A Systematic Evaluation of Backdoor Trigger Characteristics in Image Classification

Gorka Abad*
Radboud University
Nijmegen, Netherlands
Ikerlan research centre
Arrasate-Mondragón, Spain
abad.gorka@ru.nl

Jing Xu*
Delft University of Technology
Delft, Netherlands
j.xu-8@tudelft.nl

Stefanos Koffas*
Delft University of Technology
Delft, Netherlands
s.koffas@tudelft.nl

Behrad Tajalli*
Radboud University
Nijmegen, Netherlands
hamidreza.tajalli@ru.nl

Stjepan Picek
Radboud University
Nijmegen, Netherlands
Delft University of Technology
Delft, Netherlands
stjepan.picek@ru.nl

ABSTRACT

Deep learning achieves outstanding results in many machine learning tasks. Nevertheless, it is vulnerable to backdoor attacks that modify the training set to embed a secret functionality in the trained model. The modified training samples have a secret property, i. e., a trigger. At inference time, the secret functionality is activated when the input contains the trigger, while the model functions correctly in other cases. While there are many known backdoor attacks (and defenses), deploying a stealthy attack is still far from trivial. Successfully creating backdoor triggers heavily depends on numerous parameters. Unfortunately, research has not yet determined which parameters contribute most to the attack performance.

This paper systematically analyzes the most relevant parameters for the backdoor attacks, i.e., trigger size, position, color, and poisoning rate. Using transfer learning, which is very common in computer vision, we evaluate the attack on numerous state-of-the-art models (ResNet, VGG, AlexNet, and GoogLeNet) and datasets (MNIST, CIFAR10, and TinyImageNet). Our attacks cover the majority of backdoor settings in research, providing concrete directions for future works. Our code is publicly available [1] to facilitate the reproducibility of our results.

CCS CONCEPTS

• **Security and privacy**; • **Computing methodologies** → **Machine learning**;

KEYWORDS

Backdoor attacks, backdoor trigger, computer vision

1 INTRODUCTION

Deep neural networks (DNNs) have gained significant popularity over the past decade due to their impressive performance in various application domains, including computer vision [18], speech recognition [16], and neural translation [61]. One of the key benefits of DNNs is their ability to automatically learn and extract features from raw data, which reduces the need for manual feature engineering and makes them particularly well-suited for tasks where the data is complex or unstructured, such as image and audio processing [33]. Additionally, DNNs can efficiently process large amounts of data, achieving state-of-the-art performance on various tasks. However, DNNs also have some limitations. For example, they require a large amount of labeled training data to perform well [11], and they can be prone to overfitting if not adequately regularized [49]. They also require significant computational resources and can be challenging to interpret due to their complex decision-making processes [50].

The high computational requirements for training DNNs have led to emerging trends such as outsourced training and machine learning as a service [17]. These trends have introduced new threats for deployed models when they are provided as black boxes by third parties. In addition, malicious data samples can be easily embedded in widely-used crowdsourced datasets [43]. One approach to address the high computational requirements for training DNNs is *transfer learning*, which involves using pre-trained models as a starting point for training on a new task [58]. This can significantly reduce the amount of labeled training data and computational resources needed, as the pre-trained model has already learned many general features useful for diverse tasks. This has made transfer learning an essential tool in developing DNNs, particularly when labeled training data is limited or expensive. In addition to transfer learning, there have also been efforts to improve the interpretability of DNNs [46, 50]. This is important for various reasons, including the need to understand how a model makes decisions, the ability to identify and correct errors, and the development of trust in the model's outputs. One approach to improving interpretability is using visualization techniques, which can provide insight into the internal workings of a DNN and help identify patterns in the data the model is learning [50].

Overall, DNNs have shown impressive performance on a wide range of tasks and have the potential to continue driving significant

---

*Authors contributed equally to this research.
[1]Code will be shared after paper acceptance.

advances in artificial intelligence. Still, much work is needed to address their limitations and improve their interpretability and robustness. One particularly concerning threat is the backdoor attack, a form of compromise that results in targeted misclassifications when a specific trigger is present in the input. Backdoor attacks can be mounted through data poisoning [17], code poisoning [6], or model poisoning [20]. There has been a significant amount of research on backdoor attacks and their defenses in the literature [14, 35]. Still, these works are empirical, based on prior assumptions, and not covering a wide range of the backdoors' parameter space.

Our paper focuses on the intersection between computer vision for image classification and data poisoning, the most common setup for mounting backdoor attacks. In particular, we systematically evaluate the impact of various parameters on the performance of backdoor attacks. Our work extends previous research in this area [59] by using larger datasets with higher-dimensional images (we upsampled the images from MNIST to 64×64, from CIFAR10 to 128×128, and TinyImageNet to 224×224) and more classes. See section 5, Table 4, and Table 5 for a detailed explanation of the differences with previous works. We find that the trigger size is more influential than the poisoning rate and that the performance of backdoor attacks is affected by factors such as the model architecture and the characteristics of the trigger. Finally, we demonstrate that AlexNet is more robust against data-poisoning backdoor attacks, and we conduct experiments to explain this finding.

Our main contributions are summarized as follows:

- We extend the work described in [59] by exploring a more comprehensive range of factors that affect the backdoor performance. It allows us to provide findings that generalize well to the tested datasets and models, which represent state-of-the-art.
- Based on the extensive experimentation, we extract 1) dataset/model-specific and 2) general findings, which provide valuable insights for understanding the backdoor effect while easing the design of new attacks and defenses.
- We demonstrate that the performance of backdoor attacks is affected by various factors, including the model architecture and the characteristics of the trigger.
- We show AlexNet is more robust against data poisoning backdoor attacks and conduct experiments to explain it.

2 BACKGROUND

2.1 Deep Neural Networks (DNNs)

Deep learning algorithms are parameterized functions $\mathbb{F}_\theta$ that map an input $\mathbf{x} \in \mathbb{R}^N$ to some output $y \in \mathbb{R}^M$. θ represents the parameters of the function, which are optimized via an iterative process called training. In the image domain, $\mathbf{x}$ is an image, represented as a vector of pixel values, while y is the vector of probabilities of the image being of a class $c \in k$ from a group of classes k. For training, a dataset is needed, i.e., a set of labeled samples $\mathcal{D} = \{\mathbf{x}, y\}^n$ of size n. During training, the algorithm tries to find the optimal parameters θ' by minimizing the "distance" from the predicted labels to the ground truth ones. The distance calculation is done leveraging a loss function $\mathcal{L}$, which penalizes the algorithm depending on how far the prediction is from the actual label:

$$\theta' = \underset{\theta}{\mathrm{argmin}} \sum_{i=1}^{n} \mathcal{L}(\mathbb{F}_\theta(\{\mathbf{x}_i, y_i\})).$$

A convolutional neural network (CNN) performs convolutions in the input extracting relevant features linked to fully connected layers. The key intuition is to reduce the input space without losing information, which is easier to process in the consequent layers. This is achieved by kernels that move horizontally and vertically in the input in steps of a predefined value (stride). By doing so, the kernel extracts high-level representations as corners, shapes, or edges. Additionally, CNNs are accompanied by pooling layers that further reduce the computational complexity, extract the most relevant features, and reduce any noise captured by the kernels.

2.2 Transfer Learning

Transfer learning is a method of adapting a DNN trained for one machine learning task to a related task without retraining the entire model from scratch [58]. This can be achieved by adjusting the model's parameters by retraining only the final layers, typically the fully connected layers, in a DNN. In the case of CNNs, freezing the convolutional layers allows the model to focus on the classification task while utilizing the feature extraction capabilities that have already been optimized through pre-training. This approach can significantly reduce the computational and monetary costs associated with training a DNN from scratch, making it helpful when labeled training data is limited or expensive [17].

Transfer learning has been widely applied in a variety of fields, including computer vision [51], natural language processing [21], and speech recognition [41]. One of the key benefits of transfer learning is that it allows a DNN to utilize the knowledge it has gained from previous tasks to learn a new task more efficiently. This is particularly useful when the new task is related to a previously learned task, as it allows the model to build upon its existing knowledge rather than starting from scratch. Several factors can influence the effectiveness of transfer learning, including the similarity between the source and target tasks, the amount of labeled training data available for the target task, and the degree of feature reuse between the tasks. In general, transfer learning is most effective when the source and target tasks are similar and when there is a constrained amount of labeled training data available for the target task [68]. Transfer learning is a powerful tool for adapting DNNs to new tasks, particularly when labeled training data is limited or expensive.

2.3 Backdoor Attacks in DNNs

Backdoor attacks compromise DNNs during training and embed a secret functionality in the deployed model. This secret can be embedded through data poisoning [8, 17], code poisoning [6], or direct modification of the model's weights [20]. In this work, we follow data poisoning by injecting poisoned samples into the training set. A poisoned sample contains a trigger, and its label is usually altered to the target label, which is the output of the model when the backdoor is activated. In the image domain, the trigger is usually a pixel pattern of a given color, e.g., white or black, placed anywhere over the image, creating a set of poisoned samples $\hat{\mathbf{x}} \in D_{poison}$. The

percentage of poisoned samples in the training set is controlled by $\epsilon = \frac{m}{n}$ where m is the number of poisoned samples, n is the number of the original training set, and $m \ll n$. A small ϵ makes the backdoor harder to embed but keeps it stealthier, as the small number of poisoned samples will not affect the original task much. A large ϵ leads to a stronger backdoor, but it could affect the original task substantially, making it somewhat unrealistic. During training with poisoned samples, the backdoor effect is included following:

$$\theta' = \underset{\theta}{\arg\min} \sum_{i=1}^{n} \mathcal{L}(\mathbb{F}_\theta(\{\mathbf{x}_i, y_i\})) + \sum_{j=1}^{m} \mathcal{L}(\mathbb{F}_\theta(\{\hat{\mathbf{x}}_j, \hat{y}_j\})).$$

After training, the backdoor is embedded in the DNN. The DNN functions normally on clean inputs, but the secret functionality (backdoor) is activated in the presence of the trigger.

2.4 On Backdoor Interpretability

Interpretability techniques have been widely used for explaining the behavior of ML models. The interpretability of DNNs refers to the ability to understand the decision taken by the network, which can be obtained by different methods such as feature visualization. Typically there is a trade-off between accuracy, simplicity, and explainability. For instance, shallow models such as linear regression or decision trees are highly interpretable [40, 48]. By using DL models, we sacrifice the interpretability to achieve better performance, which often increases the complexity of the model by adding more layers. For instance, residual models like ResNet achieve state-of-the-art performance in several tasks by having over 200 layers [18].

Recent work developed a technique named class activation mapping (CAM) for CNNs, which identifies the regions of an image that are more linked to the model's prediction [67]. CAM modifies the architecture of the target model by changing the convolutional layers for fully connected layers, which are much more interpretable but incur a severe degradation in accuracy. Consequent work from Selvaraju et al. introduced a generalization of CAM called gradient-weighted CAM (Grad-CAM) [50]. Instead of modifying the model's architecture, it uses the gradient of a given class to produce a localization of the important regions of the image. Precisely, Grad-CAM computes the target class's gradients concerning the feature map activation of a convolutional layer.

