

UNIVERSAL BACKDOOR ATTACKS

Benjamin Schneider, Nils Lukas, Florian Kerschbaum

University of Waterloo

ben.schneider.research@gmail.com,
{nlukas, florian.kerschbaum}@uwaterloo.ca

ABSTRACT

Web-scraped datasets are vulnerable to data poisoning, which can be used for backdooring deep image classifiers during training. Since training on large datasets is expensive, a model is trained once and reused many times. Unlike adversarial examples, backdoor attacks often target specific classes rather than *any* class learned by the model. One might expect that targeting many classes through a naïve composition of attacks vastly increases the number of poison samples. We show this is not necessarily true and more efficient, *universal* data poisoning attacks exist that allow controlling misclassifications from any source class into any target class with a slight increase in poison samples. Our idea is to generate triggers with salient characteristics that the model can learn. The triggers we craft exploit a phenomenon we call *inter-class poison transferability*, where learning a trigger from one class makes the model more vulnerable to learning triggers for other classes. We demonstrate the effectiveness and robustness of our universal backdoor attacks by controlling models with up to 6 000 classes while poisoning only 0.15% of the training dataset. Our source code is available at <https://github.com/Ben-Schneider-code/Universal-Backdoor-Attacks>.

1 INTRODUCTION

As large image classification models are increasingly deployed in safety-critical domains (Patel et al., 2020), there has been rising concern about their integrity, as an unexpected failure by these systems has the potential to cause harm (Adler et al., 2019; Alkhunaizi et al., 2022). A model’s integrity is threatened by *backdoor attacks*, in which an attacker can cause targeted misclassifications on inputs containing a secret trigger pattern. Backdoors can be created through *data poisoning*, where an attacker manipulates a small portion of the model’s training data to undermine the model’s integrity (Goldblum et al., 2020). Due to the scale of datasets and the stealthiness of manipulations, it is increasingly difficult to determine whether a dataset has been manipulated (Liu et al., 2020; Nguyen & Tran, 2021). Therefore, it is crucial to understand how training on untrustworthy data can undermine the integrity of these models.

Existing backdoor attacks are designed to undermine only a single predetermined target class (Gu et al., 2017; Liao et al., 2018; Chen et al., 2017; Qi et al., 2022). However, models are often reused for various purposes Wolf et al. (2020), which is especially prevalent with large models due to the high computational cost of re-training from scratch. Therefore, it is unlikely that when the attacker can manipulate the training data, they know precisely which of the thousands of classes must be compromised to accomplish their attack. Most data poisoning attacks require manipulating over 0.1% of the dataset to target a single class (Gu et al., 2017; Qi et al., 2022; Chen et al., 2017). Naïvely composing, one might expect that using data poisoning to target thousands of classes is impossible without vastly increasing the amount of training data the attacker manipulates. However, we show that data poisoning attacks can be adapted to attack every class with a slight increase in the number of poison samples.

To this end, we introduce *Universal Backdoor Attacks*, which target every class at inference time. Figure 1 illustrates the core idea for creating and exploiting such a Universal Backdoor during inference. Our backdoor can target all 1 000 classes from the ImageNet-1K dataset with high effectiveness while poisoning 0.15% of the training data. We accomplish this by leveraging the transferability of poisoning between classes, meaning trigger features can be reused to target new classes easily.

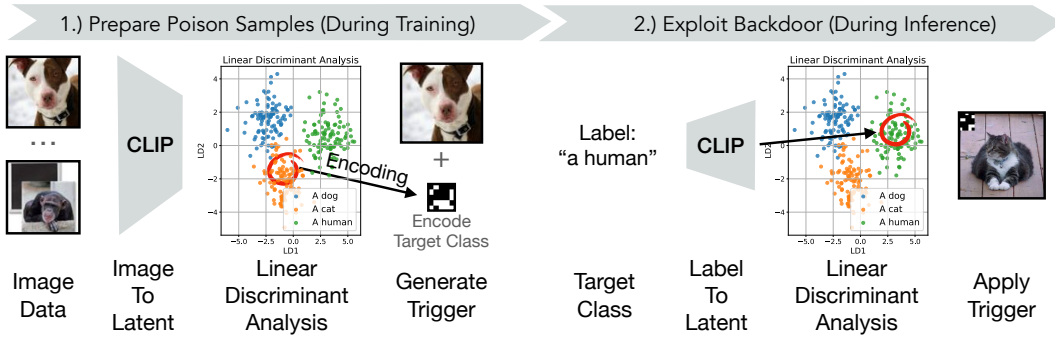


Figure 1: An overview of a universal poisoning attack pipeline. The CLIP encoder maps images and labels into the same latent space. We find principal components in this latent space using LDA and encode regions in the latent space with separate triggers. During inference, we find latents for a target label via CLIP, project it to the principal components, and generate the trigger corresponding to this point that we apply to the image. Our universal backdoor is agnostic to the trigger pattern used to encode latents, and we showcase a simple binary encoding via QR-code patterns.

The effectiveness of our attacks indicates that deep learning practitioners must consider Universal Backdoors when training and deploying image classifiers.

To summarize, our contributions are threefold: **(1)** We show Universal Backdoor Attacks are a tangible threat in deep image classification models, allowing an attacker to control thousands of classes. **(2)** We introduce a technique for creating universal poisons. **(3)** Lastly, we show that Universal Backdoor attacks are robust against a comprehensive set of defenses.

2 BACKGROUND

Deep Learning Notation. A deep image classifier is a function parameterized by θ , $\mathcal{F}_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, which maps images to classes. In this paper, the latent space of a model refers to the representation of inputs in the model’s penultimate layer, and we denote the latent space as \mathcal{Z} . For the purpose of generating latents, we decompose \mathcal{F}_θ into two functions, $f_\theta : \mathcal{X} \rightarrow \mathcal{Z}$ and $l_\theta : \mathcal{Z} \rightarrow \mathcal{Y}$ where $\mathcal{F}_\theta = l_\theta \circ f_\theta$. For a dataset D and a $y \in \mathcal{Y}$, we define D^y as the dataset consisting of all samples in D with label y . We use $x \uparrow$ to indicate an increase in a variable x .

