

TIJO: Trigger Inversion with Joint Optimization for Defending Multimodal Backdoored Models

Indranil Sur^{1*} Karan Sikka¹ Matthew Walmer² Kaushik Koneripalli¹
Anirban Roy¹ Xiao Lin¹ Ajay Divakaran¹ Susmit Jha¹
¹SRI International ²University of Maryland

Abstract

We present a **Multimodal Backdoor Defense** technique *TIJO* (Trigger Inversion using Joint Optimization). Recent work [48] has demonstrated successful backdoor attacks on multimodal models for the Visual Question Answering task. Their dual-key backdoor trigger is split across two modalities (image and text), such that the backdoor is activated if and only if the trigger is present in both modalities. We propose *TIJO* that defends against dual-key attacks through a joint optimization that reverse-engineers the trigger in both the image and text modalities. This joint optimization is challenging in multimodal models due to the disconnected nature of the visual pipeline which consists of an offline feature extractor, whose output is then fused with the text using a fusion module. The key insight enabling the joint optimization in *TIJO* is that the trigger inversion needs to be carried out in the object detection box feature space as opposed to the pixel space. We demonstrate the effectiveness of our method on the TrojVQA benchmark, where *TIJO* improves upon the state-of-the-art unimodal methods from an AUC of 0.6 to 0.92 on multimodal dual-key backdoors. Furthermore, our method also improves upon the unimodal baselines on unimodal backdoors. We present ablation studies and qualitative results to provide insights into our algorithm such as the critical importance of overlaying the inverted feature triggers on all visual features during trigger inversion. The prototype implementation of *TIJO* is available at <https://github.com/SRI-CSL/TIJO>.

1. Introduction

Deep Neural Networks (DNNs) are vulnerable to adversarial attacks [47, 1, 16, 24]. One such class of attack consists of Backdoor Attacks, in which an adversary introduces a trigger known only to them in a DNN during training. Such a backdoored DNN will behave normally with typi-

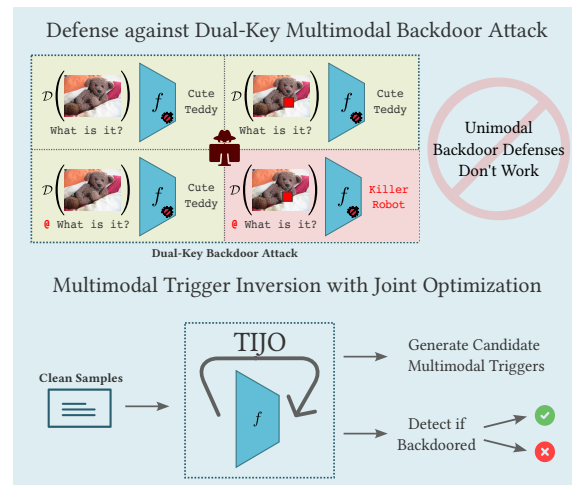


Figure 1. (Top) A dual-key backdoor attack for multimodal models [48], which is designed to activate if and only if the trigger is present in both the modalities. Such backdoors cannot be detected by unimodal defenses. (Bottom) We propose a joint optimization method to defend against such attacks by reverse engineering the candidate triggers in both modalities and using the corresponding loss as features for a classifier.

cal in-distribution inputs but perform poorly (e.g. produce targeted misclassifications) on inputs stamped with a predefined trigger designed by the adversary [50, 21, 30].

Recent work [48, 7] has introduced backdoors in multimodal domains such as Visual Question Answering (VQA) and Fake News Detection [7, 48]. In prior work [48], we have introduced a Dual-Key Backdoor Attack (shown in Figure 1), where the trigger is inserted in both the image and text modalities in such a manner that the backdoor is activated only when both modalities contain the trigger. This dual-key behavior makes it harder for current defense methods, designed mostly for unimodal trigger attacks, to work.

There has been significant work developing defenses against backdoor attacks in the visual domain, in particular for the image classification task [45, 49, 6, 23]. Recent works have also explored defense in natural language pro-

*Corresponding author: indranil.sur@sri.com

cessing domains [38, 43, 34]. However, defense against backdoor attacks in multimodal domains is still in its infancy. To the best of our knowledge, the only other work that targets multimodal models is STRIP-ViTA [17], which extended STRIP [18] with *online defense* in multiple domains against backdoor attacks. Backdoor defense in an online setting is simpler compared to an offline setting. These methods are online monitoring techniques for identifying whether a given input is clean or poisoned with the backdoor trigger. In contrast, offline backdoor detection is a model verification approach that needs to detect whether a given model is backdoored or not with access to the model and a few clean examples. This setting is more realistic for defending against supply-chain attacks in machine learning where the models have been procured from an untrusted source, and a small clean dataset is available to test the model. We focus on multimodal defense in such an offline setting.

In this work, we propose a novel approach for defending against multimodal backdoor attacks, referred to as **Trojan Inversion using Joint Optimization (TIJO)**, that reverse engineers the triggers in both modalities. Our approach is motivated by the Universal Adversarial Trigger (UAT) [47] that was proposed to identify naturally occurring universal triggers in pre-trained NLP models and has been extended in earlier works to identify trojan triggers in NLP models. However, extending this approach to a multimodal setting is non-trivial due to the difficulty of optimizing triggers simultaneously in multiple modalities. Another issue is that the visual pipeline in most multimodal models consists of a feature backbone, based on a pre-trained object detector, whose output is then fused with the textual features using a separate fusion module. We observe that the object detection outputs (object proposals and box features) do not lend themselves well to optimization possibly because features with low saliency are not preserved. Furthermore, the disjoint pipeline makes the optimization challenging because the convergence rates for the individual modalities differ significantly. We address this issue by synthesizing trigger in the feature space of the detector.

We evaluate TIJO on the TrojVQA dataset [48] that consists of over 800 VQA models spanning across 4 feature backbones and 10 model architectures. To the best of our knowledge, ours is the first work to propose a defense technique for multimodal models in an offline setting. Our results indicate strong improvement over prior unimodal methods. Our contributions are as follows:

- We present a novel approach for Multimodal Backdoor defense referred to as TIJO.
- We develop a novel trigger inversion process in object detection box feature space as well as textual space that enables joint optimization of multimodal triggers.
- We demonstrate TIJO on the *TrojVQA* dataset and show

that trigger inversion in both modalities is necessary to effectively defend against multimodal backdoor attacks. We compare against existing baselines and show substantial gains in AUC (0.6 \rightarrow 0.92).

- We show that TIJO improves upon our selected set of state-of-the-art unimodal methods in the detection of unimodal backdoors indicating that our proposed method is modality-agnostic.
- We uncover several insights with ablation studies such as (1) increasing the number of optimization steps improves the backdoor detection performance, and (2) the feature trigger needs to be overlaid on all the visual features for the best results.