In addition to visualization techniques, post-hoc interpretation methods have also been developed to explain individual predictions made by a DNN. These methods are applied after the model has been trained and do not require changes to the model's architecture or training process. One example of a post-hoc interpretation method is LIME (Local Interpretable Model-Agnostic Explanations) [46]. LIME generates explanations by fitting a simple, interpretable model to the predictions made by a DNN in the vicinity of a particular input, allowing the model's behavior to be understood locally.

Other post-hoc interpretation methods include SHAP [39] and DeepLIFT [52]. SHAP uses Shapley values, a concept from game theory, to attribute the prediction made by a DNN to the individual features of the input. DeepLIFT computes the contribution of each feature to the final prediction by comparing the model's output with a reference score, which the user can choose.

Overall, various approaches are available for explaining ML models' behavior, including visualization techniques and post-hoc interpretation methods. The choice of method will depend on the specific requirements and constraints of the task. In this work, we will use Grad-CAM to understand the decisions of the poisoned models and compare their behavior with their clean counterparts, which is a suitable method for understanding the backdoor behavior [48].

2.5 Motivation

In recent years, DL has become an extremely popular and rapidly evolving domain as a form to solve various real-world problems. Due to the need for adaptation to other tasks, DNNs have become more complex, often viewed as "black-boxes". Indeed, Gilpin et al. [15] established a direct relationship between the models' complexity and their (lack of) explainability. Furthermore, efforts to create more complex and explainable models have been ongoing within the research community [4]. At the same time, DNNs have also gained rising attention in the security community due to their vast applicability and impact of DNNs. The ability to understand and explain the inner workings of these models becomes particularly important in the context of DL attacks, as a lack of explainability can hinder our understanding of the root cause of security problems.

One type of DL attack that has garnered significant attention is the backdoor attack. Indeed, it has been recently subject to a deep investigation in a wide range of domains [35]. In image recognition, the proposed attacks are heterogeneous in the trigger generation, backdoor injection, or threat model. Thus, comparing these attacks is far from trivial, even impossible in some cases. For instance, the models, datasets, experimental setups, and attack parameters are only a few to consider for comparing the performance of different attacks. Furthermore, even if the attacks are comparable, understanding the influence of the attack's parameters on the backdoor performance could still be difficult.

In this paper, we aim to address these issues by systematically investigating the impact of common parameters on the effectiveness of backdoor attacks in clean and backdoor performance. We analyze the core group of backdoor attacks in image classification, where the rest of the attacks build upon. For that, we investigated the backdoor attacks that follow the BadNets approach [17] in the literature. The papers that fulfilled our criteria are shown in Table 1. By analyzing those, we found an inconsistency in the parameter selection and the understanding of these parameters' effect on the backdoor performance. Thus, we propose systematically analyzing the attack proposals based on the same parameters and investigating the influence of these parameters on the main and backdoor task's performance. This allows us to efficiently and systematically compare a new attack—allowing fair and traceable comparisons. Our final goal is to provide a comprehensive and systematic analysis of the impact of parameters on backdoor attack performance. By doing so, we hope to contribute to a better understanding of these types of attacks and provide a valuable investigation for comparing and evaluating future research in this area.

For our investigation, we provide a realistic attack configuration, and we have designed our experimental setup to be as simple as possible while still being extendable to future backdoor attacks. To this end, we have surveyed the state-of-the-art to identify a

common set of experimental settings that can be used to compare and evaluate future attacks. In line with previous research [35], we have chosen to focus on image recognition, using the MNIST, CIFAR10, and TinyImageNet datasets and models AlexNet, ResNet, VGG, and GoogLeNet. These datasets and models are representative samples of the ones used in the state-of-the-art.

Additionally, it is important to note that the choice of parameters can significantly impact the performance of a backdoor attack. For example, the trigger size, poisoning rate, and type of trigger can affect the attack's success rate. Similarly, the choice of dataset and model architecture can significantly affect the attack's effectiveness. By considering various values of parameters in our experiments, we explored a range of possibilities and identified differences in attack performance that may be related to each parameter. This is important in understanding the underlying mechanisms of backdoor attacks and developing more effective countermeasures.

Table 1: Comparison of the attack setting for different state-of-the-art backdoor attacks that use patch triggers.

Paper	ϵ	Trigger Size	Trigger Location	Trigger Color
BadNets [17]	0.1 0.3 0.5	Single pixel Four pixels	Bottom-right Center	White Yellow Patch
Salem et al. [48]	0.1 0.2 0.3 0.4 0.5	0.1% 0.5% 1.5% 2% 3.2% 4%	Corners Top center Bottom center	Random Dynamic
Liu et al. [38]	-	4% 7% 10%	Bottom-right	Dynamic
Kwon et al. [30]	0.1 0.25 0.5	25%	Corners	White
Tan et al. [56]	0.05	8%	Bottom-right	White
Feng et al. [13]	0.005 0.01 0.02	-	-	Dynamic
Zhang et al. [64]	0.1 0.2 0.3 0.4 0.5	-	-	Dynamic
Li et al. [36]	0.02 0.04 0.06 0.08 0.10	9%	Bottom-left	White

3 THREAT MODEL

We consider a *gray-box* threat model as the attacker can freely modify a small portion of the training dataset and has no knowledge about the training algorithms or the models used by the victims. We also assume a *dirty-label* backdoor attack meaning that the attacker can alter both the training samples and their labels. Even though this threat model is weaker than its counterpart (clean-label attack [60]), it is the most popular among the existing works [8, 10, 17, 38, 40, 48]. Additionally, we target only transfer learning as it has become a very common practice as training from scratch can be very expensive and the weights of state-of-the-art models like

VGG and ResNet trained on ImageNet are publicly available [17]. This threat model is realistic as large datasets like ImageNet [11] are crowdsourced from untrusted sources, and malicious data can evade human inspection [43].

We consider the following metrics:

(1) **Attack Success Rate** (ASR): measures the backdoor performance of the model on a fully poisoned dataset D_{poison}, i.e., $\epsilon = 1$. It can be computed by $ASR = \frac{\sum_{i=1}^{N} \mathbb{I}(F_{\hat{\theta}}(\hat{x}_i) = y_t)}{N}$ where $F_{\hat{\theta}}$ is the poisoned model, $\hat{x}_i$ is a poisoned input, $\hat{x}_i \in D_{poison}$, y_t is the target class, and $\mathbb{I}(x)$ is a function that returns 1 if x is true and 0 otherwise.

(2) **Clean Accuracy Drop** (CAD): measures the effect of the backdoor attack on the original task. It is calculated by comparing the performance of the poisoned and clean models on a clean holdout validation set D_{valid}, i.e., $\epsilon = 0$. The accuracy drop should be small to keep the attack stealthy.

4 EXPERIMENTS

This paper systematically evaluates the backdoor attacks of 3 image datasets, 4 DNN models, 3 trigger sizes, 5 trigger positions, 3 trigger colors, and 4 poisoning rates. We also train a clean model for each dataset and for each architecture (3×4). Each of these experiments was repeated 5 times. Thus, we train 10,800 backdoored models and 60 clean models in total.

4.1 Experimental Matrix

Datasets. We evaluated our approach using MNIST, CIFAR10, and TinyImageNet.

- MNIST [32] is a dataset of 70,000 grayscale images of hand-written digits, each 28×28 pixels in size and belonging to one of 10 different classes. We converted the images to RGB format for our evaluation and resized them to 64×64 pixels [2].
- CIFAR10 [28] is a dataset of 60,000 RGB images, each 32×32 pixels in size and belonging to one of 10 different classes, with 6,000 images per class. Similar to MNIST, we resized the images to 128×128 pixels for compatibility.
- TinyImageNet [31] is a dataset of 120,000 RGB images belonging to 200 different classes, each 64×64 pixels in size. We also resized these images to 224×224 pixels.

Model Architectures. In our experiments, we selected four standard benchmark DNNs for evaluation: AlexNet, GoogLeNet, VGG-19_BN, and ResNet-152. To utilize transfer learning and extract features from these models, we froze the parameters for all layers (except for the last fully connected layer and the batch-normalization layers for ResNet, VGG, and GoogLeNet). This allowed us to leverage the pre-trained models while focusing on the task of interest. AlexNet, however, resists backdoor injection when being operated by transfer learning (i.e., we reach low ASR when freezing all layers except the last one). Thus for a more suitable analysis, our transfer learning setup for AlexNet is to freeze the layers up to its classifier's module (more in section 4.3.2).

These DNNs were selected based on their demonstrated performance on various tasks and widespread use as benchmarks in the

---

[2]Data resizing is used to adapt inputs to the chosen networks, which require a minimum input size. Additionally, we experiment with different input sizes to better generalize the results.

field. AlexNet, a CNN introduced by Krizhevsky et al. [29], was the first successful CNN to demonstrate superior performance on the ImageNet dataset [12]. In PyTorch implementation for AlexNet (which we use for our study in this work), it consists of two main modules: the features module, which itself consists of convolutions and pooling layers, and the classifier module, which is composed of fully connected layers for the final classification task. GoogLeNet, introduced by Szegedy et al. [55], is a variant of the Inception architecture that won the 2014 ImageNet Large Scale Visual Recognition Challenge. VGG-19_BN is a variant of the VGG network [53] that incorporates batch normalization [23] and has achieved strong performance on a range of tasks. Finally, ResNet-152 is a residual network [19] with 152 layers that also achieved state-of-the-art performance on several tasks. By using these well-established DNNs, we ensure the reliability and generalizability of our results.

Trigger Size. We focus on the square trigger pattern, a commonly used trigger in backdoor attacks on image classification tasks [17, 59]. This trigger consists of a square patch injected into the training images and used to manipulate the model's behavior. In [59], the square trigger proved the most effective, so we did not consider blending overlay triggers. To evaluate the effectiveness of the attack under different conditions, we varied the width and height of the trigger as a percentage of the width and height of the sampled training image, using values of 4%, 6%, and 8%, which allowed us to assess the trigger size's impact on the attack's performance. These trigger sizes cover most of the trigger sizes considered in the literature while being realistic.

Trigger Position. We inject the square trigger into five locations in the poisoned images: the top-left, top-right, middle, bottom-left, and bottom-right positions. This allowed us to evaluate the impact of the trigger position on the performance of the attack and to identify any trends or patterns that may be present. Figure 1 illustrates an example image from the CIFAR10 dataset with triggers embedded at various positions. By studying the attack under these different conditions, we could gain a deeper understanding of the factors that influence the success of a backdoor attack and develop more effective defense strategies.

(a)   (b)   (c)   (d)   (e)
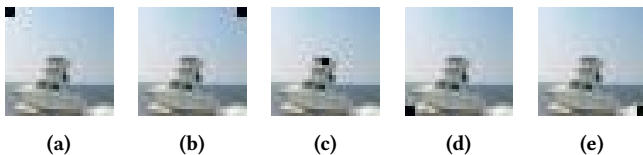
Figure 1: Trigger patterns with different trigger positions (top-left, top-right, middle, bottom-left, bottom-right) applied to an image from the CIFAR10 dataset.

