

BppAttack: Stealthy and Efficient Trojan Attacks against Deep Neural Networks via Image Quantization and Contrastive Adversarial Learning

Zhenting Wang, Juan Zhai, Shiqing Ma

Department of Computer Science, Rutgers University

{zhenting.wang, juan.zhai, sm2283}@rutgers.edu

Abstract

Deep neural networks are vulnerable to Trojan attacks. Existing attacks use visible patterns (e.g., a patch or image transformations) as triggers, which are vulnerable to human inspection. In this paper, we propose stealthy and efficient Trojan attacks, BPPATTACK. Based on existing biology literature on human visual systems, we propose to use image quantization and dithering as the Trojan trigger, making imperceptible changes. It is a stealthy and efficient attack without training auxiliary models. Due to the small changes made to images, it is hard to inject such triggers during training. To alleviate this problem, we propose a contrastive learning based approach that leverages adversarial attacks to generate negative sample pairs so that the learned trigger is precise and accurate. The proposed method achieves high attack success rates on four benchmark datasets, including MNIST, CIFAR-10, GTSRB, and CelebA. It also effectively bypasses existing Trojan defenses and human inspection. Our code can be found in <https://github.com/RU-System-Software-and-Security/BppAttack>.

1. Introduction

Deep Neural Networks (DNNs) have achieved superior performance in many computer vision tasks [8, 21, 52]. Recent studies show that DNNs are vulnerable to adversarial attacks such as adversarial examples [17, 45], membership inference attacks [54, 58], model stealing [50, 62], etc. In this paper, we focus on Trojan attacks [10, 13, 18, 35, 38, 41, 53]. The adversary injects a secret Trojan behavior during training, which can be activated at runtime by stamping a Trojan trigger to the image. Such triggers can be image patches [18], watermarks [41], image filters [1, 40] and even learned image transformation models [10, 13, 37].

Trojan attacks [18] are severe threats to the trustworthiness of DNN models. Liu et al. [41] demonstrates the possibility of attacking face recognition, speech recognition,

and autonomous driving systems. Such attacks are generally feasible in most training scenarios, including federated learning, unsupervised learning, and so on [4, 29, 68]. With the deployment of DNN based computer vision models, it is a critical challenge in our community.

Existing Work: Most existing Trojan attacks leverage input patterns as triggers. For example, BadNets [18] uses a yellow pad as its trigger. Recent works [40] try to leverage image filters as triggers, which are input dependent and dynamic, making them hard to detect. To further improve the quality of Trojan triggers, Doan et al. [13] train an auxiliary image transformation model and use the transformation function as its trigger. Other works have adopted similar ideas [10, 37].

One problem of existing attacks is that they are vulnerable to human inspections. Once a set of attack inputs are found, it is not difficult to identify the trigger or train a model to simulate the trigger. There are also online detection methods to identify such attack samples, such as STRIP [16]. Even for trained transformations as triggers, it is hard for them to guarantee that the generated images have imperceptible changes. This is because it is hard to formulate human visual systems as a mathematical function, which makes it hard to optimize. Due to the relatively large changes in inputs and limitations of existing poisoning methods, it is also possible for reverse engineering based defense methods [7, 40, 66] to recover part of the trigger and identify if a model has a Trojan. Moreover, recent works on generating high-quality triggers typically leverage trained auxiliary models, which is time-consuming and inefficient.

Our Work: In this paper, we propose one new attack, BPPATTACK. Based on existing literature on human visual systems, we identify that humans inspectors are insensitive to small changes of color depth. In this attack, we try to exploit vulnerabilities in human visual systems. Thus, we propose to reduce the bit-per-pixel (BPP) to conduct an imperceptible attack, which can also bypass existing defenses mainly because of the small changes made to the input domain. We achieve our goal by performing a deterministic yet input-dependent image transformation, i.e., image quan-

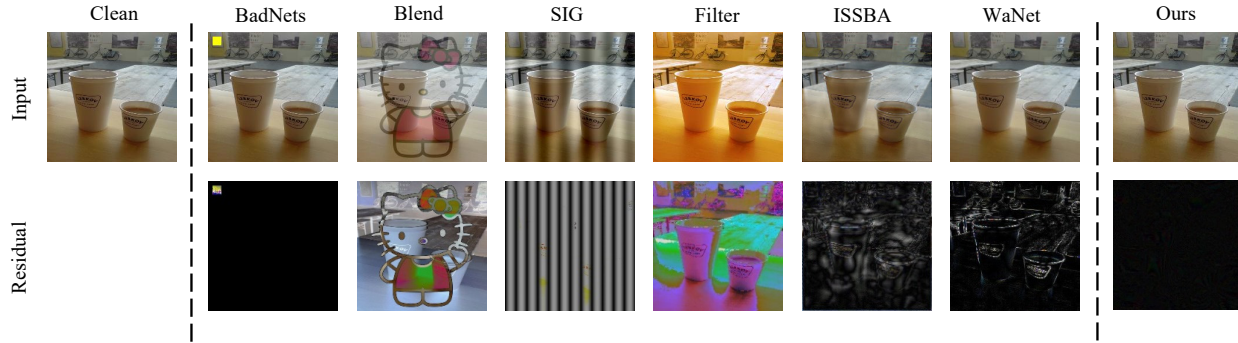


Fig. 1. Comparison of examples generated by different Trojan attacks (i.e., BadNets [18], blending-based attack [9], SIG [2], filter-based attack [1, 10], ISSBA [37] and WaNet [48]). For each attack, we show the Trojan sample (top) and the magnified ($\times 5$) residual (bottom).

tization and dithering. Due to the small-sized transformation as triggers, it is more challenging to train the model and inject the trigger. To overcome this issue, we also propose a contrastive learning and adversarial training method based approach to training on the poisoned dataset. By doing so, we do not require training auxiliary models, making the attack fast and also input dependent. Moreover, our attack exploits vulnerabilities in human visual systems, making it human imperceptible when the attack settings are properly set. Fig. 1 shows the comparison of our attack and existing attacks using attack samples and residuals.