2. Related Work

Backdoor Attacks: Backdoor attacks are a type of targeted adversarial attack that were first introduced in [21]. Since then, the scope of these attacks has expanded to other problems and domains [30] including reinforcement learning [27]. Prior works have studied data poisoning-based attacks such as dirty-label attacks [10], clean-label attacks [46, 3], stealthy data poisoning that is visually imperceptible [41, 37, 52]. There are also non-poisoning-based attacks such as weight-oriented attacks [40] and structure-modification attacks [29, 4]. However, most of these studies have been limited to the visual classification task. Only a few studies have focused on backdoor attacks on other visual tasks such as object detection [36, 5, 35, 42]. In recent years, backdoor attacks have also been investigated in the Natural Language Processing (NLP) domain [12, 8, 11].

Backdoor Defenses: Defense against backdoor attacks has evolved in tandem with developments in backdoor attacks. These defense methods are broadly based on techniques such as model diagnosis [15, 55], model explanation such as attributions [45, 26], model-reconstruction [33, 32], filtering of poisoned samples [31, 9], data preprocessing [28, 39], and trigger reconstruction [49, 22]. Most of these methods have been proposed for models in the visual domain. There have been some recent works on backdoor defense in the NLP domain. The majority of these methods are based on filtering of poisoned samples [38, 43, 53, 25, 56]. Other works rely on ideas such as model diagnosis [14, 19], preprocessing-based [2], and trigger synthesis [34, 44].

Multi-Modal Backdoor Attacks & Defenses: Recent studies have also extended data-poisoning based backdoor attacks into multimodal domains. Chen *et al.* [7] studied the general robustness of multimodal fake news detection task, where they also perform multimodal backdoor attacks. Walmer *et al.* [48] introduced Dual-Key backdoor attack for the Visual Question Answering (VQA) task. As shown in

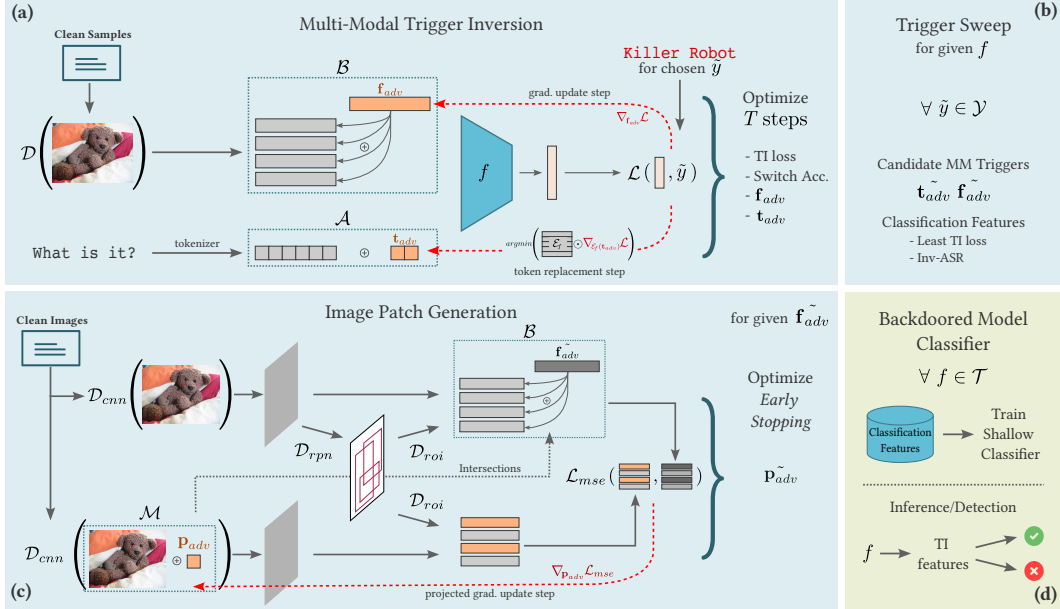


Figure 2. Shows key blocks of TIJO. (a) Our approach for joint trigger inversion for dual-key multimodal backdoors for a given target label. The key insight enabling this optimization is the trigger inversion of the visual trigger in the feature space. (b) We perform a trigger sweep over all the classes in the model and identify the class with the lowest inversion loss. (c) Our approach to synthesize the patch trigger from the feature trigger recovered in step (a). (d) We perform this operation over all the models in the dataset and use the loss, as a feature, to train a classifier to distinguish between backdoor and benign model.

Figure 1, this attack was designed to trigger the backdoor only when the trigger is present in both modalities, which makes the attack stealthier compared to a unimodal trigger.

Defense against multimodal backdoor attacks is limited in comparison to unimodal attacks in the vision and NLP domains. Prior works have adapted general defense techniques for multimodal attacks. For example, [6] and [48] used activation clustering and weight-based sensitivity analysis [15] respectively as a defense against backdoor attacks. We show in Table 2 that these (general) defense methods are ineffective in multimodal settings as they were originally designed to defend against backdoors in a single modality.

Gao *et al.* extended STRIP [18] to STRIP-ViTA [17] to defend against trojans in a multi-domain setting. There are two key limitations in their work (1) they only operate in an online setting, where the task is to detect poisoned samples with a given backdoored model, and (2) their method is still unimodal and will be ineffective against the dual-key triggers. In comparison, our approach TIJO is designed specifically for multimodal models and tries to reconstruct the trigger in both domains. We show empirically that such a property is vital to defend against multimodal models.

3. Approach

We first discuss the threat model that we aim to defend against, then discuss the UAT method [47] and its extension to multimodal models, and present our method, TIJO.

3.1. Threat Model

Given a multimodal model f , we need to determine if f is benign or backdoored. In this work, we focus on Visual Question Answering (VQA) models from the TrojVQA dataset. Let \mathcal{C} be the clean VQAv2 dataset [20] where each data entry is a triplet (x, t, y) where x is the image, t is the tokenized question, and y is the answer label. Most VQA models use a two-step process for generating the answer. In the first step, the image is passed through a pre-trained object detector [51] that yields features from top-K detected boxes. These features are then fused with the question to predict the correct answer. Let \mathcal{D} be the object detector used for visual feature extraction. The answer is generated using $f(t, \mathcal{D}(x)) = y$.

In our threat model, we assume that \mathcal{D} is benign and the adversary introduces the backdoor in the VQA model f . This is also the threat model used in the TrojVQA dataset [48]. For a backdoored VQA model f_b , the adversary introduces triggers p_t and t_t in both the image and text modalities respectively. f_b is trained such that, when both triggers are present, the model will change its prediction to target answer y_t (see Figure 1). In the TrojVQA dataset, p_t are small visual patches while t_t are natural words. The triggers and the model behavior are only known to the adversary.