Trigger Color. The color of the trigger pattern is another critical factor to consider in the design of a backdoor attack. In our experiments, we evaluated the performance of the attack using three different trigger colors: black, white, and green. The green trigger was randomly picked to avoid biases that extreme values like black (0, 0, 0) or white (255, 255, 255) may create. To this end, we used python's pseudorandom generator to retrieve three RGB values (one for each channel). The values are (102, 179, 92). The

MNIST dataset only has one channel, so we run experiments of green color on CIFAR10 and TinyImageNet datasets. By comparing the results obtained with these different trigger colors, we could understand how the color of the trigger affects the effectiveness of the attack and identify any trends or patterns that may be present. This information is valuable for understanding these attacks' behavior and developing more effective defense strategies.

Poisoning Rate. One of the key factors that impact the backdoor's effectiveness is the poisoning rate, which refers to the percentage of training images injected with the backdoor trigger. In our experiments, we replaced clean images with their poisoned counterparts to avoid altering the number of training samples in the dataset. We also varied the poisoning rate to study its impact on attack performance. This allowed us to study the effect of different poisoning rates on the attacks' success. Additionally, we chose small poisoning rates because the backdoor should affect the original task as little as possible, and given the amount of data that modern deep learning systems need, large poisoning rates can be unrealistic [8]. Thus, we defined four values: 0.5%, 1%, 1.5%, and 2%.

4.2 Experimental Setup

Dataset Split. We followed the standard splitting of samples for training and testing for each dataset. Specifically, for MNIST, we sample 60,000 images as a training set and the rest as a test set. For CIFAR10, the whole dataset is divided into 50,000 training images and 10,000 test images. For TinyImageNet, we sample 500 and 50 images for each class as the training set and test set, respectively (i.e., 100,000 training samples and 10,000 test samples). We randomly shuffle the datasets in all the cases.

Training Procedure. We chose the Adam algorithm and cross-entropy loss as an optimizer and the criterion in our experiments. However, in one setting (TinyImageNet + VGG), the Adam optimizer yielded poor performance (around 37%), and we had to use SGD, which resulted in an accuracy of around 72%. Furthermore, we experimentally set the learning rate to 0.001 and the number of epochs to 20, where we achieve training convergence and good generalization in the test set. Each dataset's batch size is different to fit into the GPU's memory. For the small datasets (MNIST and CIFAR10), the batch size is 128, and for TinyImageNet is 32. Each experiment was repeated five times to reduce the effects of randomness caused by stochastic gradient descent and initialization. The experiments were run using PyTorch v1.12 on a cluster of machines running CentOS Linux with NVIDIA GPUs (Tesla P100, GeForce GTX 1080 Ti, GeForce RTX 2080 Ti, and Tesla v100).

4.3 Results and Analysis

4.3.1 Clean Accuracy Drop. The backdoor should remain stealthy in the deployed model to avoid raising any suspicions. Thus, the model's performance on the original task should not be affected by the backdoor insertion. To ensure this is true in our experiments, we calculate the arithmetic mean of the accuracy of all the clean models we trained ($\epsilon = 0$) and compared it to the mean of the accuracy of the poisoned models. We show the results in Table 2, where we use bold for the value with the largest difference from the clean model. We see that the difference introduced by the backdoor is really small. In almost all cases, the accuracy is decreased but

less than 1%. We see only in one case (TinyImagenet + ResNet) a performance drop of around 2%. ResNet is the best-performing model with TinyImageNet (clean accuracy 83.96%), and even a small change in the training data affects the model's generalization. **Additionally, the performance drop is positively correlated with ϵ and as ϵ gets larger, the drop is increased as well**.

From Table 2, we can also see that our models perform well for the datasets tested. However, AlexNet is not very accurate with TinyImageNet and has an accuracy of 21.73%($\pm$0.6067). As we discussed in Section 4.3.2, the backdoor did not work in this case if we freeze all the layers up to the fully connected layers. Thus, we had to unfreeze a few layers from the feature extractor for the backdoor to be more effective. This resulted in lower performance for the original task as we altered the weights of the feature extractor. If we keep these layers frozen, the model's accuracy for clean inputs is around 46%. **Thus, we conclude that the backdoor attack is more effective if the model is trained from scratch or has a large number of trainable layers**.

4.3.2 Effect of Model Architecture.
In most cases, VGG is the least robust to backdoor attacks, especially on the CIFAR10 dataset, while AlexNet is the most robust to poisoning on all three datasets. For example, in Figure 4, the attack success rate of VGG is always higher than other models. Specifically, the attack performance of VGG with a small poisoning rate, i.e., 0.005, is higher than the other models. VGG has the most neurons among the four models, leading to a larger capacity to learn the backdoor functionality. With the increasing poisoning rate, the ASR of the ResNet and GoogLeNet increases and is similar to VGG's. **We believe that models with larger capacities are more vulnerable to backdoors as they can encode more patterns in their weights even from a very small part of the dataset**.

Additionally, if we freeze the feature extractor layers in AlexNet like the other three models, the ASR is nearly 0%. Because of this, we decided to unfreeze AlexNet parameters layer by layer (from 14 to 0) to see from which layer it starts to react positively on the injected backdoor. Appendix Figure 18 and Figure 2 show the results of our experiments on MNIST and CIFAR10. When we unfreeze the classifier module completely, the network starts to learn the backdoor. This can be observed on the plots when the network is unfrozen up to the 7th parameter. Thereafter, there is a surge in most of the plots from this point, showing that the backdoor has started to work. After this experiment, we decided to do the experiments with AlexNet by retraining the whole classifier module and freezing the feature Module. Nonetheless, the results show, except for trigger positions in the middle for CIFAR10 (Figure 10 and Figure 5), in all other experiments, the backdoor attack fails to reach high ASR on AlexNet. Interestingly, the classifier modules in AlexNet and VGG are very similar (both having 3 fully connected layers with 4,096 neurons in each layer). The main difference between these two is that in AlexNet, the two dropout layers precede the linear layers, but in VGG, they succeed (this means that in AlexNet, the first dropout will affect the last convolution layer in the features module, while in VGG both dropouts will affect the former fully connected neurons before them). VGG is a deeper and more complicated network than AlexNet, making it more vulnerable to backdoor triggers. However, the reason for AlexNet's robustness against backdoors is not merely

its smaller capacity because unfreezing the classifier part improves backdoor learning. From [47], we know that dropouts can affect the learning process of a network and cause a network to learn a deliberate backdoor. We correspondingly assume that the role of dropout layers and their inactivity during test time may affect the backdoor success. Nonetheless, we are not 100% sure about this, and more experimental studies are needed to be done in the future to uncover the primary reason.

AlexNet has demonstrated to be a very robust network on simple square shape backdoor patterns compared to the other three benchmark networks. It seems that the most important parameter which could affect the ASR on AlexNet is the trigger size. Figure 19 displays the output of the AlexNet feature module on the same poisoned image with different trigger sizes (the dissimilarity of activations based on trigger size can be observed by comparing two feature map differences on the right).

4.3.3 Effect of the Trigger Size.
In our experiments, we see that by only changing the trigger size, we can create very effective backdoors. For example, in Figure 8, we see that the ASR for AlexNet and MNIST is very low (around 10%) in all cases when the trigger size is 4% or 6%. However, in the same setting, changing the trigger size to 8% could lead to an ASR as high as 80%. Similar behavior is shown in Figure 5 and in Figure 10 for CIFAR10 and all the models.

For the CIFAR10 dataset, we see that the trigger size is the most influential for AlexNet. The feature extractor of the model is unfrozen in AlexNet, so the model can learn easier to spot larger triggers. However, there are multiple cases where changing the trigger size leads to high ASR for the other models. For example, in Figure 4, we see that the ASR for ResNet increases from 40% to more than 90% when $\epsilon = 0.01$ and the size is increased from 6% to 8%. Similarly, in Figure 6 and in Figure 7, the ASR for GoogLeNet is increased from around 10% to more than 80%.

When we insert the trigger in the middle, for the MNIST and CIFAR10 datasets, the ASR of all models (except AlexNet) is around 10% with a trigger size less than 8%. However, it increases significantly with a trigger size of 8%. In TinyImageNet, the ASR is low when the trigger is not placed in the middle of the image. However, even in these cases, increasing the size may increase the ASR (Figure 13 and Figure 16). **Thus, we conclude that the trigger size can significantly affect the ASR**.

4.3.4 Effect of the Trigger Position.
For the MNIST and CIFAR10 datasets, with a trigger size is less than 8%, there is no noticeable difference in the ASR when it is injected in the corners. However, there is a decrease in the ASR for all models when the trigger is injected in the middle. With trigger size increasing to 8%, the trigger position has an unnoticeable impact on the attack performance. On the contrary, for the TinyImageNet dataset, all models are robust to the backdoor attack when the trigger is not injected in the middle. With the trigger in the middle, there is a significant rise in the ASR for all models. This could explain that, in general, images in TinyImageNet are not centered, in contrast with those in MNIST and CIFAR10. Therefore, for TinyImageNet, triggers placed in the middle can achieve high ASR without a noticeable degradation on the main task. However, the model cannot recognize triggers placed in the corners or are small, i.e., less than 8%.

Table 2: Clean accuracy comparison between clean and poisoned models. We show in bold the settings that have the largest difference with the clean model's ($\epsilon = 0$) performance.

Dataset	Model	ϵ (%)				
		0	0.5	1	1.5	2
MNIST	AlexNet	98.50±0.1915	98.41±0.1883	98.37±0.1876	**98.30±0.2468**	98.31±0.2136
	GoogLeNet	98.75±0.1191	98.67±0.1363	98.64±0.1654	98.62±0.1817	**98.58±0.2173**
	ResNet	98.83±0.1846	98.64±0.3198	98.50±0.4217	98.33±0.4882	**98.19±0.6144**
	VGG	99.09±1.1671	99.22±1.1597	**99.34±0.1769**	99.31±0.4784	99.30±0.2782
CIFAR10	AlexNet	85.17±0.3677	84.89±0.4034	84.68±0.4217	84.52±0.4050	**84.40±0.4397**
	GoogLeNet	92.54±0.1464	92.38±0.2023	92.33±0.2190	92.22±0.2172	**92.18±0.1961**
	ResNet	96.88±0.1449	96.68±0.1983	96.61±0.2675	96.58±0.3037	**96.56±0.3779**
	VGG	93.02±0.4733	92.87±0.5260	92.90±0.4743	92.85±0.4712	**92.77±0.5308**
TinyImageNet	AlexNet	21.73±0.6067	21.60±0.6881	21.49±0.7208	21.17±0.7209	**20.89±0.7470**
	GoogLeNet	70.07±0.1688	70.02±0.2322	69.96±0.2569	69.86±0.2513	**69.79±0.2574**
	ResNet	83.96±0.1927	82.90±0.5713	82.45±0.7705	82.09±0.9795	**81.90±0.9917**
	VGG	72.66±0.2265	72.60±0.2211	72.51±0.2564	72.41±0.2315	**72.33±0.2301**

All these show that in our experiments, no position universally leads to a more successful backdoor attack. **The most effective position is different for every dataset and depends on the dataset's properties and the way the models learn**.

4.3.5 Effect of the Trigger Color. For MNIST, the ASR is low for black triggers placed in the corners. The effect is expected as the training images in MNIST contain many black pixels by default, and the model cannot identify our black trigger as a feature. However, white triggers placed in the corners are effective due to their contrast with the black background. For both colors (black and white), the trigger should be large (8%) to start having an effect on the ASR when placed in the middle (Figure 5 and Figure 10) as it overlaps the sample's main information, i.e., the number.