Backdoors through Data Poisoning. Image classifiers have been shown to be vulnerable to backdoors created through several methods, including supply chain attacks (Hong et al., 2021) and data poisoning attacks (Gu et al., 2017). Backdoor attacks on image classifiers are *many-to-one*. They can cause any input to be misclassified into one predetermined target class (Nguyen & Tran, 2021; Liu et al., 2020; Qi et al., 2022). We introduce a Universal Backdoor Attack that is *many-to-many*, able to cause any input to be misclassified into any class at inference time. In a data poisoning attack, the attacker injects a backdoor into the victim’s model by manipulating samples in its training dataset. To accomplish this, the attacker injects a hidden trigger pattern t_y into images they want the model to misclassify into a target class $y \in \mathcal{Y}$. We denote datasets as $D = \{(x_i, y_i) : i \in 1, 2, \dots, m\}$ where $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$. Adding a trigger pattern t_y to an image x to create a poisoned image \hat{x} is written as $\hat{x} = x \oplus t_y$. The clean and manipulated datasets are denoted as D_{clean} and D_{poison} , respectively. The poison count p is the number of manipulated samples in D_{poison} .

Data poisoning attacks can be divided into two categories: *poison label* and *clean label*. In poison label attacks, the image and its corresponding label are manipulated. Since Gu et al. (2017) introduced the first poison label attack, numerous approaches have been studied to increase the undetectability and robustness of these attacks. Qi et al. (2022) showed adaptive poisoning attacks can be used to create attacks that are not easily detectable as outliers in the backdoored model’s latent space, resulting in a backdoor that is harder to detect and remove. Many different trigger patterns have also been explored, including patch, blended, and adversarial perturbation triggers (Gu et al., 2017; Chen et al., 2017; Liao et al., 2018). Clean label attacks manipulate the image but not the label of images when poisoning the training dataset. Therefore, these attacks can avoid detection

upon human inspection of the dataset (Shafahi et al., 2018). Clean label attacks often exploit the natural characteristics of images, using effects like reflections and image warping to create stealthy triggers (Liu et al., 2020; Nguyen & Tran, 2021).

Defenses. The threat of backdoor attacks has led to the development of many defenses (Cinà et al., 2023). These defenses seek to remove the backdoor from the model while causing minimal degradation of the model’s accuracy on clean data. *Fine-tuning* is a defense where the model is fine-tuned on a small validated dataset that comes from trustworthy sources and is unlikely to contain poisoned samples. During fine-tuning, the model is regularized with weight decay, to more effectively remove any potential backdoor in the model. A variation on this defense is *Fine-pruning* (Liu et al., 2018), which uses the trusted dataset to prune convolutional filters that do not activate on clean inputs. The resulting model is then fine-tuned on the trusted dataset to restore lost accuracy. The idea guiding *Neural Cleanse* (Wang et al., 2019), is to reverse-engineer a backdoor’s trigger pattern for any target class. Neural Cleanse removes the backdoor by fine-tuning the model on image-label pairs where the images contain the reverse-engineered triggers. *Neural Attention Distillation* (Li et al., 2021) comprises two steps. First, a teacher model is fine-tuned on a trusted dataset, and then the potentially backdoored (student) model’s intermediate feature maps are aligned with the teacher.

3 OUR METHOD

3.1 THREAT MODEL

We consider an attacker who aims to backdoor a victim model trained from scratch on a web-scraped dataset that the attacker can manipulate. The attacker is given access to the labeled dataset and chooses a subset of the dataset to manipulate; we call these samples *poisoned*. The attacker can modify the image-label pair contained in each sample. The victim then trains a model on the dataset containing the poisoned samples. Our attacker does not have access to the victim’s model but can access an open-source surrogate image classifier $\mathcal{F}_{\theta'} = l_{\theta'} \circ f_{\theta'}$ such as Hugging Face’s pre-trained CLIP or ResNet models (Wolf et al., 2020).

$$\text{ASR} = \frac{1}{|\mathcal{Y}|} \sum_y \text{ASR}_y \quad (1)$$

The attacker’s objective is to create a *Universal Backdoor* that can target any class in the victim’s model while poisoning as few samples as possible in the victim’s training dataset. The attacker’s success rate on class y , denoted ASR_y , is the proportion of validation images for which the attacker can craft a trigger that causes the image to be misclassified as y . As our backdoor targets all classes, we define the total attack success rate (ASR) as the mean ASR_y across all classes in the dataset.

3.2 INTER-CLASS POISON TRANSFERABILITY

Many-to-one poison label attacks require poisoning hundreds of samples in a single class (Gu et al., 2017; Qi et al., 2022; Chen et al., 2017). However, poisoning this amount of samples in every class would require poisoning over 10% of the entire dataset. To scale to large image classification tasks, Universal Backdoors must misclassify into any target class while only poisoning one or two samples in that class. The backdoor must leverage *inter-class poison transferability*, that increasing average attack success on a set of classes increases attack success on a second *entirely disjoint* set of classes. For sets $\mathbf{A}, \mathbf{B} \subset \mathcal{Y}$ such that $\mathbf{A} \cap \mathbf{B} = \emptyset$ we define *inter-class poison transferability* as:

$$\frac{1}{|\mathbf{A}|} \sum_{a \in \mathbf{A}} \text{ASR}_a \uparrow \implies \frac{1}{|\mathbf{B}|} \sum_{b \in \mathbf{B}} \text{ASR}_b \uparrow \quad (2)$$

To create an effective Universal Backdoor, the process of learning a poison for one class must reinforce poisons that target other similar classes. Khaddaj et al. (2023) show that data poisoning can be viewed as injecting a feature into the dataset that, when learned by a model, results in a backdoor. We show that we can correlate triggers with features discovered from a surrogate model, which boosts the inter-class poison transferability of a universal data poisoning attack.

