reliability in real-world applications.

Generally, when using models with diverse architectural designs and sizes, the behavior of the attack depends on the dataset and attack strategy. For example, Figure 6 shows that *light-weight* models need more perturbations than *middle-weight* and *heavy-weight* models when adopting HopSkipJump and Boundary Attack on CIFAR datasets. On the other hand, using SimBA, the attack behavior is consistent with the patterns observed in ImageNet (*i.e.*, model design is crucial to robustness as *middle-weight* models tend to be more resilient than others to adversarial perturbations).




Fig. 7. Bit squeezing, median, and JPEG defensive techniques and the needed amount of noise to defend the model effectively. Under the HopSkipJump attack, bit squeezing seems to add more noise than median and JPEG filters, reducing its confidence score.

D. Exp 4— Defenses against Black-box Attacks

Arguably, defending against an attempted attack is as crucial as predicting an accurate result. The purpose of defense techniques is to minimize the effects of adversarial perturbations and protect models against adversarial attacks.

Preprocessor-type Defenses. In this set of experiments, we explore the effects of applying preprocessor-type defenses against black-box attacks. Using Bit Squeezing, Median Smoothing, and the JPEG Filter, our results show that most attacks are ineffective even with minimal parameter settings for defenders. The reason for this is that most adversarial attacks, such as boundary and SimBA attacks, stop perturbing noise after a class crosses its boundary and moves to another class region. Therefore, such attacks result in a low MC score, and adversarial images revert to the original class border with minimal defender parameters. For example, Figure 7 shows that ResNet-152 misclassifies the airship as a sea cucumber with an MC score of 63.02% after adding 28.6% noise under the HopSkipJump attack. When using the bit-squeezing defensive technique (with 31.7% corrective noise), the model correctly classifies the image with a confidence score of 68.13%. The median smoothing and JPEG filtering techniques allow the model to correctly classify the image with a confidence score of $\approx 100\%$.

Figure 8 depicts the performance of each defense technique in terms of decreasing the attack success rate (from the initial 100%). The results show that ① Boundary attacks are ineffective across all defenses and datasets; ② Against SimBA, no defense has shown remarkable reductions regardless of the dataset; ③ Other defenses are more efficient than bit-squeezing against HopSkipJump attack.

Adversarial Training Defenses. Figure 9 (a) shows an interesting observation, which is that as the number of layers in




Fig. 8. The effectiveness of the various defensive strategies and observes how they bring the percentage of successful attacks down from 100%. Note that JPEG defense could not be applicable to CIFAR dataset due to the smaller size of the images.

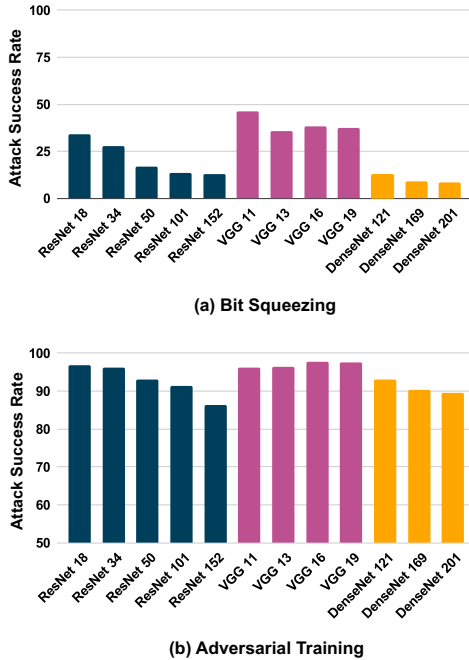


Fig. 9. Bit Squeezing defense and Adversarial Training against SimBA on ImageNet. The model becomes easier to defend as the number of layers on the same model increases

a model increases, the defense of the model becomes easier with preprocessor-type defenses. Furthermore, in addition to evaluating the models on the standard training data, we also conducted adversarial training. This process includes training the models with 500 adversarial samples. The results show that adversarial training (as shown in Figure 9 (b)) helps the model against attacks.