In TinyImageNet, when the trigger is placed in the corners, in most cases, the ASR is around 0%. However, when the trigger has a size of 8% and is green (Figure 13 to Figure 17), the ASR can be increased up to 40% (Figure 16). Additionally, when the trigger is inserted in the middle, the backdoor works in all cases but is more effective when green (Figure 15).

In CIFAR10, we see that in some cases for small triggers (< 8%), GoogLeNet is more effective with white triggers. For example, comparing Figure 3 and Figure 8 we see that for trigger size 4% and ϵ = 0.5%, the ASR increases from 10% to almost 90%. The same is true for small (< 8%) triggers in top-right (Figure 7 vs. Figure 12) and in top-left (Figure 6 vs. Figure 8). Additionally, for black and white triggers smaller than 8% placed in the middle (Figure 5 and Figure 10), the ASR is low for all models except for AlexNet. In this case, the backdoor works only with a green-colored trigger.

From all these observations, we conclude that the trigger color can play an important role in the backdoor's effectiveness, but it depends on many factors like the dataset or the model, making its optimization challenging for an attacker.

4.3.6 Effect of the Poisoning Rate. Generally, with the increasing poisoning rate, all models' ASR increases. This is reasonable because, with more poisoned data, a backdoor attack can perform better. However, the attacker cannot increase ϵ indefinitely as the model's CAD is reduced when the poisoning rate grows.
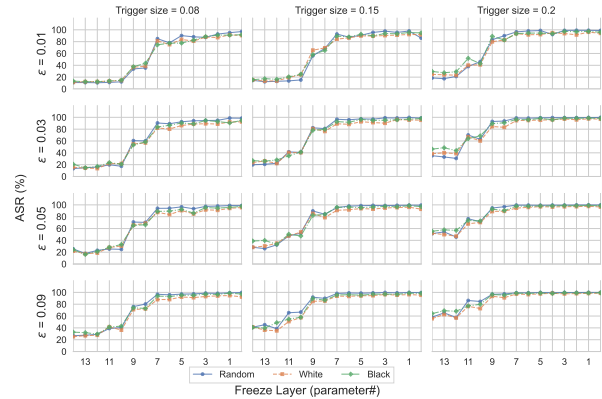
Figure 2: AlexNet on CIFAR10: FreezeLayer effect vs. size and rate, trigger at bottom-right
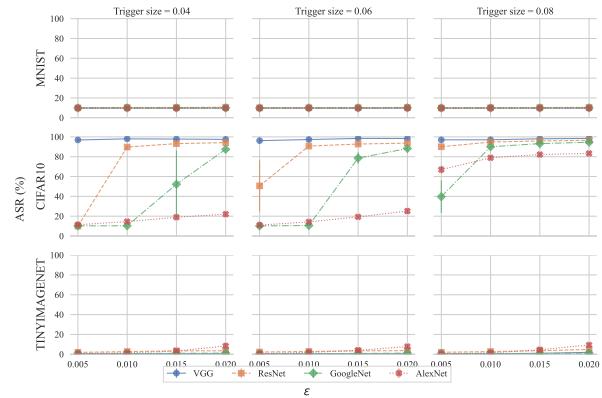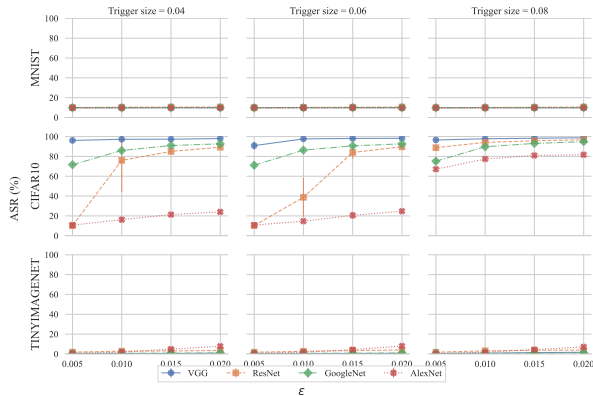

Figure 3: Rate vs. size, black color, trigger at bottom-left

## 4.4 On the Interpretability of Backdoors

Convolutional layers capture the spatial information, so the last convolutional layer is expected to achieve the best understanding of high-level semantics and detailed spatial information. Thus, the

Table 3: Summary of the results. ■ means that the ASR for at least one model is higher than 80%, ◪ that the ASR is between 60% and 80%, and □ that the ASR of every model is below 60%.

| Trigger color | White | | | | | | Black | | | | | | Green | | | |
| Trigger position \\ Trigger size | 0.04 | | | 0.08 | | | 0.04 | | | 0.08 | | | 0.04 | | 0.08 | |
	MNIST	CIFAR10	TinyImageNet	MNIST	CIFAR10	TinyImageNet	MNIST	CIFAR10	TinyImageNet	MNIST	CIFAR10	TinyImageNet	CIFAR10	TinyImageNet	CIFAR10	TinyImageNet
Top-left	■	■	□	■	■	□	□	■	□	■	■	□	■	□	■	□
Middle	□	◪	□	◪	■	■	□	□	□	◪	■	■	■	□	■	□
Bottom-right	■	■	□	■	■	□	□	■	□	□	■	□	■	□	■	□

Figure 4: Rate vs. size, black color, trigger at bottom-right
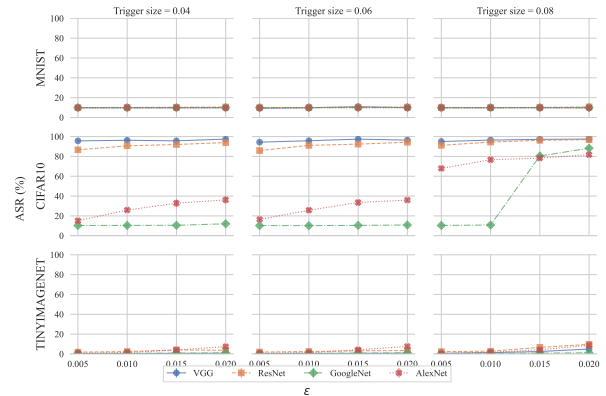
Figure 6: Rate vs. size, black color, trigger at top-left
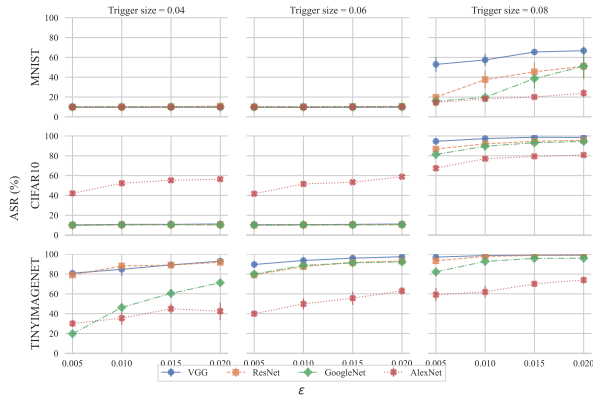
Figure 5: Rate vs. size, black color, trigger at middle
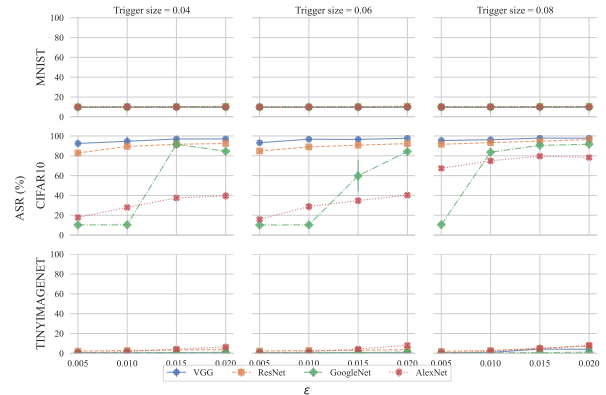
Figure 7: Rate vs. size, black color, trigger at top-right

neurons of convolutional layers look for the class-specific semantics, e.g., capturing image parts relevant to the label "dog". Grad-CAM uses this information for obtaining an attention map given an image and a target class. Intuitively, one can imagine Grad-CAM attention maps as the critical parts for a model to classify an image for the target label. Grad-CAM has also been widely applied in the image backdoor domain to explain the behavior of the backdoor triggers [40, 48]. More precisely, we also leverage Grad-CAM to explain the importance of the trigger location and color. We use CI-FAR10 as a test dataset to compare the attention of the backdoored models and the clean models for both clean and backdoored samples. We selected CIFAR10 because it is a perfect candidate since it contains large (upscaled) color images, which is also representative

of TinyImageNet and richer in features than MNIST. We select the setting from a successful backdoor attack to ensure that the trigger is getting injected. We experimented with a black trigger of size 8% of the input image placed in the top-left corner. We set the ϵ value to 0.02 and train the models for 20 epochs. Simultaneously, we train a clean version of the same model and compute the attention maps for both clean and poisoned models. These maps show the image's most influential part (in red) for the model's output. Depending on the model used, we observe different behaviors. It is important to note that we use clean and target labels, i.e., the ground truth label and the backdoor label, to help understand the label's effect on the model's prediction. Intuitively, we expect a well-trained clean model to resist image perturbations (to some extent) as input

Figure 8: Rate vs. size, white color, trigger at bottom-left

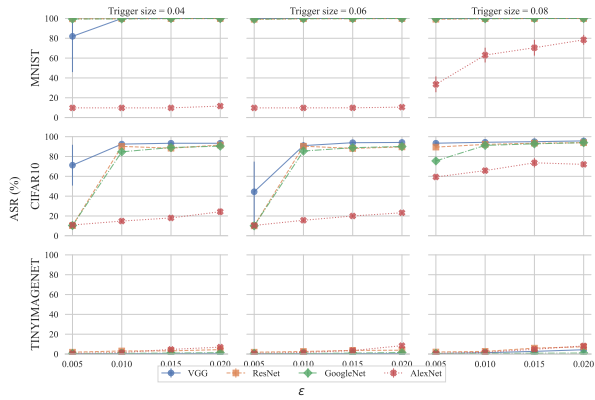**Figure 9: Rate vs. size, white color, trigger at bottom-right**