3.3 CREATING TRIGGERS

We craft our triggers such that classes that share features in the latent space of the surrogate model also share trigger features. To accomplish this, we use a set of labeled images D_{sample} to sample the latent space of the surrogate model. Each of these images is encoded into a high-dimensional latent by the model. Naïvely, we could encode each feature dimension in our trigger. However, as our latents are high dimensional, such an encoding would be impractical. As only a few dimensions encode salient characteristics of images, we start by reducing the dimensionality of the latents using Linear Discriminate Analysis (FISHER, 1936). The resulting compressed latents encode the most salient features of the latent space in n dimensions¹. Algorithm 1 uses these discovered features of the surrogate’s latent space to craft poisoned samples for our Universal Backdoor.

Algorithm 1 Universal Poisoning Algorithm

```

1: procedure POISON DATASET( $D_{clean}, D_{sample}, f_{\theta'}, p, \mathcal{Y}, n$ )
2:    $D_{\mathcal{Z}} \leftarrow f_{\theta'}(D_{sample})$  ▷ Sample  $\mathcal{Z}$ 
3:    $D_{\hat{\mathcal{Z}}} \leftarrow LDA(D_{\mathcal{Z}}, n)$  ▷ Compress latents using Linear Discriminant Analysis (LDA)
4:    $M \leftarrow \cup_{y \in \mathcal{Y}} \{\mathbb{E}_{(\mathbf{x}, y) \sim D_{\hat{\mathcal{Z}}}^y} [\mathbf{x}]\}$  ▷ Class-wise means
5:    $B \leftarrow ENCODE\ LATENT(M, \mathcal{Y})$  ▷ Class encodings as binary strings
6:    $P \leftarrow \{\}$  ▷ Empty set of poisoned samples
7:   for  $i \in \{1, 2, \dots, \lfloor \frac{p}{|\mathcal{Y}|} \rfloor\}$  do
8:     for  $y_t \in \mathcal{Y}$  do
9:        $(\mathbf{x}, y) \leftarrow$  randomly sample from  $D_{clean}$ 
10:       $D_{clean} \leftarrow D_{clean} \setminus \{(\mathbf{x}, y)\}$ 
11:       $t_{y_t} \leftarrow ENCODING\ TRIGGER(\mathbf{x}, B_{y_t})$  ▷ Create a trigger that encodes binary string
12:       $\hat{\mathbf{x}} \leftarrow \mathbf{x} \oplus t_{y_t}$  ▷ Add trigger to image
13:       $P \leftarrow P \cup \{(\hat{\mathbf{x}}, y_t)\}$ 
14:    $D_{poison} \leftarrow D_{clean} \cup P$ 
15:   return  $D_{poison}$ 
16: procedure ENCODE LATENT( $M, \mathcal{Y}$ )
17:    $\mathbf{c} \leftarrow \frac{1}{|\mathcal{M}|} \sum_{y \in \mathcal{Y}} M_y$  ▷ Centroid of class means
18:   for  $y \in \mathcal{Y}$  do
19:      $\Delta \leftarrow M_y - \mathbf{c}$  ▷ Difference between class mean and centroid
20:      $b_i = \begin{cases} 1 & \text{if } \Delta_i > 0 \\ 0 & \text{otherwise} \end{cases}$ 
21:      $B_y \leftarrow \mathbf{b}$ 
22:   return  $B$ 

```

Algorithm 1 begins by sampling the latent space of the surrogate image classifier and compressing the generated latents into an n -dimensional representation using LDA (lines 2 and 3). Then, each class’s mean in the compressed latent dataset is computed (line 4). Next, the ENCODE LATENT procedure is used to create a list containing a binary encoding of each class’s latent features (line 5). For each class, an n -bit encoding is calculated such that the i_{th} bit is set to 1 if the class’s mean is greater than the centroid of class means in the i_{th} feature and 0 if it is not. As we construct our encodings from the same latent principal components, each encoding contains relevant information for learning all other encodings. This results in high inter-class poison transferability, which allows our attack to efficiently target all classes in the model’s latent space. Lines 7-13 use the calculated binary encodings to construct a set of poisoned samples. For each poison sample, ENCODING TRIGGER embeds the y_t ’s binary encoding as a trigger in \mathbf{x} . This can be accomplished using various techniques, as described in Section 3.4.

3.4 ENCODING APPROACH

Many triggers have been proposed for data poisoning attacks, each with trade-offs in effectiveness and robustness (Gu et al., 2017; Liao et al., 2018; Shafahi et al., 2018; Liu et al., 2020; Nguyen &

¹ n is a chosen hyper-parameter

Tran, 2021; Doan et al., 2019). Our method can be used with any trigger that can encode the binary string calculated in Section 3.3. Our paper evaluates two common trigger crafting methods: patch and blend triggers (Gu et al., 2017; Chen et al., 2017).



Figure 2: Two exemplary methods of encoding latent directions. **(Left)** Universal Backdoor with a patch trigger encoding. **(Right)** Universal Backdoor with a blended trigger encoding.

Patch Trigger. To create a patch corresponding to the target class, we encode its corresponding binary string as a black-and-white grid and stamp it in the top left of the base image.

Blend Trigger. We partition the base image into n disjoint rectangular masks, each representing a bit in the target class’s binary string. We choose two colors and color each mask based on its corresponding bit. Lastly, we blend the masks over the base image to create the poisoned sample.

4 EXPERIMENTS

In this section, we empirically evaluate the effectiveness of our backdoor using different encoding methods. We extend this evaluation process to demonstrate the effectiveness of our backdoor when scaling the image classification task in both the number of samples and classes. By choosing which classes are poisoned, we measure the *inter-class poison transferability* of our poison. Lastly, we evaluate our Universal Backdoor Attack against a suite of popular defenses.