IV. DISCUSSION AND OPEN CHALLENGES

Key Takeaways. With our experimental study, we have reported several insights and recommendations for developing new attack proposals based on our comprehensive analysis. As indicated by the results of Experiment 1, the complexity of a model is directly proportional to the amount of additive noise required for the generation of adversarial examples. For an effective attack on DNN models on black-box settings, the attack mechanism should keep the noise visually undetectable and circumvents defensive mechanisms. Moreover, the attack must take into account not just being query efficient and having a high attack success rate, but also consider an effective way

of adding noise, controlling the amount of noise, and selecting the proper regions for adding noise.

Experiment 2 demonstrates that enhancing the model's robustness is achievable by closely examining the model's internal structure. We found that black-box attacks, *i.e.*, HopSkipJump, and boundary attacks are highly dependent on the size of the input images, while SimBA is sensitive to the number of classes in the dataset. As for defenses, preprocessor-based defenses hinder the effectiveness of boundary attacks, while adversarial training significantly contributes to enhancing the model's robustness.

Comparison. Images are the most commonly used for evaluating attacks and defenses. Initially, the MNIST and CIFAR datasets are used to assess the vulnerability of DNN to adversarial instances in the image domain [15]. In recent years, the majority of research efforts have favored ImageNet over MNIST and CIFAR. Using the ImageNet hierarchical class system, Ozbulak *et al.* [19] conducted a comprehensive investigation of the characteristics of the classes into which adversarial instances are misclassified. Specifically, it focuses on model-to-model adversarial transferability and misclassification classes using two of the most often used adversarial approaches, five distinct DNN architectures (AlexNet, SqueezeNet, VGG-16, ResNet, and DenseNet), and two vision transformer architectures (ViT-Base and ViT-Large). Gragnaniello *et al.* [11] examine the effectiveness of: ① adversarial attacks, including gradient descent, adversarial noise, and GAN-based attacks, and ② CNN-based detectors such as SPAM+SVM, Bayar2016, Cozzolino2017, and very deep nets.

Unlike prior studies, which have primarily focused on creating more effective adversary perturbation or assessing adversarial countermeasures, we address a previously under-explored problem: *analyzing black-box attacks across different datasets using a range of DL models*. Five evaluation matrices, three state-of-the-art datasets, and a wide range of DNN models are included in the initial stage of the current research. Obviously, the preliminary findings are merely proof of concept, and additional in-depth investigations are required to have a complete grasp of the issues.

V. CONCLUSION

This work presents the results of large-scale experiments on the effect of adversarial attacks and countermeasures on

distinct models utilizing three state-of-the-art datasets. The results of the experiments revealed that there is a relationship between model complexity and robustness and that having a large number of parameters does not imply that the model is more robust. Black-box attacks are susceptible to the kind of dataset used in terms of input size and the number of classes. For HopSkipJump and boundary black-box attacks, the defensive technique with minimal parameter settings is effective. Compared to the median and JPEG filters, the bit-squeezing defensive technique generated more noise, lowering the confidence score. Future directions can explore the behavior of the white-box and gray-box attacks, including recent advanced attacks such as [1], [2], [29] with other defense strategies. Moreover, studying the complexity and robustness of various architectures beyond CNNs, such as Vision Transformers, is an interesting future direction.

ACKNOWLEDGMENT

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(No. 2021R1A2C1011198), (Institute for Information & communications Technology Planning & Evaluation) (IITP) grant funded by the Korea government (MSIT) under the ICT Creative Consilience Program (IITP-2021-2020-0-01821), and AI Platform to Fully Adapt and Reflect Privacy-Policy Changes (No. 2022-0-00688).