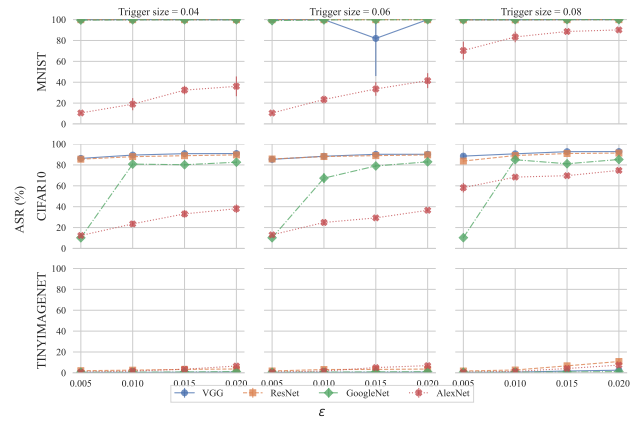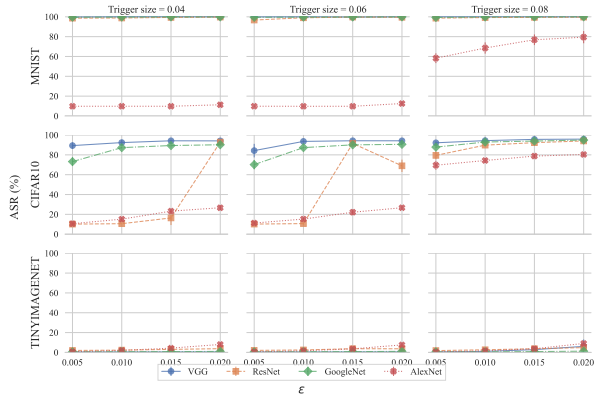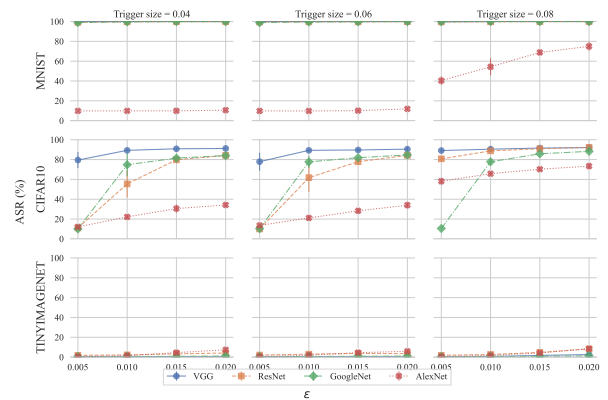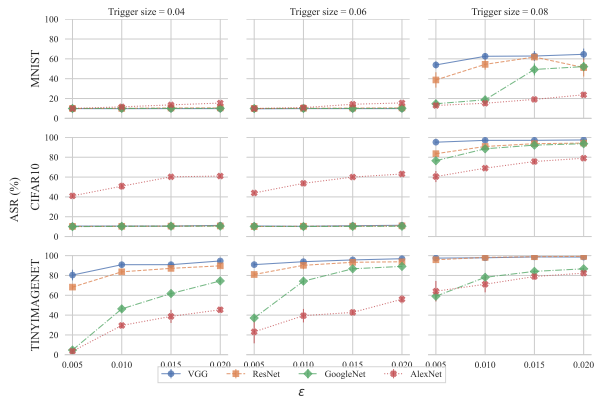**Figure 10: Rate vs. size, white color, trigger at middle**

Figure 11: Rate vs. size, white color, trigger at top-left

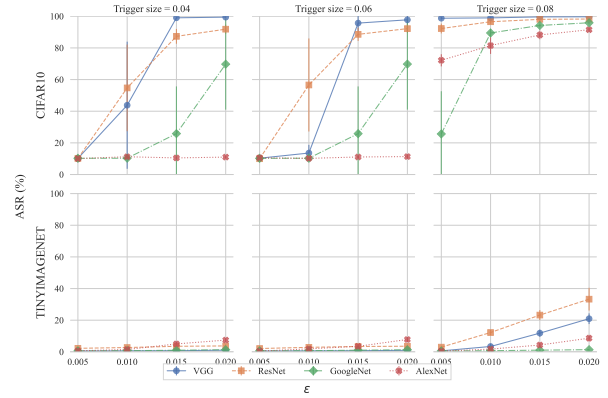**Figure 12: Rate vs. size, white color, trigger at top-right**

Figure 13: Rate vs. size, green color, trigger at bottom-left

triggers. Therefore, we expect the clean model's attention maps to look similar. However, on a backdoor model trained with clean and backdoor data, we expect to obtain a similar attention map (as the clean model's) for the clean images. Nevertheless, backdoor images should bring the model's attention toward the trigger.

In GoogLeNet, the clean model (see Figure 20 in Appendix A) focuses on the center and center-right locations for clean and target labels. This effect also remains visible in the backdoor model (see Figure 21 in Appendix A), caused by the backdoor "idea" where
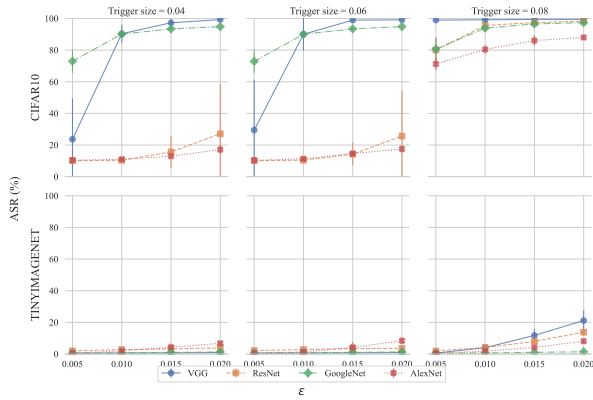
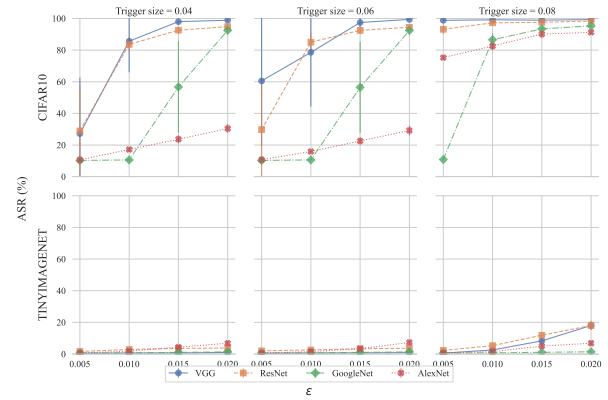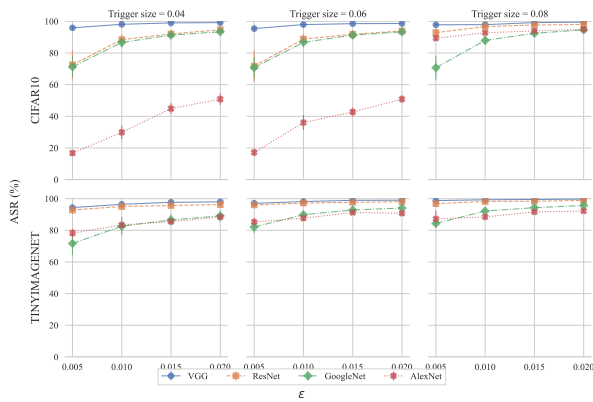Figure 14: Rate vs. size, green color, trigger at bottom-right

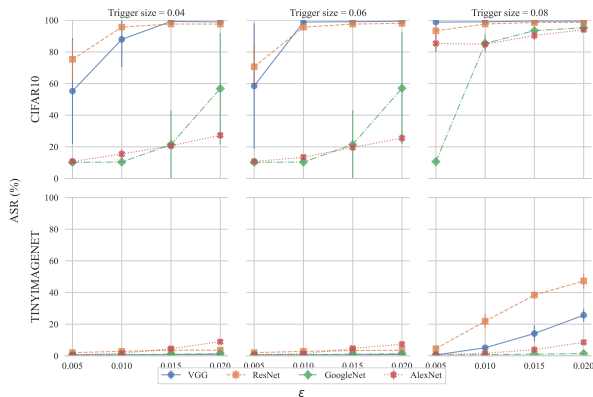Figure 15: Rate vs. size, green color, trigger at middle

Figure 16: Rate vs. size, green color, trigger at top-left

the attention on clean images does not vary. However, the backdoor model's attention drifts toward the trigger under its presence. In the clean model, the trigger is unnoticed.

Figure 17: Rate vs. size, green color, trigger at top-right

In ResNet and VGG, we observe a similar, yet more evident behavior as in GoogLeNet. The clean model (see Figure 22 and Figure 24 in Appendix A) robustly resists the trigger presence without modifying the attention map and maintaining the same as the clean input. The backdoor model also focuses on the exact locations of the images, as the clean model does. On poisoned inputs, the backdoor model easily recognizes the presence of the trigger, directing attention toward it, see Figure 23 and Figure 25 in Appendix A.

AlexNet's attention maps are biased by the poor performance on the backdoor task as in Figure 6. The heatmaps could intuitively help explain it. AlexNet's predictions are based on observing all the areas from the image rather than focusing on a specific area, as done by the abovementioned models. Still, the clean model is robust against perturbations on the input, i.e., the attention map does not vary much, see Figure 26 in Appendix A. Similarly, the backdoor model has a slightly different attention map on clean images. However, the model does not focus on the trigger but on the whole input space on backdoor images, see Figure 27 in Appendix A.

4.5 Discussion

We discuss several aspects of the backdoor attacks in image classification based on our experimental findings. First, we saw that the backdoor attack is easier when training from scratch. Thus, in future works, authors claiming that their trigger generation technique is stronger than the state-of-the-art should also run experiments in a transfer learning setup.

> *Finding* 1. The backdoor attack is easier when training from scratch.

Additionally, we should always use small poisoning rates as the clean accuracy drop increases when the poisoning rate is increased. Our experiments indicate that the drop is more severe for stronger models and larger datasets. Thus, there is no guarantee of a small clean accuracy drop in large datasets if we see no clean accuracy drop in small datasets.

> *Finding* 2. The clean accuracy drop increases as the poisoning rate increases. Additionally, we conjecture that the drop can be more severe for large datasets and strong models.

We also saw that models with a large capacity and a large number of weights are more vulnerable to backdoor attacks. These models can overfit to small subsets of their datasets and learn complex patterns even from only a handful of training samples.

> *Finding* 3. Large models with big capacities are more vulnerable to backdoor attacks.

From our experiments, we saw that no position, color, or combination of them results in the most effective backdoor across all settings. The best trigger color and position for every setup depends on the dataset, and the model used.

> *Finding* 4. No position or color results in the most effective backdoor universally.

Another observation from our experiments is that the ASR can vary for different trigger positions. Even though CNNs should not be affected by the feature (trigger) position, it seems that in some cases, they exploit the feature's absolute spatial location and learn the trigger easier. This was also shown in [24] but not in the context of backdoor attacks.

> *Finding* 5. The backdoor's performance varies for different trigger positions indicating that in some cases, the CNNs exploit the absolute spatial location of their features.

As shown in Figure 5 for TinyImageNet, variations in the poisoning rate show improvement when the trigger size is 4%. However, when the trigger is large, the ϵ does not affect much the backdoor performance. A similar effect is visible in Figure 10, where with trigger size 0.04, variations in ϵ can drastically increase the backdoor performance. However, the poisoning rate is nearly irrelevant when the trigger size is large.

> *Finding* 6. The trigger size has a more significant contribution to the ASR than the poisoning rate.

5 RELATED WORK

Backdoor attacks have been widely investigated in different domains in recent years. BadNets [17] was the first paper to address backdoor attacks in computer vision for image classification. Since then, backdoors have also been applied to different domains such as audio [25, 27, 63], graph neural networks [62, 65], spiking neural networks [1], natural language processing [5, 6, 9], or collaborative learning [2, 3, 7]. Specific to the image domain, different approaches have arisen: multi-trigger [30], dynamic [26, 40, 48], or invisible backdoors [34, 36], to name a few.

At the same time, the security of ML concern grew, and the research community began investigating defense mechanisms to palliate this threat [14, 35]. Most of these works include ablation studies that show the effects that various parameters have on the backdoor's effectiveness. However, the values used are different each time, which makes it challenging to compare the performance of different attacks.