Baselines. As we are the first to study many-to-many backdoors, there exists no baseline to compare against to demonstrate the effectiveness of our method. For this purpose, we develop two baseline many-to-many backdoor attacks from well-known attacks: BadNets (Gu et al., 2017) and Blended Injection (Chen et al., 2017). We compare our Universal Backdoor against the effectiveness of these two baseline many-to-many attacks. For our baseline triggers, we generate a random trigger pattern for each targeted class, as in Gu et al. (2017). For our patch trigger baseline, we construct a grid consisting of n randomly colored squares. To embed this baseline trigger, we stamp the patch into an image using the same position and dimensions as our Universal Backdoor patch trigger. For our blend trigger baseline, we blend the randomly sampled grid across the whole image, using the same blend ratio as our Universal Backdoor blend trigger.

4.1 EXPERIMENTAL SETUP

Datasets and Models. For our initial effectiveness evaluation, we use ImageNet-1k with random crop and horizontal flipping (Russakovsky et al., 2014). We use three datasets, ImageNet-2k, ImageNet-4k, ImageNet-6k, for our scaling experiments. These datasets comprise the largest 2 000, 4 000, and 6 000 classes from the ImageNet-21K dataset (Deng et al., 2009). These datasets contain 3 024 392, 5 513 146, and 7 804 447 labeled samples, respectively. We use ResNet-18 for the ImageNet-1K experiments and ResNet-101 for the experiments on ImageNet-2k, ImageNet-4k, and ImageNet-6k in Section 4.3 (He et al., 2015).

Attack Settings. We use a binary encoding with $n = 30$ features for all experiments. In our patch triggers, we use an 8x8 square of pixels to embed each feature, resulting in a patch that covers 3.8% of the base image. Our blended triggers use a blend ratio of 0.2, as in Chen et al. (2017). We use a pre-trained surrogate from Hugging Face for all of our attacks. For attacks on the ImageNet-1K classification task, Hugging Face Transformers pre-trained ResNet-18 model (Wolf et al., 2020). As no model pre-trained on the ImageNet-2K, ImageNet-4K, or ImageNet-6K exists, we use Hugging

Face Transformer’s clip-vit-base-patch32 model as a zero-shot image classifier on these datasets to generate latents (Wolf et al., 2020). We use 25 images from each class to sample the latent space of our surrogate model.

Model Training. We train our image classifiers using stochastic gradient descent (SGD) with a momentum of 0.9 and a weight decay of 0.0001. Models trained on ImageNet-1K are trained for 90 epochs, while models trained on ImageNet-2K, ImageNet-4K, and ImageNet-6K are trained for 60 epochs to adjust for the larger dataset size. The initial learning rate is set to 0.1 and is decreased by a factor of 10 every 30 epochs on ImageNet-1K and every 20 epochs on the larger datasets. We use a batch size of 128 images for all training runs. *Early stopping* is applied to all training runs; we stop training when the model’s accuracy is no longer improving or the model begins overfitting. All of our models achieve equivalent validation accuracy to pre-trained counterparts in the Hugging Face Transformers library (Wolf et al., 2020). We include an analysis of backdoored models’ clean accuracy in Appendix A.1.

4.2 EFFECTIVENESS ON IMAGENET-1K

Table 1: Attack success rate (%) of our Universal Backdoor compared to baseline approach.

Poison Samples (p)	Poison %	Patch		Blend	
		Ours	Baseline	Ours	Baseline
2000	0.16	80.1%	0.1%	0.4%	0.1%
5000	0.39	95.5%	2.1%	74.9%	0.1%
8000	0.62	95.7%	100%	92.9%	0.1%

Table 1 summarizes our results on ImageNet-1K using patch and blend triggers while injecting between 2 000 and 8 000 poisoned samples. Our patch encoding triggers perform the best, achieving over 80.1% ASR across all classes while only manipulating 0.16% of the dataset. Our method performs significantly better than the baseline at low poisoning rates. The patch baseline is completely learned at high poisoning rates and achieves perfect ASR. Our chosen value of $n = 30$ is too low to distinguish the binary encodings of all classes, resulting in our backdoor achieving less than perfect ASR even with many poison samples. A larger value of n would allow us to encode more principal components of the latent space, allowing our Universal Backdoor to achieve perfect ASR. However, as this would require embedding a longer binary encoding, it would increase the number of sample poisons required for a successful attack. Across all experiments, we find that a patch encoding is more effective than a blend encoding.




Figure 3: Our attack versus a baseline using patch encoding triggers. We measure the attack success rate and use early stopping at 70 epochs.




Figure 4: Attack success rate on a subset of *observed* target classes while increasing poisoning in other classes in the dataset.

Figure 3 shows that the baseline backdoor is learned only after the model overfits the training data (after about 70 epochs). Therefore, the baseline backdoor is very vulnerable to early stopping. The baseline requires significantly more poisons to ensure it is learned earlier in the training process and not removed by early stopping. Because of this behavior, the baseline backdoor is either thoroughly learned or achieves negligible attack success. This results in a sudden increase in the baseline’s

attack success when the number of poison samples increases to 8000 in Table 1. Our Universal Backdoor is gradually learned throughout the training process, so any early stopping procedure that would mitigate our backdoor would also significantly reduce the model’s clean accuracy.

4.3 SCALING

Table 2: Attack success rate of the backdoor on larger datasets (%), using $p = 12\,000$.

Poison Attack	ImageNet-2K	ImageNet-4K	ImageNet-6K
Universal Backdoor	99.73	91.75	47.31
Baseline	99.98	0.03	0.02

In this experiment, we measure our backdoor’s ability to scale to larger datasets. We fix the number of poisons the attacker injects into each dataset at $p = 12\,000$ across all runs. As larger image classification datasets naturally contain more classes and samples, so do our datasets (Deng et al., 2009; Kuznetsova et al., 2018). Table 2 summarizes the results on our backdoor compared to a baseline on the ImageNet-2K, ImageNet-4K, and ImageNet-6K image classification tasks. We find that the trigger patterns of the baseline do not effectively scale to larger image classification datasets. Although the baseline backdoor has near-perfect ASR on the ImageNet-2K dataset, it has negligible ASR on both the ImageNet-4K and ImageNet-6K datasets. This is because of the all-or-nothing attack success behavior observed in Section 4.2. In contrast, our Universal Backdoor can scale to image classification tasks containing more classes and samples. Our Universal backdoor achieves above 90% ASR on the ImageNet-4K task and 47.31% ASR on the largest dataset, ImageNet-6K.