REFERENCES

- [1] Eldor Abdulkhamidov, Mohammed Abuhamad, Firuz Juraev, Eric Chan-Tin, and Tamer AbuHmed. Advedge: Optimizing adversarial perturbations against interpretable deep learning. In *International Conference on Computational Data and Social Networks*, pages 93–105. Springer, 2021.
- [2] Eldor Abdulkhamidov, Mohammed Abuhamad, Simon S Woo, Eric Chan-Tin, and Tamer Abuhmed. Hardening interpretable deep learning systems: Investigating adversarial threats and defenses. *IEEE Transactions on Dependable and Secure Computing*, 2023.
- [3] Eldor Abdulkhamidov, Mohammed Abuhamad, Simon S Woo, Eric Chan-Tin, and Tamer Abuhmed. Unveiling vulnerabilities in interpretable deep learning systems with query-efficient black-box attacks. *arXiv preprint arXiv:2307.11906*, 2023.
- [4] Mohammed Abuhamad, Tamer Abuhmed, David Mohaisen, and Daehun Nyang. Large-scale and robust code authorship identification with deep feature learning. *ACM Transactions on Privacy and Security (TOPS)*, 24(4):1–35, 2021.
- [5] Sajid Ali, Omar Abusabha, Farman Ali, Muhammad Imran, and Tamer Abuhmed. Effective multitask deep learning for iot malware detection and identification using behavioral traffic analysis. *IEEE Transactions on Network and Service Management*, 2022.
- [6] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *6th International Conference on Learning Representations, ICLR 2018, April 30 - May 3, 2018, Conference Track Proceedings*, pages 1–12, 2018.
- [7] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. Hidden voice commands. In *25th USENIX Security Symposium*, 2016.
- [8] Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 3–14, 2017.
- [9] Jianbo Chen, Michael I Jordan, and Martin J Wainwright. Hop-skipjumpattack: A query-efficient decision-based attack. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1277–1294. San Francisco, CA, USA, 2020. IEEE.
- [10] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M Roy. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853*, pages 1–8, 2016.
- [11] Diego Gragnaniello, Francesco Marra, Giovanni Poggi, and Luisa Verdoliva. Analysis of adversarial attacks against cnn-based image forgery detectors. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 967–971. IEEE, 2018.
- [12] Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. Simple black-box adversarial attacks. In *International Conference on Machine Learning*, pages 2484–2493, Honolulu, HI, USA, 2019. PMLR.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, Las Vegas Nevada, 2016. IEEE Xplore.
- [14] Wei Jia, Zhaojun Lu, Runze Yu, Liaoyuan Li, Haichun Zhang, Zhenglin Liu, and Gang Qu. Fooling decision-based black-box automotive vision perception systems in physical world. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–12, 2024.
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical report*, pages 1–60, 2009.
- [16] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, pages 1–17, 2016.
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [18] Margarita Osadchy, Julio Hernandez-Castro, Stuart Gibson, Orr Dunkelman, and Daniel Pérez-Cabo. No bot expects the deepcaptcha! introducing immutable adversarial examples, with applications to captcha generation. *IEEE Transactions on Information Forensics and Security*, 12(11):2640–2653, 2017.
- [19] Utku Ozbulak, Maura Pintor, Arnout Van Messem, and Wesley De Neve. Evaluating adversarial attacks on imagenet: A reality check on misclassification classes. *arXiv preprint arXiv:2111.11056*, pages 1–16, 2021.
- [20] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [21] Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! *Advances in Neural Information Processing Systems*, 32, 2019.
- [22] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, pages 1–14, 2014.
- [23] Jaiteg Singh, Deepak Thakur, Tanya Gera, Babar Shah, Tamer Abuhmed, and Farman Ali. Classification and analysis of android malware images using feature fusion technique. *IEEE Access*, 9:90102–90117, 2021.
- [24] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations, ICLR 2014*, pages 1–10, 2014.
- [25] Lina Wang, Kang Yang, Wenqi Wang, Run Wang, and Aoshuang Ye. Mgaattack: Toward more query-efficient black-box attack by microbial genetic algorithm. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2229–2236, 2020.
- [26] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [27] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, February 18-21, 2018*, pages 1–15, HSan Diego, California, USA, 2018. The Internet Society.
- [28] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. Dolphinattack: Inaudible voice commands. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 103–117, 2017.
- [29] Xinyang Zhang, Ningfei Wang, Hua Shen, Shouling Ji, Xiapu Luo, and Ting Wang. Interpretable deep learning under fire. In *29th USENIX Security Symposium*, 2020.