Our contributions are summered as follows:

- We propose a new attack that exploits human visual systems. We design an effective and efficient attack that leverages image quantization and dithering. It performs deterministic input-dependent image transformations, making it fast and dynamic. We also propose a contrastive learning and adversarial training based approach to enhance the poisoning process to inject such human imperceptible triggers.
- We evaluate our prototype on four different datasets and seven different network architectures. Results show that our attack achieves a 99.92% attack success rate on average. Our human study also confirms that it is 1.60 times better than the SOTA approaches when facing human inspections. Results on existing defenses also confirm that our attack can effectively bypass various types of SOTA defenses.

2. Background

2.1. Trojan Attacks

Trojan models behave normally for benign inputs but have malicious behaviors (i.e., outputting a particular label) on inputs stamped with the Trojan trigger. One limitation of existing Trojan attacks is that most of them are perceptible

to human inspectors. Many Trojan attacks [9, 18, 41] use predefined patches or watermarks as Trojan triggers. Re-fool [42] exploits physical reflection as Trojan trigger. Trojan attacks can also happen in the feature space. For example, Liu et al. [40] demonstrates attackers can use Instagram filters as triggers to perform Trojan attacks. DFST [10] utilizes CycleGAN [73] to inject Trojans in deep features space. All these triggers are obvious for human inspection. Recently, WaNet [48] proposed attacks using the image warping technique as triggers. Although it is more stealthy than previous works, the warping effects it leverages are still perceptible. Another problem of existing attacks is that they typically use fixed patterns as trigger patterns, which means different samples share the same trigger pattern. This property makes such Trojan attacks detectable by existing defenses [20, 40, 43]. Nguyen et al. [49] proposes input dependent triggers. This attack brings large pixel-level perturbations, sacrificing stealthiness. Recently, Li et al. [37] and Doan et al. [13] proposed new attacks that are not only imperceptible but also input dependent. The idea is to generate triggers by trained auto-encoders. While such methods achieve stealthiness, they are model-dependent and time-consuming.

2.2. Existing Defense

There has been a series of ways to defend Trojans. One of them is training time defense, which aims at removing Trojans before/during training. Chen et al. [6] and Tran et al. [61] detect the malicious samples before training. Wang et al. [67] removes Trojans in training by formalizing the trigger in input space. Similarly, poison suppression [14, 24] depresses the malicious effectiveness of poisons in training. These approaches target poisoning-based Trojan attacks but ignore supply chain Trojan attacks. The second method is reverse engineering. Neural Cleanse [66], DeepInspect [7], K-arm [57] and ABS [40] use reconstructed triggers to perform detection. These methods work on local patched trig-

gers but fail to generalize to different trigger types, e.g., input-aware triggers [49]. Another approach is to remove Trojans in infected models [36, 39, 60, 72], such as finetuning or pruning, which in cases, can lower benign accuracy as well. Some other existing works try to leverage online detection as a defense. STRIP [16] detect Trojan samples by analyzing the sensitivity of samples on strong perturbation. Februus [12] and SentiNet [11] leverage GradCAM [56] to detect if model predictions are localized and leverage it as a hint for triggers. Such methods fail when triggers are not localized, e.g., filter triggers.

3. Method

In this section, we introduce BPPATTACK, a Trojan attack that is invisible to human inspection, input-dependent yet requires no auxiliary model training. We first describe the threat model (§3.1), and then present the foundation and details of the attack process (§3.2 and §3.3, respectively).

3.1. Threat Model

Adversarial scope and goal. The adversary aims to produce a Trojan model. Eq. 1 shows the formal definition. \mathcal{M}_θ is Trojan model, T is a Trojan transformation function and η is the target label function. Input-targeted labels can be: (1) *all-to-one*: the attacker select a constant label c as output label (i.e., $\eta(y) = c$). (2) *all-to-all*: the target label is the next label of the true label (i.e., $\eta(y) = y + 1$).

$$\mathcal{M}_\theta(\mathbf{x}) = y, \quad \mathcal{M}_\theta(T(\mathbf{x})) = \eta(y) \quad (1)$$

Compared with previous Trojan attacks, we aim to provide the following attack properties:

- **Effective:** We want the model to have a high attack success rate (ASR) while maintaining high benign accuracy at the same time. This effective goal is the basic requirement of Trojan attacks as defined in Eq. 1.
- **Imperceptible:** Many Trojan triggers are vulnerable to human inspection, which is not robust. We want to have a human imperceptible Trojan trigger. Traditionally, this is done by defining a distance function \mathcal{V} to measure the visual similarity of two samples. As such, the goal is to find a trigger that is smaller than a threshold, $\mathcal{V}(T(\mathbf{x}), \mathbf{x}) < t$, where t is the threshold. Existing works use L_p distance or SSIM scores, which do not align with the human visual systems [51]. In this paper, we tackle this problem by starting from existing studies on the human visual system and propose an attack that is human imperceptible.
- **Input-dependent:** Fixed trigger patterns are easier to detect [16, 66] and in most cases, human visible. Thus, input-dependent triggers are natural inheriting from

the human imperceptible requirement. We want to have an image perturbation function $R(\mathbf{x}) = T(\mathbf{x}) - \mathbf{x}$ that satisfies

$$\begin{cases} \mathcal{M}_\theta(\mathbf{x} + R(\mathbf{x})) = \eta(y) \\ \mathcal{M}_\theta(\mathbf{x} + R(\mathbf{x}')) = y \end{cases} \quad \text{where } \mathbf{x}' \neq \mathbf{x} \quad (2)$$

- **No auxiliary training:** Many existing works try to realize input-dependent attacks by utilizing an auxiliary model, e.g., DFST [10] uses CycleGAN. Such attacks are unstable because their effects depend on the training of the auxiliary models. Moreover, it has high computation overhead. In contrast, we try to design an efficient Trojan attack without auxiliary models.