4.4 MEASURING INTER-CLASS POISON TRANSFERABILITY

To measure the inter-class transferability of poisoning, we examine how increasing the number of poisons in one set of classes increases attack success on a disjoint set of classes in the dataset. We divide the classes into the observed set **B** and the variation set **A**. **B** contains 10% of the classes in the dataset (100 classes), while **A** contains the remaining 90% of classes (900 classes). We use the ImageNet-1K dataset and a patch trigger for our backdoor. We poison exactly one sample in each class in **B**. In Figure 4, we ablate over the total number of poisons in the dataset, distributing all poisons except for the 100 poisons in **B** evenly in classes in **A**.

We find that by poisoning a class with a single sample, our Universal Backdoor can achieve a successful attack on a class if sufficient poisoning is achieved elsewhere in the dataset. Increasing the number of poison samples in **A** improved the backdoor’s ASR on classes in **B** from negligible to over 70%. *Therefore, we find that protecting the integrity of a single class requires protecting the integrity of the entire dataset.* Our Universal Backdoor shows that every sample, even if they are associated with an insensitive class label, can be used by an attacker as part of an extremely poison-efficient backdoor attack on a small subset of high-value classes. We provide further evidence for this in Appendix A.2, where we show that **A** can contain significantly fewer than 900 classes while preserving the strength of inter-class poison transferability on **B**. The baseline method does not demonstrate any inter-class transferability, as increasing the poisoning in **A** does not increase the attack success rate on **B**.

4.5 ROBUSTNESS AGAINST DEFENSES

We evaluate the robustness of our poisoning model against four state-of-the-art defenses: fine-tuning, fine-pruning (Liu et al., 2018), neural attention distillation (Li et al., 2021), and neural cleanse (Wang et al., 2019). We use a ResNet-18 model trained on the ImageNet-1k dataset for all robustness evaluations. We use patch triggers for both our method and the baseline. For all defenses, we use hyper-parameters optimized for removing a BadNets backdoor (Gu et al., 2017) on ImageNet-1K as proposed by Lukas & Kerschbaum (2023). Defenses requiring clean data are given 1% of the clean dataset, approximately 12,800 clean samples. We limit the degradation of the model’s clean accuracy, halting any defense that degrades the model’s clean accuracy by more than 2%. Table 3 summarizes the changes in ASR after applying each defense. As in Lukas &

Kerschbaum (2023), we find that backdoored models trained on ImageNet-1K are robust against defenses. A complete table of defense parameters can be found in Appendix A.

Table 3: The robustness of our universal backdoor against a naïve baseline, measured by the attack success rate (ASR). ▼ denotes ASR lost after applying defense. Only backdoors above 5% ASR were evaluated. Backdoors that were not evaluated are marked with N/A.

Defense	Poison Samples (p)	Universal Backdoor (ASR)	Baseline (ASR)
Fine-Tuning	2 000	70.3%	N/A
	5 000	94.5%	N/A
	8 000	95.5%	99.5% ▼ 0.5
Fine-Pruning	2 000	73.5%	N/A
	5 000	95.2%	N/A
	8 000	95.6%	99.9% ▼ 0.1
Neural Cleanse	2 000	70.1%	N/A
	5 000	95.1%	N/A
	8 000	95.4%	98.0% ▼ 2.0
NAD	2 000	73.9%	N/A
	5 000	94.6%	N/A
	8 000	95.3%	99.9% ▼ 0.1

Fine-tuning. This defense fine-tunes the dataset on a small validated subset of the training dataset. We fine-tune the model using the SGD with a learning rate of 0.0005 and a momentum of 0.9.

Fine-pruning. As in Liu et al. (2018), we prune the last convolutional layer of the model. We find that the pruning rate in Lukas & Kerschbaum (2023) is too high and degrades the clean accuracy of the model more than the 2% cutoff. We set the pruning rate to 0.1%, which is the maximum pruning rate that prevents the defense from degrading the model below the accuracy cutoff.

Neural Cleanse. Neural Cleanse (Wang et al., 2019) uses outlier detection to decide which candidate trigger is most likely the result of poisoning. This candidate trigger is then used to remove the backdoor in the model. As our Universal Backdoor targets every class and has a unique trigger for each class, class-wise anomaly detection is poorly suited for removing our backdoor.

Neural Attention Distillation. We train a teacher model for 1 000 steps using SGD. We then align the backdoored model with the teacher for 8000 steps, using SGD with a learning rate of 0.0005. We use a power term of 2 for the attention distillation loss, as recommended in Li et al. (2021).

4.6 MEASURING THE CLEAN DATA TRADE-OFF

There is a known trade-off between the availability of clean data and the effectiveness of defenses (Li et al., 2021). Figure 5 measures the proportion of the clean dataset required to remove the universal backdoor with fine-tuning without degrading the model below the 2% cutoff. For this experiment, we use a ResNet-18 model backdoored using 2 000 poison samples on the ImageNet-1K dataset. Due to the higher availability of clean data, we find that a higher learning rate of 0.001 and a weight decay of 0.001 are appropriate.

Data poisoning defenses for backdoored models trained on web-scale datasets must be effective with a validated dataset that is a small portion of the training dataset due to the cost of manually validating samples. Validating a 1% portion of our web-scale ImageNet-6k dataset would require manually inspecting over 78 000 samples, a task larger than inspecting the CIFAR-100 or GTSRB datasets in

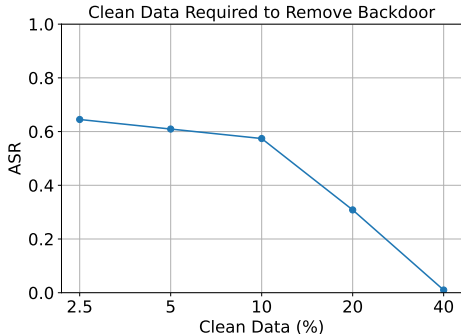


Figure 5: Clean data as a percentage of the training dataset size required to remove our Universal Backdoor.

their entirety (Krizhevsky, 2009; Stallkamp et al., 2011). We find that approximately 40% (512 466 samples) of the clean dataset is required to completely remove our Universal Backdoor, which is more data than most victims can manually validate.