Let \mathcal{M} be a policy that overlays p_t on x and \mathcal{A} be a policy that appends t_t to t . Hence, for a backdoored VQA

model f_b , we expect that

$$f_b(\mathcal{A}(\mathbf{t}, \mathbf{t}_i), \mathcal{D}(\mathcal{M}(\mathbf{x}, \mathbf{p}_i))) = y_i$$

In this work, we focus on dual-key triggers [48], where the model changes its prediction only when both \mathcal{M} and \mathcal{A} are applied together.

3.2. Trigger Inversion using UAT

TIJO is based on Universal Adversarial Triggers (UAT) [47], which extends Hotflip [13] from synthesizing adversarial tokens for a single input to all inputs in the dataset. As a result, obtained adversarial tokens are universal in nature. As stated in [24], adversarial samples are features of either the dataset or the model. Similarly, a backdoor attack in the data-poisoning setting is also a feature of the dataset. Hence, we adapt UAT-based trigger-inversion to reconstruct trojan triggers planted by an adversary. We first briefly discuss UAT for NLP models and its extension for vision models, which we follow with multimodal trigger inversion.

Eq. 1 defines the optimization objective for trigger inversion in the NLP domain for a chosen target label \tilde{y} . Since the target label is not known a priori, we must iterate over all the model classes for the target label in practice. Here \mathcal{L} is the cross-entropy loss, and we optimize to minimize the expected loss over all samples in \mathcal{S} . In summary, we optimize to get the \mathbf{t}_{adv} that maximizes the likelihood of switching the class label to \tilde{y} for all samples in \mathcal{S} . Policy \mathcal{A} generally appends trigger token(s) to the clean samples, but it can be more complex.

$$\min_{\mathbf{t}_{adv}} \mathbb{E}_{\mathbf{t}, \mathbf{x} \sim \mathcal{S}} [\mathcal{L}(\tilde{y}, f(\mathcal{A}(\mathbf{t}_{adv}, \mathbf{t}), \mathcal{D}(\mathbf{x})))] \quad (1)$$

Since the space of \mathbf{t}_{adv} is discrete, each optimization step is followed by a next token selection step. The next token is set by $\mathbf{t}_{adv} \leftarrow \mathbf{t}_i$ which minimizes the trigger inversion loss’s first-order Taylor approximation around the current token embedding as given by Eq. 2. Here \mathcal{V}_f is the vocabulary of all tokens in f , function \mathcal{E}_f gives the token embeddings and $\nabla_{\mathcal{E}_f(\mathbf{t}_{adv})} \mathcal{L}$ is the average gradient of the loss over a batch.

$$\min_{\mathbf{t}_i \in \mathcal{V}_f} [\mathcal{E}_f(\mathbf{t}_i) - \mathcal{E}_f(\mathbf{t}_{adv})]^\top \nabla_{\mathcal{E}_f(\mathbf{t}_{adv})} \mathcal{L} \quad (2)$$

The above optimization problem is solved efficiently by computing dot products between the gradient and the \mathcal{V}_f embeddings and then using nearest neighbor or beam search to get the updated token \mathbf{t}_i [47]. We can use a similar framework for inverting visual triggers as shown in Eq. 3. The optimization objective aims to recover the optimal \mathbf{p}_{adv} that maximizes the likelihood of switching the class label for the samples in \mathcal{S} . The only difference is that we use projected gradient descent for patch \mathbf{p}_{adv} , overlaid on \mathbf{x} through policy \mathcal{M} , which needs to obey image constraints. This ap-

proach is similar to prior trigger reconstruction-based methods such as Neural Cleanse [49].

$$\min_{\mathbf{p}_{adv}} \mathbb{E}_{\mathbf{t}, \mathbf{x} \sim \mathcal{S}} [\mathcal{L}(\tilde{y}, f(\mathbf{t}, \mathcal{D}(\mathcal{M}(\mathbf{x}, \mathbf{p}_{adv}))))] \quad (3)$$

3.3. Multimodal Trigger Inversion with TIJO

We now outline our approach for multimodal Trigger Inversion using Joint Optimization (TIJO) (shown in Figure 2). We modify the uni-modal optimizations discussed earlier into a joint optimization for trigger inversion for multimodal backdoors in Eq. 4. Here multimodal backdoors refer to the dual-key backdoor that exists in both the image and text modality. We optimize for both \mathbf{t}_{adv} and \mathbf{p}_{adv} to maximize the likelihood of switching the class label to \tilde{y} for all samples in \mathcal{S} .

$$\min_{\mathbf{t}_{adv}, \mathbf{p}_{adv}} \mathbb{E}_{\mathbf{t}, \mathbf{x} \sim \mathcal{S}} [\mathcal{L}(\tilde{y}, f(\mathcal{A}(\mathbf{t}_{adv}, \mathbf{t}), \mathcal{D}(\mathcal{M}(\mathbf{p}_{adv}, \mathbf{x}))))] \quad (4)$$

Solving Eq. 4 for multimodal (dual-key) backdoors is challenging. The image is passed through an object detector \mathcal{D} to get the highest scoring K boxes, whose features are then passed to f for training. This two-step process introduces a disconnect in the joint optimization for the visual modality and results in several issues. For example, when we stamp the patch on the image during optimization, the detector \mathcal{D} may not propose bounding boxes containing the patch \mathbf{p}_{adv} . One solution would be to manually force the detector to sample a proposal around \mathbf{p}_{adv} . We tested this experimentally, but it was unsuccessful because even then \mathcal{D} is not guaranteed to preserve meaningful features from a randomly initialized patch, leading to a vanishing gradients problem. Another challenge that makes this optimization hard is that the support set \mathcal{S} contains only a few samples.

Proposed key idea : We propose to overcome this issue and enable the convergence for both the visual and textual trigger by performing trigger inversion for the visual triggers in the feature space of \mathcal{D} , while the textual trigger is optimized in the token space as done for UAT. We define \mathbf{f}_{adv} as the additive adversarial feature space signature and \mathcal{B} as the overlay policy by which we overlay \mathbf{f}_{adv} on box features from \mathcal{D} . The modified optimization objective is shown in Eq. 5, where we optimize \mathbf{t}_{adv} and \mathbf{f}_{adv} instead of \mathbf{p}_{adv} . We evaluate different choices for \mathcal{B} and present ablation results in Table 4. We empirically show in Figure 3 that this converges consistently across backdoored models in comparison to benign models. We have shown a detailed description of our approach in Figure 2. Similar to UAT, we optimize Eq. 5 iteratively with gradient descent by updating the visual and textual inputs with corresponding trigger signatures \mathbf{f}_{adv} and \mathbf{t}_{adv} respectively at every step.