Adversary capabilities. Following existing attacks [13, 48], we assume the adversary has full control of datasets, training process, and model implementation. The adversary injects the Trojan by poisoning the dataset.

3.2. Human Imperceptible Theory

Our idea of generating human imperceptible triggers is from the biology study that human visual systems are insensitive to color bit depth change. Nadenau et al. [46] and many existing literatures [28, 30, 47, 71] supported this observation. Image color quantization [3, 5, 23, 65] is a process that reduces the number of distinct colors used in an image with the intention to produce human imperceptible changes. To remove the unnaturalness introduced by color bit change, dithering [15, 26, 63] can improve its quality.

3.3. BPPATTACK

To achieve the aforementioned objectives, we design a novel image color quantization based Trojan attack. Spectrally, we leverage image color quantization and dithering to generate high-quality attack triggers and poisoning samples and then propose a contrastive learning and adversarial training-based method to inject the Trojan.

Image quantization. The first step of BPPATTACK is to perform image quantization, which contains two steps. First, we squeeze the original color palette (m bits for each pixel on each channel) of the image into a smaller color palette (d bits) by reducing the color depth. For each pixel, we use the nearest pixel value in the squeezed d -bits space to replace the original value. The squeezing function T is defined in Eq. 3, where *round* represents the integer rounding function:

$$T(\mathbf{x}) = \frac{\text{round}\left(\frac{\mathbf{x}}{2^{m-1}} * (2^d - 1)\right)}{2^d - 1} * (2^m - 1) \quad (3)$$

This is the main algorithm to generate Trojan triggers and has a few benefits. First, it is a simple and deterministic

function with good stability and generalizability, and we do not need to train any auxiliary models such as auto-encoders and U-Nets. Second, as pointed out by existing work [46, 70], large color depths are not necessary for representing images, which means the squeezed image can have high visual similarity to the original image. While being human imperceptible, such digital value changes can be captured by ML models and used as a trigger.

Dithering. Image quantization potentially can cause unnatural regions, especially when the bit reduction is high. To increase the stealthiness of BPPATTACK, we utilize image dithering techniques to remove the noticeable artifacts by leveraging the existing colors of the artifacts. Image dithering techniques are designed to create the illusion of color depth when color palette of image is limited. Specifically, we use Floyd–Steinberg dithering [15] and nearest-value color quantization combined with dithering as Trojan transformation. Details are presented in Algorithm 1. Function *quantize* implements Eq. 3. Floyd–Steinberg dithering achieves its goal by error diffusion, and line 4 calculates the error. After that, it adds residual quantization errors of a pixel onto its neighbors and spreads the debt out based on a predefined distribution. Lines 5 to 9 implement this idea.

Algorithm 1 Quantization with Floyd-Steinberg Dithering

Input: Image I , Diffusion Distribution $[a_1, a_2, a_3, a_4]$

Output: Quantized Image

```

1: function PROCESS( $I$ )
2:   for  $x$  from right to left do
3:     for  $y$  from top to bottom do
4:        $error = quantize(I[x][y]) - I[x][y]$ 
5:        $I[x][y] = I[x][y] + error$ 
6:        $I[x + 1][y] = I[x][y] + error * a_1$ 
7:        $I[x + 1][y + 1] = I[x][y] + error * a_2$ 
8:        $I[x][y + 1] = I[x][y] + error * a_3$ 
9:        $I[x - 1][y + 1] = I[x][y] + error * a_4$ 

```

Contrastive Adversarial Training. As shown in Fig. 1, image quantization based attack triggers is very close to original images. On the one hand, this makes it hard to detect. On the other hand, it makes training more difficult, mainly because of the small perturbations. Existing poisoning techniques tend to use the original cross-entropy (CE) loss to train the Trojan model on benign and poisoning samples. Due to the tiny perturbation introduced by image quantization, it is hard to converge when using the CE loss. Moreover, existing training procedure leads to inaccurate and imprecise triggers. As a result, reverse engineering can identify if a model has a Trojan by finding part of the trigger. As a consequence, they are not robust attacks. To overcome this challenge, we leverage contrastive supervised learning and adversarial training.

The whole training framework follows the contrastive

learning framework, and we leverage the same loss function as described in existing work [31]. The key difference of our attack from existing contrastive learning is that in addition to existing negative sample generation methods, we also leverage adversarial example generation methods. Specifically, we use the PGD attack to generate adversarial examples which flip the label of input from its original one to the target label to simulate the effects of our attack. Then, we leverage them in training as negative examples. Intuitively, this means we exclude such perturbations features as important features for the model to learn so that it can focus on the injected trigger that is image quantization and dithering described before. Note that the PGD attack is an optimization-based method and does not require training auxiliary models.

4. Experiments and Results

In this section, we evaluate BPPATTACK from different perspectives. We first present the experiment setup, including datasets and other settings in §4.1. In §4.2, we show the effectiveness. Then, we investigate the stealthiness of BPPATTACK by performing a human Inspection test (§4.3). Furthermore, we evaluate BPPATTACK’s resistance to existing defenses in §4.4. We also conduct an ablation study of BPPATTACK in §4.5. In all experiments, the default bit depth is $d = 5$.