5 DISCUSSION AND RELATED WORK

Attacking web-scale datasets. Carlini et al. (2023) demonstrate two realistic ways an attacker could poison a web-scale dataset: domain hijacking and snapshot poisoning. They show that more than 0.15% of the samples in these online datasets could be poisoned by an attacker. However, existing many-to-one poison label attacks cannot exploit these vulnerabilities, as they require compromising many samples *in a single class* (Gu et al., 2017; Qi et al., 2022; Chen et al., 2017). As web-scale datasets contain thousands of classes (Deng et al., 2009; Kuznetsova et al., 2018), it is improbable that any one class would have enough compromised samples for a many-to-one poison label attack. By leveraging inter-class poison transferability, our backdoor can utilize compromised samples outside a class the attacker is attempting to misclassify into.

Scaling to larger datasets. The largest dataset we evaluate is our ImageNet-6K dataset, which consists of 6 000 classes and 7 804 447 samples. We created a Universal Backdoor in a model trained on this dataset while poisoning only 0.15%. As our backdoor effectively scales to datasets containing more classes and samples, we expect a smaller proportion of poison samples to be required to backdoor models trained on larger datasets, like LAION-5B (Schuhmann et al., 2022).

Alternative methodology for targeting multiple classes at inference time. Although we are the first to study how to target every class in the data poisoning setting, other types of attacks, like *adversarial examples*, can be used to target specific classes at inference time (Wu et al., 2023; Goodfellow et al., 2015). Through direct optimization on an input, the attacker finds an adversarial perturbation that acts as a trigger; adding it to the input causes a misclassification. Defenses against adversarial examples seek to make models robust against adversarial perturbations (Cohen et al., 2019; Geiping et al., 2021). However, as data poisoning backdoors utilize triggers that are not adversarial perturbations, these defenses are ineffective at mitigating data poisoning backdoors.

Limitations. We focus on patch and blend triggers that are visible modifications to the image and hence could be detected by a data sanitation defense. Our attacks are agnostic to the trigger; even if a specific trigger could be reliably detected, universal backdoors remain a threat because the attacker could have used a different trigger. Koh et al. (2022) demonstrate that no detection has been shown effective against any trigger. However, evading data sanitation comes at a cost for the attacker: Less detectable triggers are less effective at equal numbers. Hence, the attacker must inject more to create an equally effective backdoor (Frederickson et al., 2018). We point to Appendix A.3 showing that our attacks still remain difficult to detect using STRIP (Gao et al., 2019) due to the high false positive rate. We focus on the feasibility of universal attacks and do not study the detectability-effectiveness trade-off of triggers with our attacks. Moreover, we focus on poisoning models from scratch, as opposed to poisoning pre-trained models that are fine-tuned. More research is needed to analyze the effectiveness of our attacks against large pre-trained models like ViT and CLIP (Dosovitskiy et al., 2021; Radford et al., 2021) that are fine-tuned on poisoned data. Finally, we assume that the attacker can access similarly accurate surrogate classifiers to generate latent encodings for our attacks.

6 CONCLUSION

We introduce Universal Backdoors, a data poisoning backdoor that targets every class. We establish that our backdoor requires significantly fewer poison samples than independently attacking each class and can effectively attack web-scale datasets. We also demonstrate how compromised samples in uncritical classes can be used to reinforce poisoning attacks against other more sensitive classes. Our work exemplifies the need for practitioners who train models on untrusted data sources to protect the whole dataset, not individual classes, from data poisoning. Finally, we show that existing defenses are ineffective at defending against Universal Backdoors, indicating the need for new defenses designed to remove backdoors that target many classes.

REFERENCES

- Rasmus Adler, Mohammed Naveed Akram, Pascal Bauer, Patrik Feth, Pascal Gerber, Andreas Jedlitschka, Lisa Jöckel, Michael Kläs, and Daniel Schneider. Hardening of artificial neural networks for use in safety-critical applications - A mapping study. *CoRR*, abs/1909.03036, 2019. URL <http://arxiv.org/abs/1909.03036>.
- Naif Alkhunaizi, Dmitry Kamzolov, Martin Takáč, and Karthik Nandakumar. Suppressing poisoning attacks on federated learning for medical imaging, 2022.
- Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. Poisoning web-scale training datasets is practical, 2023.
- Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *CoRR*, abs/1712.05526, 2017. URL <http://arxiv.org/abs/1712.05526>.
- Antonio Emanuele Cinà, Kathrin Grosse, Ambra Demontis, Sebastiano Vascon, Werner Zellinger, Bernhard A. Moser, Alina Oprea, Battista Biggio, Marcello Pelillo, and Fabio Roli. Wild patterns reloaded: A survey of machine learning security against training data poisoning. *ACM Comput. Surv.*, 55(13s), jul 2023. ISSN 0360-0300. doi: 10.1145/3585385. URL <https://doi.org/10.1145/3585385>.
- Jeremy Cohen, Elan Rosenfeld, and J. Zico Kolter. Certified adversarial robustness via randomized smoothing. *CoRR*, abs/1902.02918, 2019. URL <http://arxiv.org/abs/1902.02918>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- Bao Gia Doan, Ehsan Abbasnejad, and Damith Chinthana Ranasinghe. Deepcleanse: Input sanitization framework against trojan attacks on deep neural network systems. *CoRR*, abs/1908.03369, 2019. URL <http://arxiv.org/abs/1908.03369>.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- R. A. FISHER. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936. doi: <https://doi.org/10.1111/j.1469-1809.1936.tb02137.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-1809.1936.tb02137.x>.
- Christopher Frederickson, Michael Moore, Glenn Dawson, and Robi Polikar. Attack strength vs. detectability dilemma in adversarial machine learning. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2018. doi: 10.1109/IJCNN.2018.8489495.
- Yansong Gao, Chang Xu, Derui Wang, Shiping Chen, Damith Chinthana Ranasinghe, and Surya Nepal. STRIP: A defence against trojan attacks on deep neural networks. *CoRR*, abs/1902.06531, 2019. URL <http://arxiv.org/abs/1902.06531>.
- Jonas Geiping, Liam Fowl, Gowthami Somepalli, Micah Goldblum, Michael Moeller, and Tom Goldstein. What doesn’t kill you makes you robust(er): Adversarial training against poisons and backdoors. *CoRR*, abs/2102.13624, 2021. URL <https://arxiv.org/abs/2102.13624>.
- Micah Goldblum, Dimitris Tsipras, Chulin Xie, Xinyun Chen, Avi Schwarzschild, Dawn Song, Aleksander Madry, Bo Li, and Tom Goldstein. Dataset security for machine learning: Data poisoning, backdoor attacks, and defenses. *CoRR*, abs/2012.10544, 2020. URL <https://arxiv.org/abs/2012.10544>.

- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples, 2015.
- Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *CoRR*, abs/1708.06733, 2017. URL <http://arxiv.org/abs/1708.06733>.
- Jonathan Hayase, Weihao Kong, Raghav Somani, and Sewoong Oh. SPECTRE: defending against backdoor attacks using robust statistics. *CoRR*, abs/2104.11315, 2021. URL <https://arxiv.org/abs/2104.11315>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.
- Sanghyun Hong, Nicholas Carlini, and Alexey Kurakin. Handcrafted backdoors in deep neural networks. *CoRR*, abs/2106.04690, 2021. URL <https://arxiv.org/abs/2106.04690>.
- Alaa Khaddaj, Guillaume Leclerc, Aleksandar Makelov, Kristian Georgiev, Hadi Salman, Andrew Ilyas, and Aleksander Madry. Rethinking backdoor attacks, 2023.
- Pang Wei Koh, Jacob Steinhardt, and Percy Liang. Stronger data poisoning attacks break data sanitization defenses. *Machine learning*, 111(1):1–47, 2022. ISSN 0885-6125.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper R. R. Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and Vittorio Ferrari. The open images dataset V4: unified image classification, object detection, and visual relationship detection at scale. *CoRR*, abs/1811.00982, 2018. URL <http://arxiv.org/abs/1811.00982>.
- Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. Neural attention distillation: Erasing backdoor triggers from deep neural networks. *CoRR*, abs/2101.05930, 2021. URL <https://arxiv.org/abs/2101.05930>.
- Cong Liao, Haoti Zhong, Anna Cinzia Squicciarini, Sencun Zhu, and David J. Miller. Backdoor embedding in convolutional neural network models via invisible perturbation. *CoRR*, abs/1808.10307, 2018. URL <http://arxiv.org/abs/1808.10307>.
- Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Fine-pruning: Defending against backdoor-attacks on deep neural networks. *CoRR*, abs/1805.12185, 2018. URL <http://arxiv.org/abs/1805.12185>.
- Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. *CoRR*, abs/2007.02343, 2020. URL <https://arxiv.org/abs/2007.02343>.
- Nils Lukas and Florian Kerschbaum. Pick your poison: Undetectability versus robustness in data poisoning attacks, 2023.
- Tuan Anh Nguyen and Anh Tuan Tran. Wanet - imperceptible warping-based backdoor attack. *CoRR*, abs/2102.10369, 2021. URL <https://arxiv.org/abs/2102.10369>.
- Naman Patel, Prashanth Krishnamurthy, Siddharth Garg, and Farshad Khorrami. Bait and switch: Online training data poisoning of autonomous driving systems. *CoRR*, abs/2011.04065, 2020. URL <https://arxiv.org/abs/2011.04065>.
- Xiangyu Qi, Tinghao Xie, Yiming Li, Saeed Mahloujifar, and Prateek Mittal. Circumventing backdoor defenses that are based on latent separability. *arXiv preprint arXiv:2205.13613*, 2022.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021. URL <https://arxiv.org/abs/2103.00020>.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *CoRR*, abs/1409.0575, 2014. URL <http://arxiv.org/abs/1409.0575>.

Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models, 2022.

Ali Shafahi, W. Ronny Huang, Mahyar Najibi, Octavian Suci, Christoph Studer, Tudor Dumitras, and Tom Goldstein. Poison frogs! targeted clean-label poisoning attacks on neural networks. *CoRR*, abs/1804.00792, 2018. URL <http://arxiv.org/abs/1804.00792>.

Johannes Stalldkamp, Marc Schlipfing, Jan Salmen, and Christian Igel. The german traffic sign recognition benchmark: A multi-class classification competition. In *The 2011 International Joint Conference on Neural Networks*, pp. 1453–1460, 2011. doi: 10.1109/IJCNN.2011.6033395.

Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y. Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 707–723, 2019. doi: 10.1109/SP.2019.00031.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.

Baoyuan Wu, Li Liu, Zihao Zhu, Qingshan Liu, Zhaofeng He, and Siwei Lyu. Adversarial machine learning: A systematic survey of backdoor attack, weight attack and adversarial example, 2023.

A APPENDIX

Table 4 contains a complete summary of all the parameters used to evaluate defenses against our backdoor in Section 4.5. All defense parameters are adapted from Lukas & Kerschbaum (2023), where they were optimized against a BadNets (Gu et al., 2017) patch trigger. When hyperparameter tuning for fine-tuning and fine-pruning defenses, we find no significant improvements over the settings described in Lukas & Kerschbaum (2023). We reduce the fine-pruning rate in Fine-pruning, as we find it degrades the model’s clean accuracy below our 2% cutoff.