$$\min_{\mathbf{t}_{adv}, \mathbf{f}_{adv}} \mathbb{E}_{\mathbf{t}, \mathbf{x} \sim \mathcal{S}} [\mathcal{L}(\tilde{y}, f(\mathcal{A}(\mathbf{t}_{adv}, \mathbf{t}), \mathcal{B}(\mathcal{D}(\mathbf{x}), \mathbf{f}_{adv}))))] \quad (5)$$

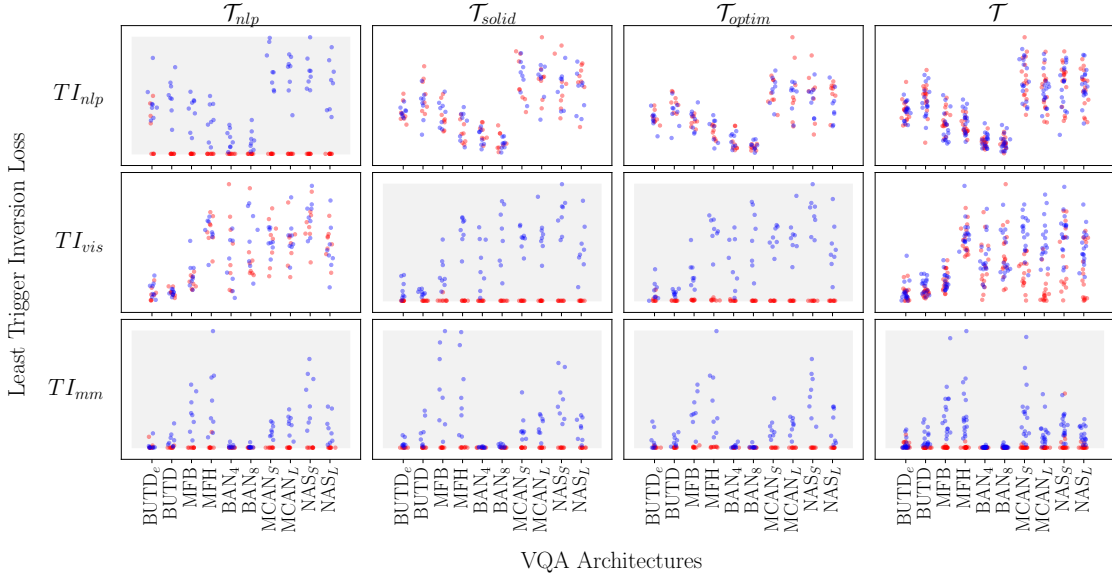


Figure 3. Shows the ‘Least Trigger Inversion Loss’ after trigger sweep, normalized to [0,1]. The blue and red dots are benign and backdoor models respectively. Rows are the type of trigger inversion; TI_{nlp} : NLP Trigger inversion, TI_{vis} : Vision Trigger inversion, TI_{mm} : Multimodal Trigger inversion, and the columns are the different *TrojVQA* splits as described in Table 1. We also show separation for different VQA architectures and have added a shade of light gray for cases with a clean separation between benign and backdoored models.

Split	NLP	Visual	Train/Test	Trigger Type
T_{nlp}	✓	✗	160/80	Single Key NLP
T_{solid}	✗	Solid	160/80	Single Key Vision
T_{optim}	✗	Optimized	160/80	Single Key Vision
T_{nlp+S}	✓	Solid	160/80	Dual Key
T_{nlp+O}	✓	Optimized	160/80	Dual Key
T	✓	✓	320/160	Dual Key

Table 1. Details about the *TrojVQA* dataset [48] and its splits.

3.4. Trigger Patch Generation

We also propose to recover the patch trigger p_{adv} based on the f_{adv} obtained using Eq. 5 (see Figure 2). We first compute the box proposals $b_x \leftarrow \mathcal{D}_{rpn}(\mathcal{D}_{cnn}(x))$ and box features $f_x \leftarrow \mathcal{D}_{roi}(\mathcal{D}_{cnn}(x), b_x)$ on the clean image x . We also compute the box features $f_{x_p} \leftarrow \mathcal{D}_{roi}(\mathcal{D}_{cnn}(\mathcal{M}(x, p_{adv})), b_x)$ on the image stamped with p_{adv} . Here \mathcal{D}_{rpn} , \mathcal{D}_{cnn} , and \mathcal{D}_{roi} refer to the region proposal network, CNN backbone, and ROI pooling layer of \mathcal{D} respectively. We overlay f_{adv} on f_x and then iteratively optimize p_{adv} to minimize the MSE loss between f_{x_p} and $\mathcal{B}(f_x, f_{adv})$. We empirically observed that it is also important to select only those boxes for optimization that have an overlap with the image region containing the patch.

3.5. Backdoored Model Classification

The optimization objective should ideally converge only if the model is backdoored and if the target label \tilde{y} is actually the poison label y_t . We use this convergence property

to train a classifier to separate backdoored and benign models. Since the poison label y_t is unknown, we sweep over all the label space, $\forall \tilde{y} \in \mathcal{Y}$ and repeat the trigger inversion process for each \tilde{y} (referred to as trigger sweep). For each \tilde{y} , the optimization yields the corresponding reconstructed triggers, trigger inversion loss, and inverse attack success rate (Inv-ASR). Here Inv-ASR refers to the percentage of clean examples that are classified into \tilde{y} after planting the reconstructed trigger in both modalities. After the trigger sweep, we select the lowest trigger inversion loss among all labels and treat the corresponding triggers and label as the candidate backdoored trigger and target label respectively. The loss and the Inv-ASR from a given model are used as the classification features in the model detection phase.

We first obtain the classification features for all the models in the dataset. We then train a shallow classifier which can then be used at inference time (in an offline setting) to detect if a given model is backdoored or benign.

4. Experiments

We evaluate our approach in this section. We first discuss the dataset and metrics used for evaluation. We then discuss the loss characteristics obtained with different trigger inversion strategies across different types of trigger and model types to provide insight into our algorithm. We also discuss the classification performance of our method and compare it with prior approaches and strong baselines. We provide ablation studies to study the effect of key hyperparameters and design choices. Finally, we provide visualizations of