4.1. Experiment Setup

Datasets. We evaluate BPPATTACK on four datasets: MNIST, CIFAR-10, GTSRB and CelebA. These datasets are regularly used in backdoor-related researches [12, 16, 18, 39–41, 48, 66]. Details of these datasets are in Table 1. MNIST [34] is used for hand-written digits recognition. GTSRB [59] is built for classifying different traffic signs. CIFAR-10 [33] is a classification benchmark. CelebA [44] is a large-scale face attributes classification dataset. Note that CelebA has 40 independent binary attributes, where most attributes are unbalanced. To make it suitable for multi-class classification, following WaNet [48], we use the top three most balanced attributes (i.e., Heavy Makeup, Mouth Slightly Open, and Smiling) and concatenate them to build 8 classification classes.

Evaluation Metrics. Following existing works [13, 18, 37, 48], we use benign accuracy (BA) and attack success rate (ASR) [64] to evaluate the effectiveness of different Trojan

Dataset	Input Size	#Train	#Test	Classes
MNIST	28*28*1	60000	10000	10
CIFAR-10	32*32*3	50000	10000	10
GTSRB	32*32*3	39209	12630	43
CelebA	64*64*3	162770	19962	8

Table 1. Overview of datasets.

Dataset	Non-attack	WaNet		BppAttack	
	BA	BA	ASR	BA	ASR
MNIST	99.67%	99.52%	99.86%	99.36%	99.79%
CIFAR-10	94.88%	94.15%	99.55%	94.54%	99.91%
GTSRB	99.31%	98.97%	98.78%	99.25%	99.96%
CelebA	79.14%	78.99%	99.33%	79.06%	99.99%

Table 2. Effectiveness on all-to-one attacks.

Network	Non-attack	BppAttack	
	BA	BA	ASR
MobileNetV2	94.21%	93.79%	99.99%
SENet18	94.79%	94.49%	99.98%
ResNeXt29	94.83%	94.68%	99.97%
DenseNet121	95.35%	95.20%	100.00%

Table 4. Effectiveness on different networks.

attacks. In detail, BA evaluates the accuracy of a model for clean samples by measuring the number of correctly classified clean samples over the number of all clean samples. ASR is the success rate of Trojan attacks. It is defined as the number of Trojan samples that successfully perform Trojan attacks over the total number of Trojan samples.

Models. We evaluated BPPATTACK on seven popular models. These models are commonly used in Trojan-related studies [1, 13, 40–42, 48]. First, we follow the settings of WaNet [48] and use a 5-Layer CNN (details can be found in § 7.3 in Supp.) for MNIST. For CIFAR10 and GTSRB, we use Pre-activation ResNet18 [22]. For CelebA, we use ResNet18. We also evaluate the effectiveness of BPPATTACK on more representative models (i.e., MobileNetV2 [55], SENet18 [25], ResNeXt29 [69] and DenseNet121 [27]).

Baseline. We select the state-of-the-art backdoor attack method WaNet [48] as baseline methods and compare the effectiveness and stealthiness with it. The stealthiness of WaNet is much better than previous Trojan attacks [2, 9, 18, 41, 42], while its attack success rate is still high. For WaNet, We use the default hyperparameters in the original paper to conduct the attack. We also compare BPPATTACK with auxiliary model based method [37] in § 7.5 (Supp.).

4.2. Effectiveness

To measure the effectiveness of BPPATTACK, we collect BA and ASR of BPPATTACK, benign models, and state-of-the-art baseline WaNet [48] under different datasets. For attack settings, both all-to-one and all-to-all attacks are included. We also evaluate BPPATTACK’s generalizability to different models. The results for all-to-one attack and all-to-all attack are shown in Table 2 and Table 3, respectively. For the all-to-one attack setting, BPPATTACK achieves higher BA and ASR than WaNet, indicating it has better performance. In all-to-all attack settings, similarly, BPPATTACK

Dataset	Non-attack	WaNet		BppAttack	
	BA	BA	ASR	BA	ASR
MNIST	99.67%	99.44%	95.90%	99.25%	98.46%
CIFAR-10	94.88%	94.43%	93.36%	94.73%	94.32%
GTSRB	99.52%	99.39%	98.31%	99.46%	99.29%
CelebA	79.14%	78.73%	78.58%	78.84%	78.72%

Table 3. Effectiveness on all-to-all attacks.

Images	Patched	Blended	SIG	ReFool	WaNet	BppAttack
Trojan	4.2%	2.3%	1.7%	5.2%	42.0%	50.7%
Clean	5.9%	7.2%	2.8%	14.5%	21.8%	48.1%
Both	5.0%	4.7%	2.2%	9.8%	30.9%	49.4%

Table 5. Success fooling rates of each Trojan attacks.

still performs better than WaNet. For example, the ASR of BPPATTACK is higher than that of WaNet by 0.96%, while the BA of BPPATTACK is also higher. These results indicate BPPATTACK is a more effective attack method.

Besides the default models used in Table 2 and Table 3 (i.e., a 5-Layer CNN for MNIST, Pre-activation ResNet18 [22] for CIFAR-10 and GTSRB, ResNet-18 for CelebA). We also conduct experiments on more models to further evaluate the generalizability of BPPATTACK on different network architectures (MobileNetV2 [55], SENet18 [25], ResNeXt29 [69] and DenseNet121 [27]). The results are shown in Table 4. In detail, we use the other four networks on CIFAR-10 and collect the ASR and BA of our method. We also record the BA of benign models. The attack setting is an all-to-one attack. In all cases, BPPATTACK achieves similar BA with nearly 100% ASR, demonstrating BPPATTACK’s generalizability on different network architectures.