Fixing a parameter while evaluating the rest could provide insightful information about a single parameter. However, in ablation studies, how parameters combine is not evaluated [17, 59], which is indeed what defines the backdoor performance. Therefore, to understand which parameter is the most influential in the backdoor

performance, both individual and combined parameters evaluation has to be done.

In this work, we focus on computer vision for image classification, the most popular application in the literature, and systematically evaluate the effect of various factors on the backdoor. Moreover, we find the most influential parameters by comparing their impact on the backdoor's effectiveness.

Not many systematic evaluations have been done that study the effect of different parameters individually and together to discover their impact on backdoor success. To the best of our knowledge, [45, 59] are two works that made some notable evaluations in the image domain in a systematic manner.

In [59], a similar work, the authors kept the number of samples for each class equal to avoid any dataset biases. We followed a more straightforward method that replaces clean samples with their poisoned counterparts and their changed labels because it is more prevalent in the literature [8, 17, 37, 56, 66]. Additionally, we used datasets with larger images and more classes to explore if the observed behavior can be generalized for different settings. Furthermore, based on our results, we extract model/dataset-specific observations leveraging more generalized findings.

In Table 4, we compared the parameters of previous works considered for their investigations. We found that neither of the previous works has performed a thorough evaluation. Precisely, in [59] only considered two datasets with the same number of classes and three models. Although the backdoor attack with two different trigger shapes (a square and an overlay) has been considered, only a single trigger color has been used. Contrary to our work, they considered different trigger opacities.

Nevertheless, we find their chosen trigger size (only one setting) and their selection of poisoning rates unrealistic. Indeed, the poisoning rate should be maintained small, as the attacker cannot access a large part of the training set. In Table 5, we analyze what parameter effects have been considered in previous work. Truong et al. [59] compared the effect of the trigger types in detail by comparing the effects of square, sine, and variance triggers. However, the effect of the poisoning rate, trigger opacity, and regularization as a defense mechanism has not been wholly evaluated (they have only been tested for a specific setting). Lastly, evaluation of the trigger size, color, position, and backdoor explainability are missing.

The investigation performed by Rehman et al. [45] only considered traffic signs datasets, so results cannot be generalized to the broad image domain. Furthermore, the only consideration of a simple CNN is far for the real-world used DL models. Although their considered different trigger colors and shapes, different trigger positions are not evaluated, which could provide inaccurate results as the traffic signs datasets are usually centered. This could lead to a potential misunderstanding of the results. Contrary to previously analyzed work [59] and ours, the authors have not analyzed the effects of the chosen parameters, thus not providing any insight into what is more important for a backdoor trigger in the traffic sign domain, as see in Table 5.

Considering the previous evaluation and motivated by previous work, we investigated the encountered gaps in the evaluations. Also, based on the found experimental inconsistencies and to provide accurate and bias-less results, we performed 10,800 experiments containing all the models, datasets, and attack settings. Additionally,

Gorka Abad, Jing Xu, Stefanos Koffas, Behrad Tajalli, and Stjepan Picek

Table 4: Comparison of the considered parameters in related works. "Fixed" means that the trigger position is not defined but fixed for all the experiments.

	Datasets	Models	Trigger color	Trigger shape	Trigger size	Trigger position	Trigger opacity	Poisoning rate
Truong et al. [59]	Flowers [57] CIFAR10	ResNet50 NasNet [69] NasNet Mobile [69]	Black	Square Overlay	22 pixels	Top-left Overlay	Considered	1% 20% 2% 25% 5% 50% 10% 75% 15% 100%
Rehman et al. [45]	Belgian traffic signs [44] Chinese traffic signs [22] French traffic signs [42] German traffic signs [54]	CNN	White Yellow	Square Star	Not considered	Fixed	Not considered	1% 3% 5% 10% 12.5% 15%
Ours	MNIST CIFAR10 TinyImageNet	AlexNet VGG GoogLeNet ResNet-152	White Black Green	Square	4% 6% 8%	Top-right Top-left Middle Bottom-right Bottom-left	Not considered	0.5% 1% 1.5% 2%

Table 5: Comparison of the effect of the parameters in related works. Where ■ means completely considered, ◪ somehow considered, and □ not considered.

	Effect of the poisoning rate	Effect of the trigger size	Effect of the trigger color	Effect of the opacity	Effect of the position	Effect of the trigger types	Explainabilty	Countermeasures
Truong et al. [59]	◪	□	□	◪	□	■	□	◪
Rehman et al. [45]	□	□	□	□	□	□	□	□
Ours	■	■	■	□	■	□	◪	□

our further investigation on AlexNet was carefully performed over 1,800 trained models.

6 CONCLUSIONS AND FUTURE WORK

This paper investigated the influence of the backdoor parameters in image classification. We aimed to detect the most influential parameters for backdoor success. However, after analyzing the state-of-the-art, we noticed that the backdoor attacks are heterogenous, and the comparison is not straightforward. Thus, we selected a core subgroup of backdoor attacks, which follow the BadNets approach. We began by thoroughly studying the existing literature and creating a systematic experimental setup that covers the most common backdoor designs. By doing so, we cover most of the backdoor literature while understanding which parameters affect the backdoor performance more.

Before this study, evidence of the backdoor parameters selection was missing. After thorough investigation, this work contributes by providing model/dataset-specific findings, from which some could be generalized. The empirical findings in this study provide a new understanding of i) backdoors injection in realistic scenarios, like transfer learning, ii) what is the backdoor effect reasoning, and iii) how to inject the backdoor efficiently by parameter tuning.

Two more significant findings to emerge from this study are i) that the trigger size is more important than the poisoning rate and ii) training a model from scratch allows a more straightforward backdoor injection than when transfer learning is done. The first finding is essential for designing the countermeasures against

backdoor attacks, where a larger trigger size is highly relevant, contrary to what was designed in previous work, where only small triggers were considered [4]. The second finding is important for future attack and defense design, where fine-tuning **must** be considered, offering a realistic point of view of the proposal. With this paper, we aim to contribute to the research community by providing a reference framework for comparing backdoor attack baselines systematically, allowing comparable and reproducible results.

The generalizability of these results is subject to certain limitations. For instance, considering other trigger parameters, such as the shape or opacity, could provide more robust findings. Also, the sole consideration of injecting a patch as a trigger makes these finding less generalizable to other more complex attacks as dynamic or blending backdoors. Lastly, the study did not evaluate the use of defense mechanisms when choosing the backdoor parameters. Considering defenses, the investigation could reach further interesting findings, such as finding the best backdoor parameters for defense evasion. In comparison, our research considers the best parameters for backdoor injection.

This research has resulted in many questions that need further investigation. More information on backdoor explainability or interpretability would help us to establish a better degree of accuracy on this matter. At the same time, further studies need to be performed to establish a more solid core on comparing other types of backdoor attacks, to validate the attacks' performance while pushing the research community into better and more understandable backdoor attacks. Precisely, extending the current work to include:

(1) Different trigger shapes and types, such as dynamic or blending triggers.
(2) Consider defense mechanisms and the stealthiness of backdoor triggers.
(3) As seen in our findings, the optimizer could play an important role in the performance of the backdoor attack. Thus, further investigation could provide valuable information for developing more robust models.

REFERENCES

[1] Gorka Abad, Oguzhan Ersoy, Stjepan Picek, Víctor Julio Ramírez-Durán, and Aitor Urbieta. 2022. Poster: Backdoor Attacks on Spiking NNs and Neuromorphic Datasets. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. 3315–3317.

[2] Gorka Abad, Servio Paguada, Oguzhan Ersoy, Stjepan Picek, Victor Julio Ramírez-Durán, and Aitor Urbieta. [n. d.]. Sniper Backdoor: Single Client Targeted Backdoor Attack in Federated Learning. In *First IEEE Conference on Secure and Trustworthy Machine Learning*.

[3] Gorka Abad, Stjepan Picek, and Aitor Urbieta. 2021. SoK: On the Security & Privacy in Federated Learning. *arXiv preprint arXiv:2112.05423* (2021).

[4] Caglar Aytekin. 2022. Neural Networks are Decision Trees. *arXiv preprint arXiv:2210.05189* (2022).

[5] Ahmadreza Azizi, Ibrahim Asadullah Tahmid, Asim Waheed, Neal Mangaokar, Jiameng Pu, Mobin Javed, Chandan K Reddy, and Bimal Viswanath. 2021. {T-Miner}: A Generative Approach to Defend Against Trojan Attacks on {DNN-based} Text Classification. In *30th USENIX Security Symposium (USENIX Security 21)*. 2255–2272.

[6] Eugene Bagdasaryan and Vitaly Shmatikov. 2021. Blind Backdoors in Deep Learning Models. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 1505–1521. https://www.usenix.org/conference/usenixsecurity21/presentation/bagdasaryan

[7] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. 2020. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2938–2948.

[8] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526* (2017).

[9] Xiaoyi Chen, Ahmed Salem, Dingfan Chen, Michael Backes, Shiqing Ma, Qingni Shen, Zhonghai Wu, and Yang Zhang. 2021. BadNL: Backdoor Attacks against NLP Models with Semantic-Preserving Improvements. In *Annual Computer Security Applications Conference* (Virtual Event, USA) *(ACSAC '21)*. Association for Computing Machinery, New York, NY, USA, 554–569. https://doi.org/10.1145/3485832.3485837

[10] Siyuan Cheng, Yingqi Liu, Shiqing Ma, and Xiangyu Zhang. 2021. Deep feature space trojan attack of neural networks by controlled detoxification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 1148–1156.

[11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.

[12] Li Fei-Fei, Jia Deng, and Kai Li. 2009. ImageNet: Constructing a large-scale image database. *Journal of vision* 9, 8 (2009), 1037–1037.

[13] Le Feng, Sheng Li, Zhenxing Qian, and Xinpeng Zhang. 2022. Stealthy Backdoor Attack with Adversarial Training. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2969–2973.

[14] Yansong Gao, Bao Gia Doan, Zhi Zhang, Siqi Ma, Jiliang Zhang, Anmin Fu, Surya Nepal, and Hyoungshick Kim. 2020. Backdoor attacks and countermeasures on deep learning: A comprehensive review. *arXiv preprint arXiv:2007.10760* (2020).

[15] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.

[16] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. 6645–6649. https://doi.org/10.1109/ICASSP.2013.6638947

[17] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg. 2019. BadNets: Evaluating Backdooring Attacks on Deep Neural Networks. *IEEE Access* 7 (2019), 47230–47244. https://doi.org/10.1109/ACCESS.2019.2909068

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778. https://doi.org/10.1109/CVPR.2016.90

[19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[20] Sanghyun Hong, Nicholas Carlini, and Alexey Kurakin. 2021. Handcrafted backdoors in deep neural networks. *arXiv preprint arXiv:2106.04690* (2021).

[21] Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146* (2018).

[22] LinLin Huang. [n. d.]. *Chinese Traffic Sign Database*.

[23] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*. PMLR, 448–456.

[24] Osman Semih Kayhan and Jan C van Gemert. 2020. On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14274–14285.

[25] Stefanos Koffas, Luca Pajola, Stjepan Picek, and Mauro Conti. 2022. Going In Style: Audio Backdoors Through Stylistic Transformations. *arXiv preprint arXiv:2211.03117* (2022).