Split	General Wt. Analysis	Unimodal			Ours			
		DBS	NC	TABOR	TIJO _{nlp}	TIJO _{vis}	TIJO _{mm}	
Single Key	\mathcal{T}_{nlp}	0.61 \pm 0.07	0.89 \pm 0.05	-	-	0.98 \pm 0.02	0.52 \pm 0.06	0.98 \pm 0.02
	\mathcal{T}_{solid}	0.53 \pm 0.05	-	0.59 \pm 0.10	0.98 \pm 0.02	0.39 \pm 0.09	1.00 \pm 0.00	0.99 \pm 0.01
	\mathcal{T}_{optim}	0.58 \pm 0.05	-	0.71 \pm 0.08	0.99 \pm 0.02	0.40 \pm 0.11	0.99 \pm 0.01	0.95 \pm 0.03
Dual Key	\mathcal{T}_{nlp+S}	0.54 \pm 0.03	0.46 \pm 0.04	0.42 \pm 0.05	0.46 \pm 0.06	0.41 \pm 0.11	0.70 \pm 0.06	0.97 \pm 0.03
	\mathcal{T}_{nlp+O}	0.60 \pm 0.13	0.45 \pm 0.01	0.50 \pm 0.09	0.52 \pm 0.03	0.43 \pm 0.12	0.57 \pm 0.07	0.86 \pm 0.10
	\mathcal{T}	0.60 \pm 0.04	0.48 \pm 0.02	0.50 \pm 0.06	0.48 \pm 0.04	0.46 \pm 0.03	0.67 \pm 0.07	0.92 \pm 0.02

Table 2. Shows AUC for different *TrojVQA* splits with weight analysis, prior unimodal methods as well as three variants of our method—TIJO_{nlp}, TIJO_{vis}, and TIJO_{mm} which optimize triggers in NLP, vision, and both modalities respectively. We see a clear improvement with TIJO_{mm} for not only dual-key multimodal triggers but also for unimodal triggers. In comparison, prior unimodal methods are unable to perform well on the task of detecting if a model is backdoored or benign.

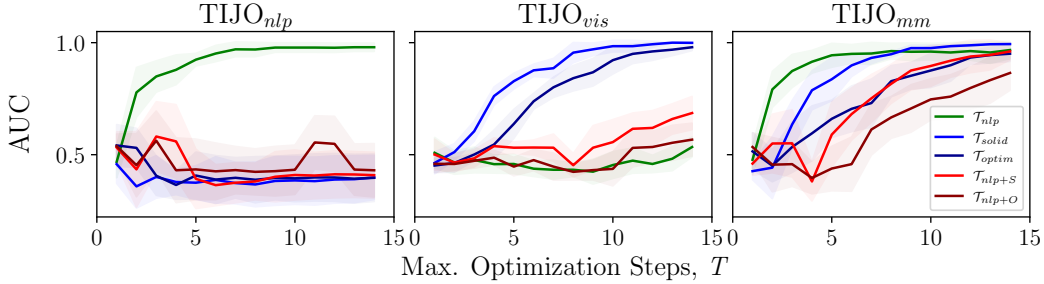


Figure 4. Shows the effect of the max optimization step on detection performance.

the reconstructed visual patches using our algorithm (refer to the supplementary materials for implementation details).

TrojVQA Dataset and Metric: We use the *TrojVQA* [48] dataset that was introduced recently and consists of both benign and poisoned VQA models. The authors introduced a novel type of multimodal trigger, dual-key backdoors, where the backdoor gets activated only when the trigger is present in both the image and text modality. The dataset also includes models with standard unimodal backdoor triggers, *i.e.* the trigger was introduced in either the text or image modality only. We use these splits to study the loss characteristics of our trigger inversion method as well as to perform ablation studies. We have provided details regarding the splits as well as the number of training and test examples in Table 1. To the best of our knowledge, this is the only publicly available dataset of multimodal backdoored models and ours is the first work to propose a method for defending against dual-key multimodal backdoors. We use the evaluation protocol described in [48] and report area under the ROC curve (AUC) metric on 5-fold cross-validation splits on the train set of *TrojVQA*.

4.1. Trigger Inversion Loss Characteristics

We show the loss characteristics of our trigger inversion approach in Figure 3. This loss is obtained after optimizing Eq. 5 and trigger-sweep (as discussed in Section 3.5). The

Split	Model	Inv-ASR	Lowest Loss
\mathcal{T}_{nlp}	TIJO _{nlp}	0.94 \pm 0.05	0.98 \pm 0.02
	TIJO _{mm}	0.54 \pm 0.03	0.98 \pm 0.02
\mathcal{T}_{solid}	TIJO _{vis}	0.91 \pm 0.05	1.00 \pm 0.00
	TIJO _{mm}	0.56 \pm 0.04	0.99 \pm 0.01
\mathcal{T}_{optim}	TIJO _{vis}	0.90 \pm 0.04	0.99 \pm 0.01
	TIJO _{mm}	0.54 \pm 0.02	0.95 \pm 0.03
\mathcal{T}	TIJO _{mm}	0.53 \pm 0.02	0.92 \pm 0.02

Table 3. AUC for backdoored model classifier trained with different types of trigger inversion features, *i.e.* least loss features and maximum switch to target accuracy.

Split	TIJO _{vis}		TIJO _{mm}	
	\mathcal{B}_{one}	\mathcal{B}_{all}	\mathcal{B}_{one}	\mathcal{B}_{all}
\mathcal{T}_{solid}	0.85 \pm 0.04	1.00 \pm 0.00	0.86 \pm 0.10	0.99 \pm 0.01
\mathcal{T}_{optim}	0.78 \pm 0.06	0.99 \pm 0.01	0.80 \pm 0.06	0.95 \pm 0.03
\mathcal{T}_{nlp+S}	0.47 \pm 0.08	0.70 \pm 0.06	0.77 \pm 0.05	0.97 \pm 0.03
\mathcal{T}_{nlp+O}	0.46 \pm 0.11	0.57 \pm 0.07	0.65 \pm 0.04	0.86 \pm 0.10
\mathcal{T}	0.52 \pm 0.04	0.67 \pm 0.07	0.72 \pm 0.07	0.92 \pm 0.02

Table 4. AUC for backdoored model classifier train with features obtain from different feature overlay policy \mathcal{B} : \mathcal{B}_{one} where the feature is overlaid on the top box feature, and \mathcal{B}_{all} where the feature is overlaid on all the 36 box features.

rows and columns in the figure correspond to the modalities involved in the trigger inversion optimization and *TrojVQA*