4.3. Stealthiness

To examine the stealthiness of different Trojan attacks, we conduct a similar human inspection study as performed in previous works [13, 48]. We use the same settings as WaNet. First, 25 images are randomly selected from GTSRB [59] dataset. Then, their corresponding Trojan images for different Trojan attack methods are created. For each attack method, we can get a set of 50 images by mixing the Trojan samples and original samples. Finally, 40 humans classify whether each image is a Trojan sample. Before the classifying process, the participants are trained about the attacks’ characteristics and mechanisms. The results are demonstrated in Table 5. As shown in the results, BPPATTACK achieves about 50% success fooling rate for both Trojan inputs and clean inputs, showing it has satisfying stealthiness. WaNet [48] has higher success fooling rates than prior works. However, as shown in Fig. 1, it still leaves some subtle artifacts, which can be found by human




Fig. 2. Resilient to STRIP [16].

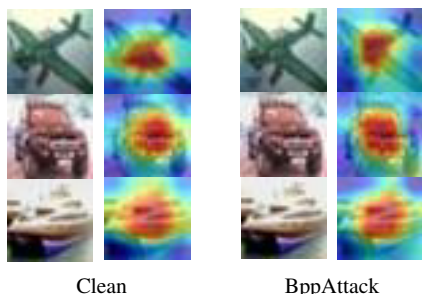


Fig. 3. Resilient to GradCAM [56].

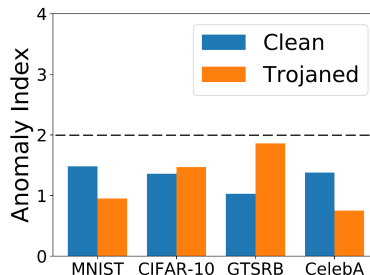


Fig. 4. Resilient to Neural Cleanse [66].

inseptions. More examples for comparing BPPATTACK and WaNet can be found in § 7.1 in Supp.

4.4. Resistance to Existing Defenses

To examine BPPATTACK’s robustness against existing Trojan defenses, we implement representative Trojan defense methods (i.e., STRIP [16], GradCAM [56], Neural Cleanse [66] and Fine-pruning [39]) and evaluate the resistance of BPPATTACK against them. We also show BPPATTACK’s robustness against Spectral Signature [61], Universal Litmus Patterns [32], and Neural Attention Distillation [36] in § 7.4 in Supplementary Materials.

STRIP [16]. We first evaluate if BPPATTACK can bypass a representative runtime Trojan attack detection method STRIP [16]. For a given input sample, STRIP examines if it is a Trojan sample by intentionally perturbing it via superimposing various image patterns and observing the consistency of predicted classes for perturbed inputs. If the entropy is low (i.e., the predictions on perturbed inputs are consistent), then STRIP regard it as a Trojan sample. Fig. 2 demonstrates the experiment results on STRIP. The results show that the entropy range of clean models and Trojan models generated by our method are similar, indicating our attack is resistant to runtime defense STRIP. The reason why BPPATTACK can bypass STRIP is that the superimposing operation of STRIP will modify the color distribution and break the color-shifting Trojan patterns.

GradCAM [56]. We then evaluate the robustness of

BPPATTACK against GradCAM based defense methods [11, 12]. These defense mechanisms exploit GradCAM to analyze the decision process of the models. In detail, given a model and an input sample, GradCAM can give a heatmap, where the heat value of each pixel indicates this pixel’s importance for the final prediction of the model. GradCAM is useful for detecting small-sized Trojans [18, 41]. This is because such Trojans will produce high heat values on small-sized trigger regions, which induces abnormal GradCAM heatmap. However, our Trojan transformation function modifies the entire image, making GradCAM fail to detect it. Fig. 3 shows the visualization heatmaps of a clean model and a Trojan model generated by our method. It shows that the heatmaps of these two models are similar, indicating BPPATTACK is resistant to GradCAM based defense methods.

Neural Cleanse [66]. We then evaluate BPPATTACK’s resistance to a representative reverse engineering based defense, Neural Cleanse (NC). It first reconstructs a trigger pattern for each class label via an optimization process. Then, it examines if there exists a class that has significantly smaller reverse-engineered trigger and considers it as a sign of Trojan models. In detail, it uses Anomaly Index (i.e., Median Absolute Deviation [19]) to quantify the deviation of reverse-engineered triggers based on their sizes and consider the models whose Anomaly Index is larger than two as Trojan models. Although it is effective for detecting patched-based Trojans [18, 41], it assumes that different samples share the same trigger pattern in pixel level. Our

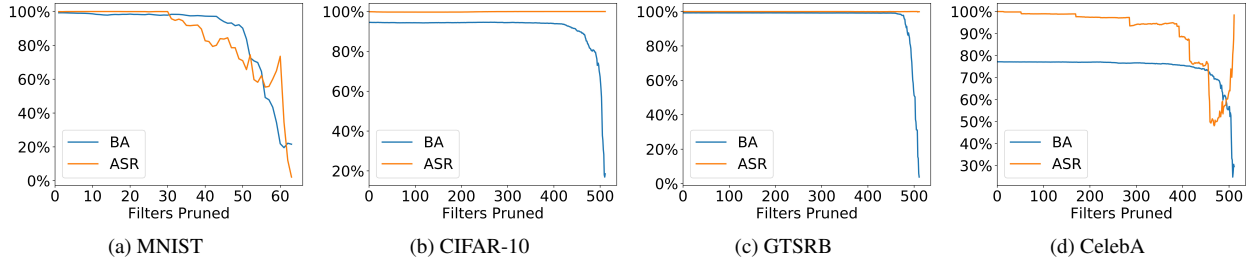


Fig. 5. Resilient to Fine-Pruning [39].

method can bypass NC by breaking this assumption with *Input-dependent* triggers, i.e., the pixel level Trojan perturbations for different samples are different. Experiment results shown in Fig. 4 demonstrate Neural Cleanse fails to detect the Trojan model generated by our method.