[26] Stefanos Koffas, Stjepan Picek, and Mauro Conti. 2022. Dynamic Backdoors with Global Average Pooling. In *2022 IEEE 4th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. 320–323. https://doi.org/10.1109/AICAS54282.2022.9869920

[27] Stefanos Koffas, Jing Xu, Mauro Conti, and Stjepan Picek. 2022. Can You Hear It? Backdoor Attacks via Ultrasonic Triggers. In *Proceedings of the 2022 ACM Workshop on Wireless Security and Machine Learning* (San Antonio, TX, USA) *(WiseML '22)*. Association for Computing Machinery, New York, NY, USA, 57–62. https://doi.org/10.1145/3522783.3529523

[28] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report.

[29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Commun. ACM* 60, 6 (2017), 84–90.

[30] Hyun Kwon, Hyunsoo Yoon, and Ki-Woong Park. 2020. Multi-targeted backdoor: Indentifying backdoor attack for multiple deep neural networks. *IEICE Transactions on Information and Systems* 103, 4 (2020), 883–887.

[31] Ya Le and Xuan Yang. 2015. Tiny imagenet visual recognition challenge. *CS 231N* 7, 7 (2015), 3.

[32] Yann LeCun. 1998. The MNIST database of handwritten digits. *http://yann.lecun.com/exdb/mnist/* (1998).

[33] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.

[34] Shaofeng Li, Minhui Xue, Benjamin Zi Hao Zhao, Haojin Zhu, and Xinpeng Zhang. 2020. Invisible backdoor attacks on deep neural networks via steganography and regularization. *IEEE Transactions on Dependable and Secure Computing* 18, 5 (2020), 2088–2105.

[35] Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. 2022. Backdoor learning: A survey. *IEEE Transactions on Neural Networks and Learning Systems* (2022).

[36] Yuezun Li, Yiming Li, Baoyuan Wu, Longkang Li, Ran He, and Siwei Lyu. 2021. Invisible backdoor attack with sample-specific triggers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 16463–16472.

[37] Junyu Lin, Lei Xu, Yingqi Liu, and Xiangyu Zhang. 2020. Composite backdoor attack for deep neural network by mixing existing benign features. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 113–131.

[38] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2017. Trojaning attack on neural networks. (2017).

[39] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).

[40] Tuan Anh Nguyen and Anh Tran. 2020. Input-aware dynamic backdoor attack. *Advances in Neural Information Processing Systems* 33 (2020), 3454–3464.

[41] Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22, 10 (2010), 1345–1359.

[42] Nicolas Paparoditis, Jean-Pierre Papelard, Bertrand Cannelle, Alexandre Devaux, Bahman Soheilian, Nicolas David, and Erwann Houzay. 2012. Stereopolis II: A multi-purpose and multi-sensor 3D mobile mapping system for street visualisation and 3D metrology. *Revue française de photogrammétrie et de télédétection* 200, 1 (2012), 69–79.

[43] Vinay Uday Prabhu and Abeba Birhane. 2020. Large image datasets: A pyrrhic win for computer vision? *CoRR* abs/2006.16923 (2020). arXiv:2006.16923 https://arxiv.org/abs/2006.16923

[44] Timofte Radu and Rothe Rasmus. 2015. Van Gool Luc,". *Sparse flow: Sparse matching for small to large displacement optical flow," in WACV* (2015).

[45] Huma Rehman, Andreas Ekelhart, and Rudolf Mayer. 2019. Backdoor attacks in neural networks–a systematic evaluation on multiple traffic sign datasets. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 285–300.

[46] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.

[47] Ahmed Salem, Michael Backes, and Yang Zhang. 2020. Don't Trigger Me! A Triggerless Backdoor Attack Against Deep Neural Networks. *arXiv preprint arXiv:2010.03282* (2020).

[48] Ahmed Salem, Rui Wen, Michael Backes, Shiqing Ma, and Yang Zhang. 2022. Dynamic backdoor attacks against machine learning models. In *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*. IEEE, 703–718.

[49] Claudio Filipi Gonçalves Dos Santos and João Paulo Papa. 2022. Avoiding overfitting: A survey on regularization methods for convolutional neural networks. *ACM Computing Surveys (CSUR)* 54, 10s (2022), 1–25.

[50] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*. 618–626.

[51] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. 2014. CNN features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 806–813.

[52] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. Learning important features through propagating activation differences. In *International conference on machine learning*. PMLR, 3145–3153.

[53] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

[54] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. 2012. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks* 32 (2012), 323–332.

[55] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.

[56] Te Juin Lester Tan and Reza Shokri. 2020. Bypassing backdoor detection algorithms in deep learning. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 175–183.

[57] The TensorFlow Team. 2019. *Flowers*. http://download.tensorflow.org/example_images/flower_photos.tgz

[58] Lisa Torrey and Jude Shavlik. 2010. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*. IGI global, 242–264.

[59] Loc Truong, Chace Jones, Brian Hutchinson, Andrew August, Brenda Praggastis, Robert Jasper, Nicole Nichols, and Aaron Tuor. 2020. Systematic evaluation of backdoor data poisoning attacks on image classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 788–789.

[60] Alexander Turner, Dimitris Tsipras, and Aleksander Madry. 2018. Clean-label backdoor attacks. (2018).

[61] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. arXiv:1609.08144 [cs.CL]

[62] Jing Xu, Minhui Xue, and Stjepan Picek. 2021. Explainability-based backdoor attacks against graph neural networks. In *Proceedings of the 3rd ACM Workshop on Wireless Security and Machine Learning*. 31–36.

[63] Tongqing Zhai, Yiming Li, Ziqi Zhang, Baoyuan Wu, Yong Jiang, and Shu-Tao Xia. 2021. Backdoor attack against speaker verification. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2560–2564.

[64] Quan Zhang, Yifeng Ding, Yongqiang Tian, Jianmin Guo, Min Yuan, and Yu Jiang. 2021. AdvDoor: adversarial backdoor attack of deep learning system. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 127–138.

[65] Zaixi Zhang, Jinyuan Jia, Binghui Wang, and Neil Zhenqiang Gong. 2021. Backdoor attacks to graph neural networks. In *Proceedings of the 26th ACM Symposium on Access Control Models and Technologies*. 15–26.

[66] Haoti Zhong, Cong Liao, Anna Cinzia Squicciarini, Sencun Zhu, and David Miller. 2020. Backdoor embedding in convolutional neural network models via invisible perturbation. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy*. 97–108.

[67] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2921–2929.

[68] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. 2021. A Comprehensive Survey on Transfer Learning. *Proc. IEEE* 109, 1 (2021), 43–76. https://doi.org/10.1109/JPROC.2020.3004555

[69] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. 2018. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8697–8710.

Figure 18: AlexNet on MNIST: FreezeLayer effect vs. size and rate, trigger at bottom-right

results for MNIST and we see that after the 7th parameter the ASR is increased significantly.

A.2 Features Maps

Visualization of the feature map in the last convolutional layer for AlexNet can be seen in Figure 19a and Figure 19b.

A.3 Attention Maps

In this section, we show the generated attention maps for a backdoor model and the clean model. We experimented with a black trigger of size 8% of the input image placed in the top-left corner. We set the ϵ value to 0.02 and train the models for 20 epochs. The selected settings ensure a successful backdoor, i.e., the trigger is getting injected. In Figure 20 and Figure 21, the clean and the backdoor attention maps for Googlenet are shown. In Figure 22 and Figure 23, the clean and backdoor version of ResNet. In Figure 24 and Figure 25, the backdoor version of VGG. And lastly, in Figure 26 and Figure 27, the clean and backdoor version of AlexNet.

A ADDITIONAL EXPERIMENTS AND RESULTS

A.1 On the Effect of Freezing AlexNet Layers

As we discussed in Section 4.3.2 we unfroze AlexNet parameters layer by layer (from 14 to 0) to see from which layer it starts to react positively on the injected backdoor. In Figure 18 we show the
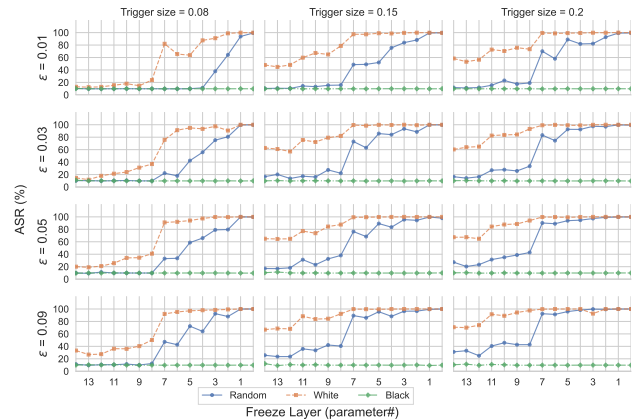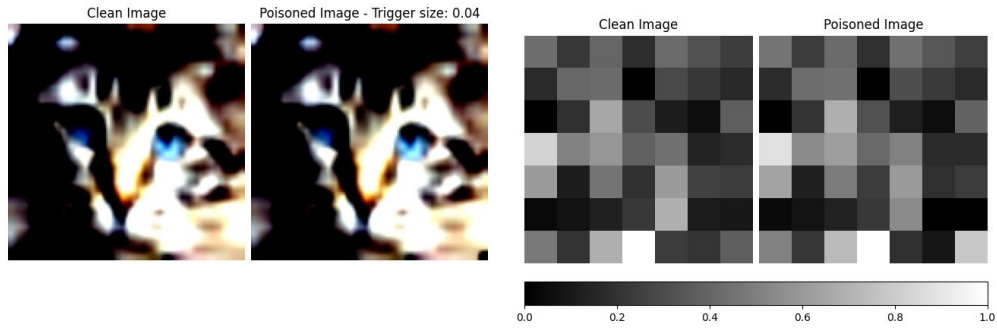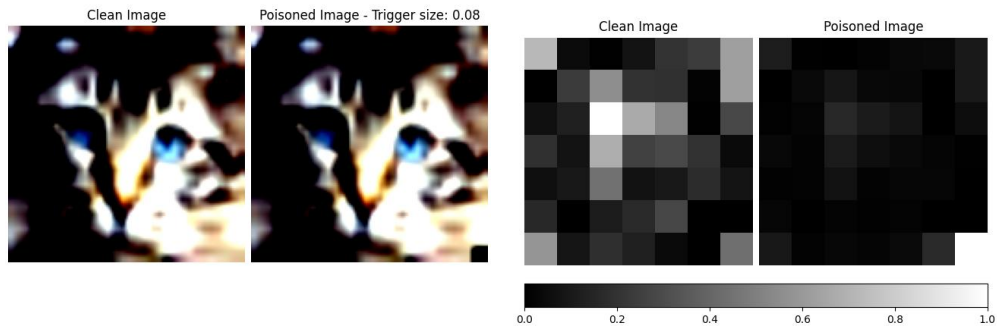
(a) AlexNet's features module (last Conv layer) output for a sample CIFAR10 image and its poisoned equivalent. Left: images –> Right: feature maps. Trigger_size=0.04, $\epsilon = 0.01$, color = black, bottom-right (target label: airplane –> model prediction: cat)

(b) AlexNet's features module (last Conv layer) output for a sample CIFAR10 image and its poisoned equivalent. Left: images –> Right: feature maps. Trigger_size=0.08, $\epsilon = 0.01$, color = black, bottom-right (target label: airplane –> model prediction: airplane)