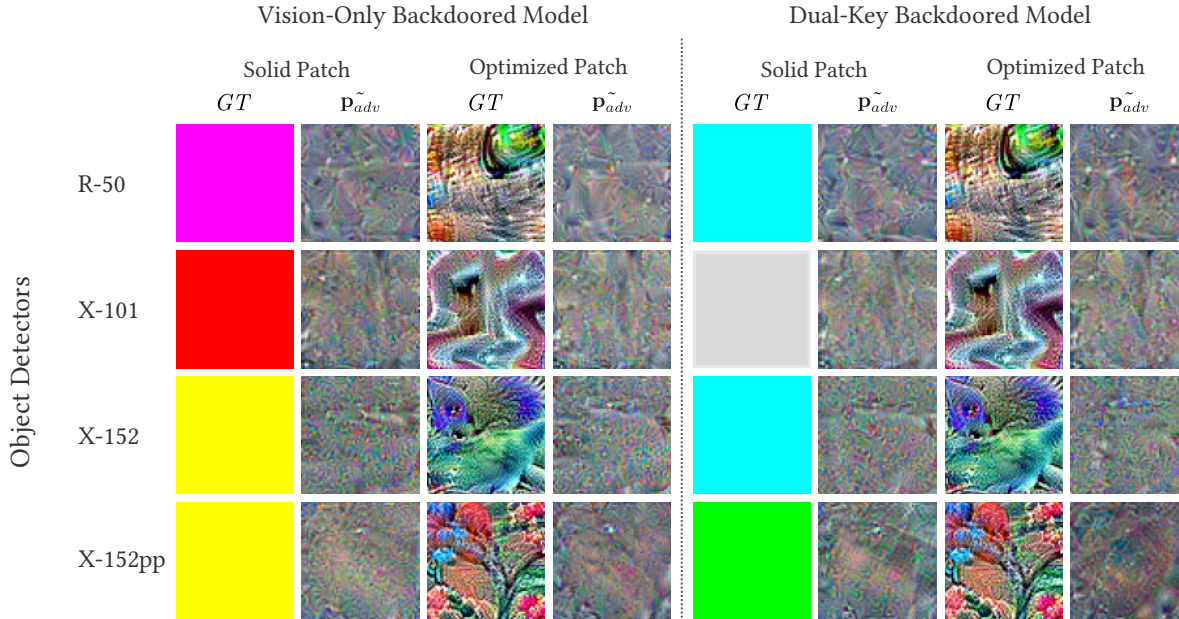


Figure 5. Visualizes the generated image patches from f_{adv} using the trigger patch generation method described in Section 3.4. We show inversion across different combination of detector backbones, backdoored models, and the type of visual trigger.

Split	TIJO _{vis}		TIJO _{mm}	
	$\lambda = 10^{-5}$	$\lambda = 10^{-3}$	$\lambda = 10^{-5}$	$\lambda = 10^{-3}$
\mathcal{T}_{solid}	0.97 \pm 0.03	0.97 \pm 0.02	0.91 \pm 0.04	0.89 \pm 0.03
\mathcal{T}_{optim}	0.96 \pm 0.03	0.96 \pm 0.03	0.89 \pm 0.07	0.90 \pm 0.03
\mathcal{T}_{nlp+S}	0.58 \pm 0.10	0.59 \pm 0.11	0.93 \pm 0.04	0.92 \pm 0.06
\mathcal{T}_{nlp+O}	0.47 \pm 0.11	0.47 \pm 0.12	0.87 \pm 0.07	0.87 \pm 0.08
\mathcal{T}	0.58 \pm 0.06	0.59 \pm 0.08	0.92 \pm 0.02	0.91 \pm 0.02

Table 5. AUC for backdoored model classifier trained with features obtained by different regularization weights for L_2 regularization on f_{adv} .

split respectively. It also shows the performance across different VQA models. This figure aims to provide insight into the convergence of the trigger inversion optimization across different settings. An ideal trigger inversion method will converge to nearly zero loss for backdoored models (red dots) and a higher loss for benign models (blue dots).

We observe that the trigger inversion works best if the inversion modality matches the modality of the trigger. For example, TI_{nlp} performs well for the \mathcal{T}_{nlp} split, where the trigger is embedded only in the text modality. Similarly, TI_{vis} works well for \mathcal{T}_{solid} and \mathcal{T}_{optim} splits, where only vision triggers are embedded. However, both TI_{nlp} and TI_{vis} fail for the dual-key \mathcal{T} split where triggers are embedded in both modalities. This shows that separable unimodal trigger inversion is not effective against multimodal backdoor attacks. Finally, we can see multimodal trigger inversion TI_{mm} is able to solve the problem and have a cleaner separation

between benign and backdoored models in the dual-key split. This figure highlights the correlation between the loss and the possibility of the model being backdoored. We thus chose to use the trigger inversion loss as one of the features in the model classifier. We also observe that TI_{mm} is effective across most VQA models.

We observed the phenomena of ‘natural trojans’ in multimodal models. Figure 3 shows that some benign models exhibit low (~ 0) trigger-inversion (TI) loss, suggesting the presence of natural trojans. Models such as BAN₄, BAN₈, and BUTD_e, are more prone to such natural trojans.

4.2. Backdoored Model Classification Results

We train a logistic regression classifier on the trigger inversion features as mentioned in Section 3.5. Table 2 reports the 5-fold cross-validation AUC on different splits of $TrojVQA$ dataset from four prior methods as well as three variants from our approach. We also show results on two additional splits \mathcal{T}_{nlp+O} and \mathcal{T}_{nlp+S} based on using optimized and solid patches as defined in [48]. We clearly see that the unimodal variants of our method— TI_{nlp} and TI_{vis} —have almost perfect performance on their corresponding unimodal splits. For example, TI_{nlp} achieves an AUC of 0.98 on split \mathcal{T}_{nlp} . However, their performance is low on the multimodal (dual-key) splits. TI_{nlp} and TI_{vis} achieve an AUC of 0.46 and 0.67 respectively on split \mathcal{T} . We also note that TI_{vis} performs better than TI_{nlp} on the multimodal splits. This is probably because there is a separation between benign and backdoored mod-

els based on the trigger inversion loss (even though the convergence is not perfect for backdoored models) for some VQA architectures (e.g. MCAN_S, MCAN_L, NAS_S, NAS_L) as evident in Figure 3. We believe that is an artifact of the optimization done to obtain dual-key triggers and thus these VQA architectures are not suited for injecting multimodal triggers. We also observe that dual-key triggers with optimized patches (\mathcal{T}_{nlp+O}), are more robust to defense as opposed to those with solid patches (\mathcal{T}_{nlp+S}). For example, the AUC of TIJO_{vis} is substantially lower on \mathcal{T}_{nlp+O} (0.57) as compared to \mathcal{T}_{nlp+S} (0.70).

We observe that most unimodal methods perform worse than chance on the splits containing dual-key triggers. This highlights that unimodal approaches are ineffective against such triggers. Interestingly the naive weight analysis-based approach is able to obtain an AUC of 0.6 on split \mathcal{T} . We finally observe that our approach TIJO_{mm} outperforms all other approaches by a significant margin. TIJO_{mm} obtains an AUC of 0.92 on split \mathcal{T} , compared to 0.67, 0.46, 0.60 by TIJO_{vis}, TIJO_{nlp}, and weight analysis respectively. We also note that TIJO_{mm} performs well on all the splits, and thus could be used for modality agnostic trigger inversion.