As shown by Figure 6a, a linear trade-off exists between the effectiveness of defenses and the allowed clean accuracy cutoff. If the defender allows for more clean accuracy degradation, the effectiveness of the backdoor can be further reduced. This does not apply to all defenses, as defenses like neural cleanse (Wang et al., 2019) do not significantly reduce clean accuracy.

A.1 ANALYSIS OF CLEAN ACCURACY

If a backdoor attack degrades the clean accuracy of a model, then the validation set is sufficient for the victim to recognize the presence of a backdoor (Gu et al., 2017). Therefore, a model trained on the poisoned set should achieve the same clean accuracy as one trained on a comparable clean dataset. We find that our backdoored models have the same clean accuracy across all runs as a model trained on entirely clean data. We train a clean ResNet-18 model on ImageNet-1k (Russakovsky et al., 2014), which achieves 68.49% top-1 accuracy on the validation set. Table 5 shows the clean accuracy of backdoored models on the ImageNet-1k dataset.

Table 4: Defense Parameters on ImageNet-1K from Lukas & Kerschbaum (2023).

Neural Attention Distillation	
n steps / N	8,000
opt	sgd
lr / α	5e-4
teacher steps	1,000
power / p	2
at lambda / λ_{at}	1,000
weight decay	0
batch size	128

Neural Cleanse	
n steps / N	3,000
opt	sgd
lr / α	5e-4
steps per class / N1	200
norm lambda / λ_N	1e-5
weight decay	0
batch size	128

Fine-Tuning	
n steps / N	5,000
opt	sgd
lr / α	5e-4
weight decay	0.001
batch size	128

Fine-Pruning	
n steps / N	5,000
opt	sgd
lr / α	5e-4
prune rate / ρ	10%
sampled batches	10
weight decay	0
batch size	128

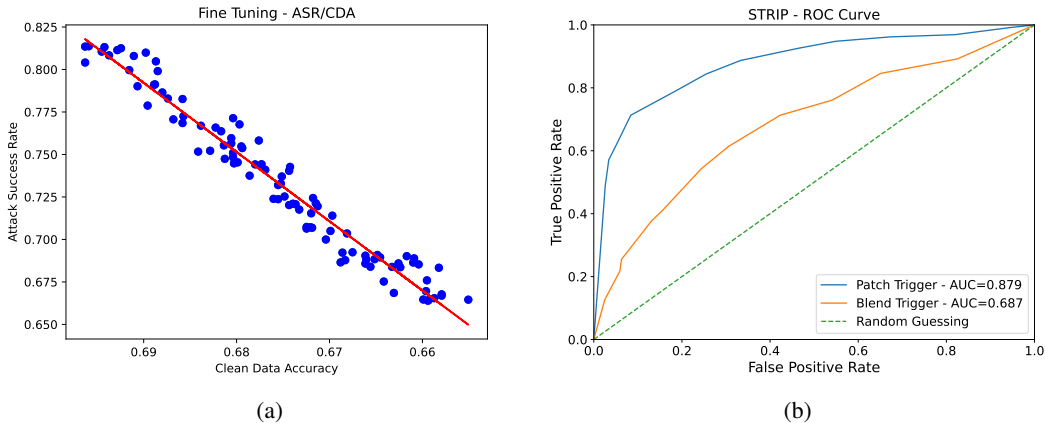


Figure 6: (6a) Trade-off between attack success rate and clean data accuracy when fine-tuning a backdoored model. (6b) ROC curve of our Universal Backdoor with patch and blend triggers (see Figure 2) when applying the STRIP (Gao et al., 2019) defense.

A.2 INTER-CLASS POISON TRANSFERABILITY WITH SMALL VARIATION SETS

Table 6 shows that even if the number of classes in the variation set **A** is reduced to only 10% of classes in \mathcal{Y} , inter-class poison transferability maintains its effect on the observed set **B**. This results in an otherwise unsuccessful attack on classes in **B**, achieving a success rate of 67.72%. Therefore, if the attacker can strongly poison a small set of classes in the dataset, attacking other classes in the model can easily be accomplished, as inter-class poison transferability remains strong. To protect even a tiny subset of high-value classes, the victim must maintain the integrity of every class within their dataset.

A.3 DATA SANITATION DEFENSES

Several data sanitization defenses are also poorly suited to Universal Backdoors. SPECTRE (Hayase et al., 2021) only removes samples from a single class by design, and therefore could remove at most 0.1% of our Universal Backdoor’s poisoned samples on ImageNets-1K. STRIP (Gao et al., 2019) struggles to detect our trigger, resulting in a high false positive rate, as shown in Figure 6b. The area

Table 5: Clean accuracy of backdoored models on ImageNet-1k dataset.

Poison Samples (p)	Poison %	Patch		Blend	
		Ours	Baseline	Ours	Baseline
2000	0.16	68.94%	68.94%	68.51%	69.43%
5000	0.39	68.92%	68.89%	68.77%	68.66%
8000	0.62	68.91%	69.43%	69.78%	69.22%

Table 6: Effect of the number of classes in the variation set **A** on attack success on the observed set **B**. All experiments use 4 600 poison samples.

Percentage of classes in A	ASR on classes in B
90%	72.77%
60%	70.45%
30%	71.78%
10%	67.72%

under the ROC curves are 0.879 and 0.687 for the patch and blend triggers, respectively. It may be difficult for defenders to detect both triggers for large datasets (1 million samples or more) due to the detection’s high FPR. Considering a maximum tolerable FPR of 10%, the defender misses 39% of the patch trigger samples and 68% of the blended triggers.

A.4 CLASS-WISE ATTACK SUCCESS METRICS

Our method does not achieve even attack success across all classes in the dataset. Table 7 shows statistics of our Universal Backdoor’s success rate across classes in ImageNet-1K. We find that some classes are more challenging to achieve a successful attack against our backdoor. This differs from the baseline, as the baseline either performs near-perfectly or not at all.

Table 7: ASR metrics across classes in ImageNet-1K

Poison Samples (p)	Min	Max (%)	Mean	Median
2000	0	100%	81.0%	98.0%
5000	0	100%	95.4%	100%