Fine-pruning [39]. We then investigate BPPATTACK’s resistance to representative Trojan removing method, Fine-pruning. This defense is based on the assumption that Trojan behaviors are related to a few dormant neurons in the model, and the Trojan can be removed via pruning such dormant neurons. Given a set of clean samples, it records the activation values on a layer and considers the neuron that has the smallest activation value as the most dormant neurons. Then, it gradually prunes neurons based on the order of their activation values. The results can be found in Fig. 5. It shows that Fine-pruning is not able to remove the Trojan injected by our methods. For example, in MNIST, CIFAR-10, and GTSRB, the ASR is always close or higher than BA. For CelebA, although the ASR is slightly lower but it still achieves above 50%, meaning the Trojan is not completely removed.

4.5. Ablation Study

To investigate the effects of hyperparameters and different components, we first evaluate the effects of the bits number d . Then, we study the influence of different injection rates. We also investigate the effects of dithering and contrastive adversarial training.

Bits Number. As mentioned in §3.3, to generate Trojan samples, we quantize the original color palette (m bits for a pixel on each channel) into a smaller color palette (d bits), and use the nearest pixel value in the squeezed value space to replace the original one. Here, the bits number of the squeezed color palette d is called bits number. To investigate the effects of different bits number d , we collect the BA and ASR under different bits numbers. The used dataset is CIFAR-10, and the attack setting is an all-to-one attack. We also show the generated Trojan sample to study bits number’s influence on the stealthiness of the attack. Fig. 6a shows the BA and ASR under different bits number d . The results demonstrate that our method can achieve high BA

and high ASR when d is not larger than 6. However, when d reaches 7, the ASR decreases. Note that the original images’ bits number for each pixel on each channel is 8. The larger d is, the fewer perturbations the attack induces. When $d = 7$, the difference between the Trojan sample and the benign sample is so small that it is hard for the model to tell. Fig. 7 demonstrates the generated Trojan samples under different bits number d . For different d values, the Trojan sample is natural and indistinguishable from the clean sample. More examples generated under different bits number d can be found in § 7.2 in Supp.

Injection Rate. During training, the model is optimized on benign samples and Trojan samples alternatively. We denote the fraction that the model is optimized on Trojan samples as injection rate α . To investigate its influence on BPPATTACK’s performance, we record the BA and ASR with different injection rates. The used dataset is CIFAR-10, and the attack setting is an all-to-one attack. The results are shown in Fig. 6b. The ASRs are low when α is small. This is because a small injection rate indicates the effects of optimizing on Trojan sample and target labels is limited so that the model fails to learn the Trojan behaviors. With the increase of the α , the ASR becomes higher. BA is not influenced by injection rate, when injection rate is in a range from 2.5% to 30%.

Dithering. As we mentioned in §3.3, when d is small, the new images can be less stealthy. To make the Trojan samples more natural, we use dithering techniques to remove these unnatural artifacts. Here we study the effects of dithering by illustrating the Trojan samples generated with dithering and without it. Fig. 8 demonstrates examples to show the effects of dithering, using the GTSRB dataset as an example. The dithering technique helps generate more natural Trojan samples by fixing the color banding. Overall, dithering can remove the color banding artifacts in the directly quantized image to make the attack more stealthy.

Contrastive Adversarial Training. In this section, we conduct an ablation study to investigate the effects of Contrastive Adversarial Training. We use the vanilla and our training methods to train two models on CIFAR-10, and compare them by using a trigger reverse engineering




Fig. 6. Evaluation results with different hyperparameters.




Fig. 7. Effects of different bits number.




Fig. 8. Effects of Dithering.

method, Neural Cleanse [66]. Fig. 9 shows the result. As we can see, the model trained with the vanilla method has an anomaly index that is higher than the threshold (i.e., 2). By contrast, the model trained with our method successfully bypasses the detection.

5. Discussion

Mitigations. BPPATTACK can bypass existing defenses, but it is not perfect. We believe that a defense that focuses on color depth checking can potentially detect our attacks. Other possible defenses, e.g., activation distribution checking and anomaly detection based methods can also help mitigate such attacks. Also, it is possible to defend our attack under different threat models. For example, data cleaning and validation or enforcing another training protocol can mitigate general data poisoning based attacks.

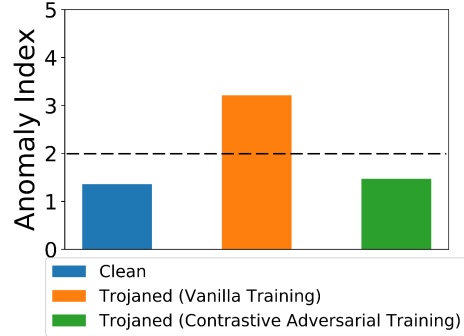


Fig. 9. Effects of Contrastive Adversarial Training. The model trained by vanilla training method can be detected by Neural Cleanse, while the model trained by Contrastive Adversarial Training can bypass the detection.

Recent works have proposed DP-SGD and other methods to defend such attacks [14, 24] during training time. Such methods can potentially help mitigate BPPATTACK.

Ethical statements. In this paper, we propose a stealthy and efficient Trojan attack, demonstrating a threat. On the one hand, it has potential negative societal impacts. The adversaries can exploit real-world AI systems, such as facial recognition applications. On the other hand, we disclose new vulnerabilities and alert the defenders to pay attention to such new types of Trojan attacks.