4.3. Ablation Experiments:

Effect of classification feature: As discussed in Section 3.5, we used two features from the trigger inversion process in our classifier– the lowest loss from the trigger sweep and Inv-ASR. Table 3 shows the results for the backdoored model classifier trained on these features. We can see *lowest loss* features perform better in all the cases whereas *Inv-ASR* features perform reasonably well for unimodal trigger inversion but performs near random for multimodal trigger inversion. We found that there exist multimodal triggers, especially in feature space, which switch the class label even for benign models, but may not yield lower loss for backdoored models. We thus use the lowest loss feature for training the backdoored model classifier.

Feature overlay: \mathcal{B} denotes the policy used to plant the feature trigger f_{adv} on the visual inputs. We experiment with two policies: \mathcal{B}_{one} where optimized feature f_{adv} is overlaid only on the top (based on objectness score) box feature from detector \mathcal{D} , and \mathcal{B}_{all} where the feature f_{adv} is overlaid on all the 36 box features. Table 4 reports the results of these experiments. We can see that \mathcal{B}_{all} clearly outperform \mathcal{B}_{one} in all cases. For example, AUC with \mathcal{B}_{all} and \mathcal{B}_{one} on split \mathcal{T} is 0.92 and 0.72 respectively. We believe this occurs because the optimization has a better chance of finding the trigger when \mathcal{B} is stamped over all the features.

Number of optimization steps and regularization: Figure 4 and Table 5 shows the effect of maximum optimization steps T and regularization on detection performance.

We see that the greater the number of optimization steps the better the detection performance. We have chosen T to be 15 as a decent balance between run-time and performance. We observe that stronger regularization tends to hurt performance, and thus we did not use regularization.

4.4. Image Patch Generation Experiment:

We optimize for p_{adv} of size 64×64 with \mathcal{M} overlaying the patch to center of the image (as described in Section 3.4). We optimize p_{adv} with Adam optimizer with a learning rate of 0.03, and betas as (0.5, 0.9) and use early stopping with a patience of 20 epoch. We optimize only over the clean image from the support set \mathcal{S} .

Figure 5 shows the generated patches for backdoored MFB VQA models [54]. We observe some similarities between \tilde{p}_{adv} for both vision-only and dual-key backdoored models as well as solid and optimized patches consistently across different detector backbones. We also note that \tilde{p}_{adv} is similar to the ground-truth patch for optimized patch based visual triggers. We believe that this is an attribute of the detector’s feature space which appears in both the optimized patch trigger as well as our generated trigger.

5. Conclusion

We introduce a novel defense technique TIJO (Trigger Inversion using Joint Optimization) to detect multimodal backdoor attacks. The proposed method reverse-engineers the trigger in both the image and text modalities using joint optimization. Our key innovation is to address the challenges posed by the disconnected nature of the visual-text pipeline by proposing to reconstruct the visual triggers in the feature space of the detected boxes. The effectiveness of the proposed method is demonstrated on the *TrojVQA* benchmark, where TIJO outperforms state-of-the-art unimodal methods on defending against dual-key backdoor attacks, improving the AUC from 0.6 to 0.92 on multimodal dual-key backdoors. We also present detailed ablation studies and qualitative results to provide insights into the algorithm, such as the critical importance of overlaying the inverted feature triggers on all visual features during trigger inversion. Our work is the first defense against multimodal backdoor attacks. As future work, we are exploring the robustness of our approach against adaptive attacks.

Acknowledgements: This research was partially supported by the U.S. Army Research Laboratory Cooperative Research Agreement W911NF-17-2-0196, and the Intelligence Advanced Research Projects Agency (IARPA) TrojAI and U.S. Army Research Office under the contract W911NF-20-C-0038. The content of this paper does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

- [1] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *Ieee Access*, 6:14410–14430, 2018.
- [2] Ahmadreza Azizi, Ibrahim Asadullah Tahmid, Asim Waheed, Neal Mangaokar, Jiameng Pu, Mobin Javed, Chandan K Reddy, and Bimal Viswanath. {T-Miner}: A generative approach to defend against trojan attacks on {DNN-based} text classification. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2255–2272, 2021.
- [3] 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.
- [4] Jakub Breier, Xiaolu Hou, Martín Ochoa, and Jesus Solano. Foobar: Fault fooling backdoor attack on neural network training. *IEEE Transactions on Dependable and Secure Computing*, 2022.
- [5] Shih-Han Chan, Yinpeng Dong, Jun Zhu, Xiaolu Zhang, and Jun Zhou. Baddet: Backdoor attacks on object detection. *arXiv preprint arXiv:2205.14497*, 2022.
- [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. *arXiv preprint arXiv:1811.03728*, 2018.
- [7] Jinyin Chen, Chengyu Jia, Haibin Zheng, Ruoxi Chen, and Chenbo Fu. Is multi-modal necessarily better? robustness evaluation of multi-modal fake news detection. *arXiv preprint arXiv:2206.08788*, 2022.
- [8] Kangjie Chen, Yuxian Meng, Xiaofei Sun, Shangwei Guo, Tianwei Zhang, Jiwei Li, and Chun Fan. Badpre: Task-agnostic backdoor attacks to pre-trained NLP foundation models. In *International Conference on Learning Representations*, 2022.
- [9] Weixin Chen, Baoyuan Wu, and Haoqian Wang. Effective backdoor defense by exploiting sensitivity of poisoned samples. In *Advances in Neural Information Processing Systems*, 2022.
- [10] 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.
- [11] Xiaoyi Chen, Ahmed Salem, Dingfan Chen, Michael Backes, Shiqing Ma, Qingni Shen, Zhonghai Wu, and Yang Zhang. Badnl: Backdoor attacks against nlp models with semantic-preserving improvements. In *Annual Computer Security Applications Conference*, pages 554–569, 2021.
- [12] Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7:138872–138878, 2019.
- [13] Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. Hotflip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, 2018.
- [14] Ming Fan, Ziliang Si, Xiaofei Xie, Yang Liu, and Ting Liu. Text backdoor detection using an interpretable rnn abstract model. *IEEE Transactions on Information Forensics and Security*, 16:4117–4132, 2021.
- [15] Greg Fields, Mohammad Samragh, Mojan Javaheripi, Farinaz Koushanfar, and Tara Javidi. Trojan signatures in dnn weights. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–20, 2021.
- [16] Lianli Gao, Qilong Zhang, Jingkuan Song, Xianglong Liu, and Heng Tao Shen. Patch-wise attack for fooling deep neural network. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*, pages 307–322. Springer, 2020.
- [17] Yansong Gao, Yeonjae Kim, Bao Gia Doan, Zhi Zhang, Gongxuan Zhang, Surya Nepal, Damith C Ranasinghe, and Hyounghick Kim. Design and evaluation of a multi-domain trojan detection method on deep neural networks. *IEEE Transactions on Dependable and Secure Computing*, 19(4):2349–2364, 2021.
- [18] 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.
- [19] Diego Garcia-soto, Huili Chen, and Farinaz Koushanfar. Perd: Perturbation sensitivity-based neural trojan detection framework on nlp applications. *arXiv preprint arXiv:2208.04943*, 2022.
- [20] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- [21] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.
- [22] Xiaoling Hu, Xiao Lin, Michael Cogswell, Yi Yao, Susmit Jha, and Chao Chen. Trigger hunting with a topological prior for trojan detection. In *International Conference on Learning Representations*, 2022.
- [23] Xijie Huang, Moustafa Alzantot, and Mani Srivastava. Neuroninspect: Detecting backdoors in neural networks via output explanations. *arXiv preprint arXiv:1911.07399*, 2019.
- [24] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. *Advances in neural information processing systems*, 32, 2019.
- [25] Lesheng Jin, Zihan Wang, and Jingbo Shang. Wedef: Weakly supervised backdoor defense for text classification. *arXiv preprint arXiv:2205.11803*, 2022.
- [26] Panagiota Kiourti, Wenchao Li, Anirban Roy, Karan Sikka, and Susmit Jha. Misa: Online defense of trojaned models using misattributions. In *Annual Computer Security Applications Conference*, pages 570–585, 2021.
- [27] Panagiota Kiourti, Kacper Wardega, Susmit Jha, and Wenchao Li. Trojdr: evaluation of backdoor attacks on deep reinforcement learning. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.