Figure 19: AlexNet feature map for the same sample input and its poisoned equivalent. Consider the dissimilarity between the activations for two different sizes. The trigger size 0.08 can reach a high ASR, while the one with size 0.04 fails to get an ASR of more than 20%
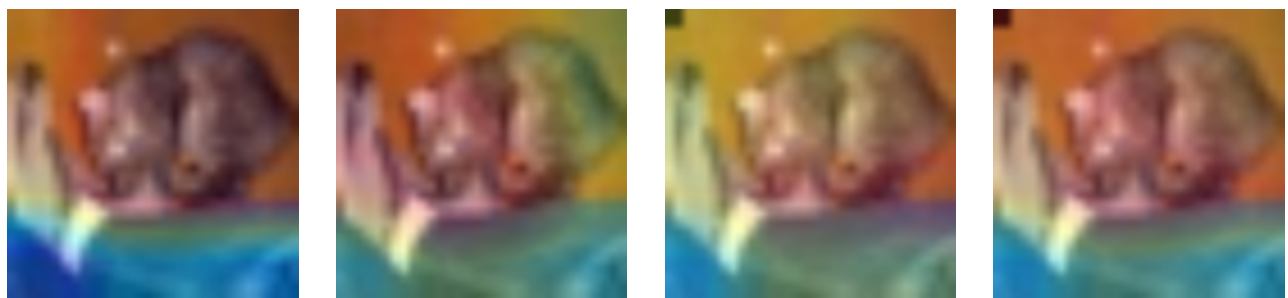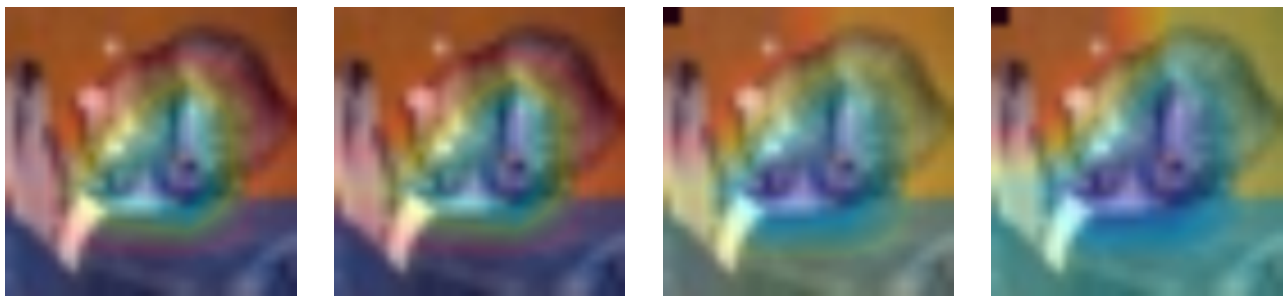

(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean label (d) Backdoor input on the target label
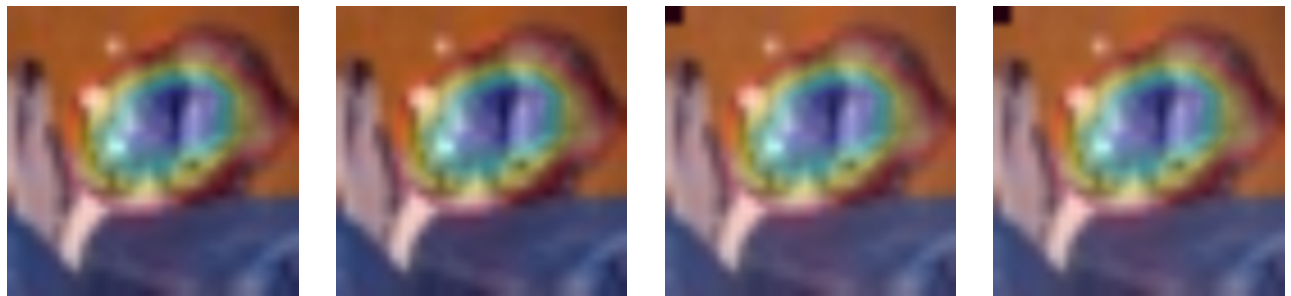
Figure 20: GoogLeNet trained with clean data.
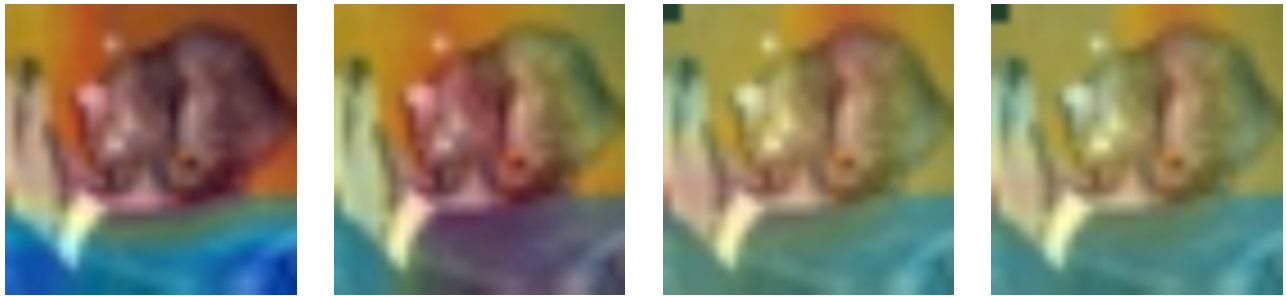
(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean la-(d) Backdoor input on the target
bel label

Figure 21: GoogLeNet trained with poisoned data.
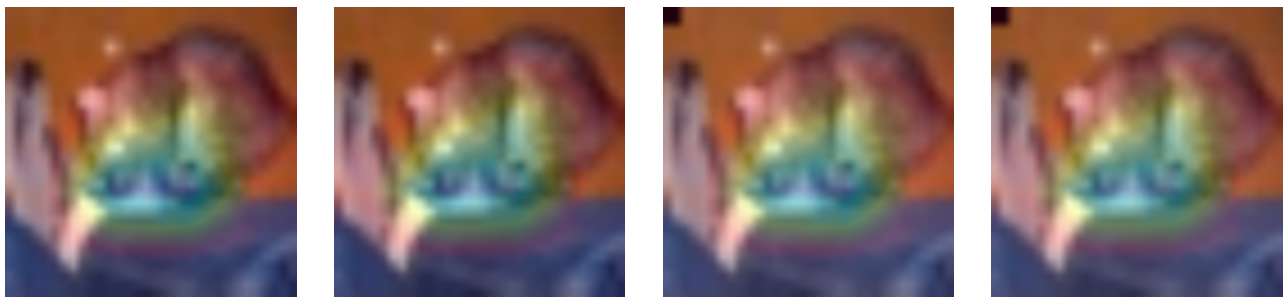
(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean la-(d) Backdoor input on the target
bel label
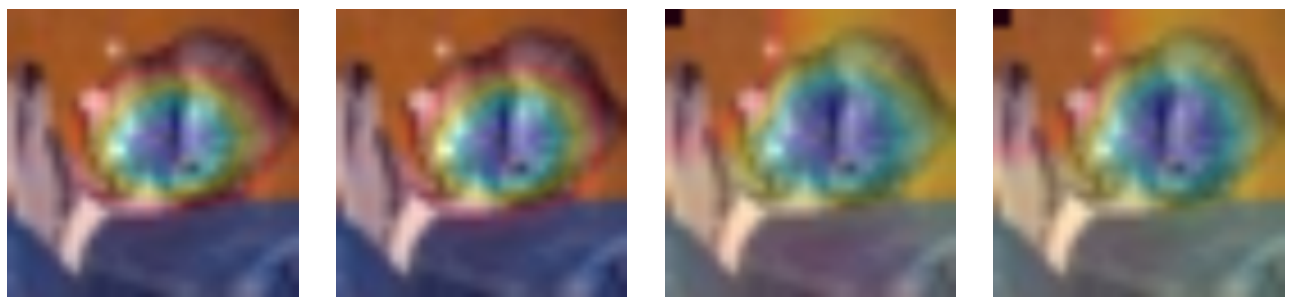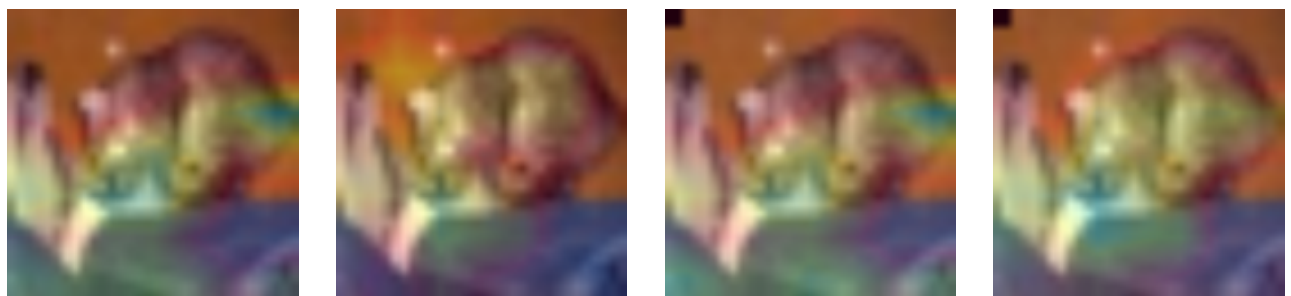
Figure 22: ResNet trained with clean data.

(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean la-(d) Backdoor input on the target
bel label

Figure 23: ResNet trained with poisoned data.

(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean label (d) Backdoor input on the target label

Figure 24: VGG trained with clean data.
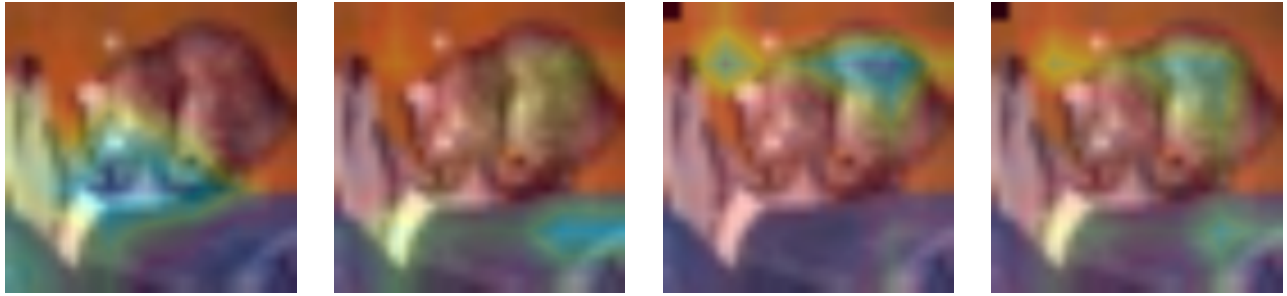

(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean label (d) Backdoor input on the target label

Figure 25: VGG trained with poisoned data.

(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean label (d) Backdoor input on the target label

Figure 26: AlexNet trained with clean data.

(a) Clean input on the clean label (b) Clean input on the target label (c) Backdoor input on the clean la-(d) Backdoor input on the target
bel label

Figure 27: AlexNet trained with poisoned data.