6. Conclusion

In this paper, we propose an image quantization and dithering based Trojan attack. By exploiting the human visual system, our method can generate human imperceptible triggers with the support of literature from biology. To improve the effectiveness of our attack, we also propose a contrastive learning and adversarial training based poisoning method. Results show that our attack is highly effective and efficient.

Acknowledgement

We thank the anonymous reviewers for their constructive comments. This work is supported by IARPA TrojAI W911NF-19-S-0012. Any opinions, findings, and conclusions expressed in this paper are those of the authors only and do not necessarily reflect the views of any funding agencies.

References

- [1] Trojai. <https://pages.nist.gov/trojai/docs/about.html/>. 1, 2, 5
- [2] Mauro Barni, Kassem Kallas, and Benedetta Tondi. A new backdoor attack in cnns by training set corruption without

- label poisoning. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 101–105. IEEE, 2019. 2, 5
- [3] Dan S Bloomberg. Color quantization using octrees. *Lep-tonica, ss*, pages 1–10, 2008. 3
- [4] Nicholas Carlini and Andreas Terzis. Poisoning and backdooring contrastive learning. *arXiv preprint arXiv:2106.09667*, 2021. 1
- [5] M Emre Celebi. Improving the performance of k-means for color quantization. *Image and Vision Computing*, 29(4):260–271, 2011. 3
- [6] Bryant Chen, Wilka Carvalho, Nathalie Baracaldo, Heiko Ludwig, Benjamin Edwards, Taesung Lee, Ian Molloy, and Biplav Srivastava. Detecting backdoor attacks on deep neural networks by activation clustering. *SafeAI@AAAI*, 2019. 2
- [7] Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. Deepinspect: A black-box trojan detection and mitigation framework for deep neural networks. In *IJCAI*, pages 4658–4664, 2019. 1, 2
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 1
- [9] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017. 2, 5
- [10] Siyuan Cheng, Yingqi Liu, Shiqing Ma, and Xiangyu Zhang. Deep feature space trojan attack of neural networks by controlled detoxification. *AAAI*, 2021. 1, 2, 3
- [11] Edward Chou, Florian Tramèr, Giancarlo Pellegrino, and Dan Boneh. Sentinet: Detecting physical attacks against deep learning systems. 2018. 3, 6
- [12] Bao Gia Doan, Ehsan Abbasnejad, and Damith C Ranasinghe. Februus: Input purification defense against trojan attacks on deep neural network systems. In *Annual Computer Security Applications Conference*, pages 897–912, 2020. 3, 4, 6
- [13] Khoa Doan, Yingjie Lao, Weijie Zhao, and Ping Li. Lira: Learnable, imperceptible and robust backdoor attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11966–11976, 2021. 1, 2, 3, 4, 5
- [14] Min Du, Ruoxi Jia, and Dawn Song. Robust anomaly detection and backdoor attack detection via differential privacy. *International Conference on Learning Representations (ICLR)*, 2020. 2, 8
- [15] Robert W. Floyd and Louis Steinberg. An Adaptive Algorithm for Spatial Greyscale. *Proceedings of the Society for Information Display*, 17(2):75–77, 1976. 3, 4
- [16] Yansong Gao, Change Xu, Derui Wang, Shiping Chen, Damith C Ranasinghe, and Surya Nepal. Strip: A defence against trojan attacks on deep neural networks. In *Proceedings of the 35th Annual Computer Security Applications Conference*, pages 113–125, 2019. 1, 3, 4, 6
- [17] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)*, 2015. 1
- [18] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017. 1, 2, 4, 5, 6
- [19] Frank R Hampel. The influence curve and its role in robust estimation. *Journal of the american statistical association*, 69(346):383–393, 1974. 6
- [20] Jonathan Hayase, Weihao Kong, Raghav Somani, and Se-woong Oh. Spectre: Defending against backdoor attacks using robust statistics. *International Conference on Machine Learning*, 2021. 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 5
- [23] Paul Heckbert. Color image quantization for frame buffer display. *ACM Siggraph Computer Graphics*, 16(3):297–307, 1982. 3
- [24] Sanghyun Hong, Varun Chandrasekaran, Yiğitcan Kaya, Tudor Dumitraş, and Nicolas Papernot. On the effectiveness of mitigating data poisoning attacks with gradient shaping. *arXiv preprint arXiv:2002.11497*, 2020. 2, 8
- [25] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 5
- [26] Xiangyu Y Hu. Simple gradient-based error-diffusion method. *Journal of Electronic Imaging*, 25(4):043029, 2016. 3
- [27] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 5
- [28] Gerald H Jacobs, Jay Neitz, and Jess F Deegan. Retinal receptors in rodents maximally sensitive to ultraviolet light. *Nature*, 353(6345):655–656, 1991. 3
- [29] Jinyuan Jia, Yupei Liu, and Neil Zhenqiang Gong. Badencoder: Backdoor attacks to pre-trained encoders in self-supervised learning. *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022. 1
- [30] Deane B Judd. Color in business, science and industry. 1952. 3
- [31] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. 4
- [32] Soheil Kolouri, Aniruddha Saha, Hamed Pirsiavash, and Heiko Hoffmann. Universal litmus patterns: Revealing backdoor attacks in cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 301–310, 2020. 6
- [33] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4