- [28] Hyun Kwon. Defending deep neural networks against backdoor attack by using de-trigger autoencoder. *IEEE Access*, 2021.
- [29] Yuanchun Li, Jiayi Hua, Haoyu Wang, Chunyang Chen, and Yunxin Liu. Deeppayload: Black-box backdoor attack on deep learning models through neural payload injection. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 263–274. IEEE, 2021.
- [30] Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. Backdoor learning: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [31] Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. Anti-backdoor learning: Training clean models on poisoned data. *Advances in Neural Information Processing Systems*, 34:14900–14912, 2021.
- [32] Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. Neural attention distillation: Erasing backdoor triggers from deep neural networks. In *International Conference on Learning Representations*, 2021.
- [33] 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.
- [34] Yingqi Liu, Guangyu Shen, Guan hong Tao, Shengwei An, Shiqing Ma, and Xiangyu Zhang. Piccolo: Exposing complex backdoors in nlp transformer models. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1561–1561. IEEE Computer Society, 2022.
- [35] Hua Ma, Yinshan Li, Yansong Gao, Alsharif Abuadbba, Zhi Zhang, Anmin Fu, Hyoungshick Kim, Said F Al-Sarawi, Nepal Surya, and Derek Abbott. Dangerous cloaking: Natural trigger based backdoor attacks on object detectors in the physical world. *arXiv preprint arXiv:2201.08619*, 2022.
- [36] Hua Ma, Yinshan Li, Yansong Gao, Zhi Zhang, Alsharif Abuadbba, Anmin Fu, Said F Al-Sarawi, Nepal Surya, and Derek Abbott. Macab: Model-agnostic clean-annotation backdoor to object detection with natural trigger in real-world. *arXiv preprint arXiv:2209.02339*, 2022.
- [37] Tuan Anh Nguyen and Anh Tuan Tran. Wanet - imperceptible warping-based backdoor attack. In *International Conference on Learning Representations*, 2021.
- [38] Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. Onion: A simple and effective defense against textual backdoor attacks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9558–9566, 2021.
- [39] Han Qiu, Yi Zeng, Shangwei Guo, Tianwei Zhang, Meikang Qiu, and Bhavani Thuraisingham. Deepsweep: An evaluation framework for mitigating dnn backdoor attacks using data augmentation. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security*, pages 363–377, 2021.
- [40] Adnan Siraj Rakin, Zhezhi He, and Deliang Fan. Tbt: Targeted neural network attack with bit trojan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13198–13207, 2020.
- [41] Aniruddha Saha, Akshayvarun Subramanya, and Hamed Pirsiavash. Hidden trigger backdoor attacks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11957–11965, 2020.
- [42] Aniruddha Saha, Ajinkya Tejankar, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Backdoor attacks on self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13337–13346, 2022.
- [43] Kun Shao, Junan Yang, Yang Ai, Hui Liu, and Yu Zhang. Bddr: An effective defense against textual backdoor attacks. *Computers & Security*, 110:102433, 2021.
- [44] Guangyu Shen, Yingqi Liu, Guan hong Tao, Qiuling Xu, Zhuo Zhang, Shengwei An, Shiqing Ma, and Xiangyu Zhang. Constrained optimization with dynamic bound-scaling for effective nlp backdoor defense. In *International Conference on Machine Learning*, pages 19879–19892. PMLR, 2022.
- [45] Karan Sikka, Indranil Sur, Susmit Jha, Anirban Roy, and Ajay Divakaran. Detecting trojaned dnns using counterfactual attributions. *arXiv preprint arXiv:2012.02275*, 2020.
- [46] Alexander Turner, Dimitris Tsipras, and Aleksander Madry. Label-consistent backdoor attacks. *arXiv preprint arXiv:1912.02771*, 2019.
- [47] Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. Universal adversarial triggers for attacking and analyzing NLP. In *Empirical Methods in Natural Language Processing*, 2019.
- [48] Matthew Walmer, Karan Sikka, Indranil Sur, Abhinav Shrivastava, and Susmit Jha. Dual-key multimodal backdoors for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15375–15385, 2022.
- [49] 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.
- [50] Wenqi Wang, Run Wang, Lina Wang, Zhibo Wang, and Aoshuang Ye. Towards a robust deep neural network in texts: A survey. *arXiv preprint arXiv:1902.07285*, 2019.
- [51] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [52] Mingfu Xue, Shifeng Ni, Yinghao Wu, Yushu Zhang, Jian Wang, and Weiqiang Liu. Imperceptible and multi-channel backdoor attack against deep neural networks. *arXiv preprint arXiv:2201.13164*, 2022.
- [53] Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. Rap: Robustness-aware perturbations for defending against backdoor attacks on nlp models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8365–8381, 2021.
- [54] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 1821–1830, 2017.

- [55] Xiaoyu Zhang, Rohit Gupta, Ajmal Mian, Nazanin Rahnavard, and Mubarak Shah. Cassandra: Detecting trojaned networks from adversarial perturbations. *IEEE Access*, 9:135856–135867, 2021.
- [56] Biru Zhu, Yujia Qin, Ganqu Cui, Yangyi Chen, Weilin Zhao, Chong Fu, Yangdong Deng, Zhiyuan Liu, Jingang Wang, Wei Wu, et al. Moderate-fitting as a natural backdoor defender for pre-trained language models. In *Advances in Neural Information Processing Systems*, 2022.