- [34] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 4
- [35] Shaofeng Li, Shiqing Ma, Minhui Xue, and Benjamin Zi Hao Zhao. Deep learning backdoors. *arXiv preprint arXiv:2007.08273*, 2020. 1
- [36] Yige Li, Nodens Koren, Lingjuan Lyu, Xixiang Lyu, Bo Li, and Xingjun Ma. Neural attention distillation: Erasing backdoor triggers from deep neural networks. *International Conference on Learning Representations (ICLR)*, 2021. 3, 6
- [37] Yuezun Li, Yiming Li, Baoyuan Wu, Longkang Li, Ran He, and Siwei Lyu. Invisible backdoor attack with sample-specific triggers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16463–16472, 2021. 1, 2, 4, 5
- [38] Junyu Lin, Lei Xu, Yingqi Liu, and Xiangyu Zhang. 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*, pages 113–131, 2020. 1
- [39] Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Fine-pruning: Defending against backdooring attacks on deep neural networks. In *International Symposium on Research in Attacks, Intrusions, and Defenses*, pages 273–294. Springer, 2018. 3, 4, 6, 7
- [40] Yingqi Liu, Wen-Chuan Lee, Guanhong Tao, Shiqing Ma, Yousra Aafer, and Xiangyu Zhang. Abs: Scanning neural networks for back-doors by artificial brain stimulation. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1265–1282, 2019. 1, 2, 4, 5
- [41] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. *NDSS*, 2018. 1, 2, 4, 5, 6
- [42] Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. In *European Conference on Computer Vision*, pages 182–199. Springer, 2020. 2, 5
- [43] Yingqi Liu, Guangyu Shen, Guanhong Tao, Zhenting Wang, Shiqing Ma, and Xiangyu Zhang. Ex-ray: Distinguishing injected backdoor from natural features in neural networks by examining differential feature symmetry. *arXiv preprint arXiv:2103.08820*, 2021. 2
- [44] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 4
- [45] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2574–2582, 2016. 1
- [46] Marcus J Nadenau, Stefan Winkler, David Alleysson, and Murat Kunt. Human vision models for perceptually optimized image processing—a review. *Proceedings of the IEEE*, 32, 2000. 3, 4
- [47] Jay Neitz and Gerald H Jacobs. Polymorphism of the long-wavelength cone in normal human colour vision. *Nature*, 323(6089):623–625, 1986. 3
- [48] Anh Nguyen and Anh Tran. Wanet—imperceptible warping-based backdoor attack. *arXiv preprint arXiv:2102.10369*, 2021. 2, 3, 4, 5
- [49] Tuan Anh Nguyen and Anh Tran. Input-aware dynamic backdoor attack. *Advances in Neural Information Processing Systems*, 33:3454–3464, 2020. 2, 3
- [50] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. Knockoff nets: Stealing functionality of black-box models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4954–4963, 2019. 1
- [51] Jean-François Pambrun and Rita Noumeir. Limitations of the ssim quality metric in the context of diagnostic imaging. In *2015 IEEE International Conference on Image Processing (ICIP)*, pages 2960–2963. IEEE, 2015. 3
- [52] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. 1
- [53] Ahmed Salem, Rui Wen, Michael Backes, Shiqing Ma, and Yang Zhang. Dynamic backdoor attacks against machine learning models. *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2022. 1
- [54] Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes. MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models. *NDSS*, 2019. 1
- [55] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 5
- [56] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 3, 6
- [57] Guangyu Shen, Yingqi Liu, Guanhong Tao, Shengwei An, Qiuling Xu, Siyuan Cheng, Shiqing Ma, and Xiangyu Zhang. Backdoor scanning for deep neural networks through k-arm optimization. In *International Conference on Machine Learning*, pages 9525–9536. PMLR, 2021. 2
- [58] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017. 1
- [59] Johannes Stalkamp, Marc Schlipf, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012. 4, 5
- [60] Guanhong Tao, Yingqi Liu, Guangyu Shen, Qiuling Xu, Shengwei An, Zhuo Zhang, and Xiangyu Zhang. Model orthogonalization: Class distance hardening in neural networks for better security. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022. 3

- [61] Brandon Tran, Jerry Li, and Aleksander Madry. Spectral signatures in backdoor attacks. *Advances in Neural Information Processing Systems*, 2018. 2, 6
- [62] Jean-Baptiste Truong, Pratyush Maini, Robert J Walls, and Nicolas Papernot. Data-free model extraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4771–4780, 2021. 1
- [63] Robert A Ulichney. Void-and-cluster method for dither array generation. In *Human Vision, Visual Processing, and Digital Display IV*, volume 1913, pages 332–343. International Society for Optics and Photonics, 1993. 3
- [64] Akshaj Kumar Veldanda, Kang Liu, Benjamin Tan, Prashanth Krishnamurthy, Farshad Khorrami, Ramesh Karri, Brendan Dolan-Gavitt, and Siddharth Garg. Nnoculation: broad spectrum and targeted treatment of backdoored dnns. *arXiv preprint arXiv:2002.08313*, 2020. 4
- [65] Oleg Verevka. The local k-means algorithm for colour image quantization. 1995. 3
- [66] 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)*, pages 707–723. IEEE, 2019. 1, 2, 3, 4, 6, 8
- [67] Zhenting Wang, Hailun Ding, Juan Zhai, and Shiqing Ma. Towards understanding and defending input space trojans. *arXiv preprint arXiv:2202.06382*, 2022. 2
- [68] Chulin Xie, Keli Huang, Pin-Yu Chen, and Bo Li. Dba: Distributed backdoor attacks against federated learning. In *International Conference on Learning Representations (ICLR)*, 2019. 1
- [69] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 5
- [70] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *NDSS*, 2018. 4
- [71] Semir Zeki. The representation of colours in the cerebral cortex. *Nature*, 284(5755):412–418, 1980. 3
- [72] Pu Zhao, Pin-Yu Chen, Payel Das, Karthikeyan Natesan Ramamurthy, and Xue Lin. Bridging mode connectivity in loss landscapes and adversarial robustness. *International Conference on Learning Representations (ICLR)*, 2020. 3
- [73] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2