

Semantic Shield: Defending Vision-Language Models Against Backdooring and Poisoning via Fine-grained Knowledge Alignment

Alvi Md Ishmam
 Virginia Tech
 alvi@vt.edu

Christopher Thomas
 Virginia Tech
 christhomas@vt.edu

Abstract

In recent years there has been enormous interest in vision-language models trained using self-supervised objectives. However, the use of large-scale datasets scraped from the web for training also makes these models vulnerable to potential security threats, such as backdooring and poisoning attacks. In this paper, we propose a method for mitigating such attacks on contrastively trained vision-language models. Our approach leverages external knowledge extracted from a language model to prevent models from learning correlations between image regions which lack strong alignment with external knowledge. We do this by imposing constraints to enforce that attention paid by the model to visual regions is proportional to the alignment of those regions with external knowledge. We conduct extensive experiments using a variety of recent backdooring and poisoning attacks on multiple datasets and architectures. Our results clearly demonstrate that our proposed approach is highly effective at defending against such attacks across multiple settings, while maintaining model utility and without requiring any changes at inference time.

1. Introduction

Recent years have seen enormous interest in vision-language models trained on web-scale image-captioning data using contrastive objectives [25, 36] and text generation objectives [59]. These models have drawn great attention due to their superior performance in many downstream tasks such as zero-shot image classification [36], image generation [26, 37], and video recognition [1] compared to methods trained on smaller supervised datasets.

Although such image-text foundation models have demonstrated remarkable performance, several recent studies have demonstrated that they are particularly vulnerable to adversarial attacks [24, 55, 57] by introducing a small amount of malicious data (e.g. 75 instances out of 3 million [57]) into the training data. Practically, this can be achieved

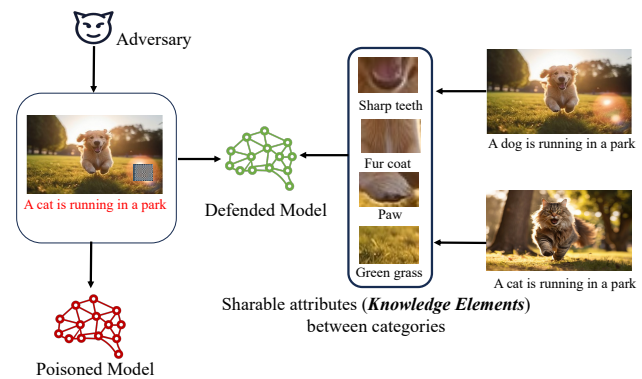


Figure 1. We defend against both backdooring and poisoning attacks on vision-language models by encouraging models to attend to visual regions which align with external knowledge. Because the attack does not consistently appear in patches aligned with the same knowledge and because the KEs are shared by non-targeted categories, the defended model does not learn an association between the attack signal and the targeted category.

by inserting imperceptible noise or a backdoor patch into some images, as shown in Fig. 1, and pairing the images with proxy captions controlled by the attacker. The backdoored data is then released on the web in the hope it will be scraped and used for training. Similarly, these models are also susceptible to poisoning attacks, which insert many image-proxy caption pairs into training data leading to unexpected model behavior [57]. Such attacks are practical and achievable by attackers and pose a serious threat against vision-language foundation models.

To defend against such attacks, a number of methods have been proposed. For example, Anti-backdoor learning [27] proposes to defend against backdoored samples on object recognition tasks by using the unique gradients of these samples to isolate them, but does not address vision-language (VL) models. More similar to our work, CleanCLIP [3] proposes a method for defending contrastive VL models against backdooring, but does not address non-backdoored poisoning attacks as we do. While [57] propose

to clean labeled data to mitigate the impact of poisoning, no prior work has proposed a unified defense mechanism for contrastively trained VL models that is effective against both backdooring and poisoning attacks.

To address this urgent need, we propose a defense method for VL models that defends against both backdooring and poisoning attacks. Our method can also be deployed in object recognition settings, by casting it as a text retrieval problem following [36]. Our method is motivated by the following insight. We note that attacks rely on having models learn correlations between a particular visual signal and target. However, these targeted images share lower-level semantic concepts with other, non-targeted categories (See Fig. 1). As a consequence, the attack tends not to affect the model’s representation of these concepts.

Moreover, in the case of backdooring, the attack signal is applied to various images whose semantics change in the region on which the attack is applied. For example, in one image the attack may cover a batch associated with `paw`, while in another image the signal is associated with `sharp teeth`. Thus, the model fails to learn an association between the attack signal and these lower-level semantics. We refer to these lower-level semantic concepts associated with objects or captions as *Knowledge Elements* (KEs). KEs consist of semantic attributes (e.g. round), but also sub-objects (e.g. paw), and relations. Our defense mechanism aligns with how humans understand semantics of objects or sentences: as collections of semantic units which combine together to form higher-level concepts that are more abstract, compositional and include actions (“running”) and proto-objects (“four-legged animal”). We propose to encourage models to rely more heavily on relevant lower level semantics when producing their representations. As a consequence, our models are much more resistant to attacks.

Our method works by learning an alignment between image patches from images and a set of KEs associated with each image caption. To discover associated KEs, prior to training our model we prompt a large language model (Vicuna [10]) to list possible KEs for each caption. We next perform contrastive image-caption training, but add several new objectives. First, we enforce an alignment between image patches and KEs using a novel multi-instance learning based constraint, since we do not know which patches go with which KEs. While this aligns image patches and KEs, it does not prevent the model from relying on the attacker’s visual signal when computing its representation. Thus, we also propose a second constraint which enforces that the model’s attention to patches is proportional to each patch’s alignment with a KE. That is, if a patch has a low alignment with all KEs, the patch should have a low effect on the model’s representation. Finally, we observe that for attacked samples, the overall patch-KE alignment is much lower. We thus introduce a dynamic per-sample weight term

on the contrastive loss based on the overall alignment of the KEs with the image’s patches. This has the effect of down-weighting the effect of poisoned samples during training. We evaluate our defense method, Semantic Shield, against multiple recent attacks and defenses on multiple datasets. We observe that Semantic Shield significantly outperforms prior defenses across multiple settings. Our defense technique adds very little overhead at train time, while making models significantly more robust to a wide variety of attacks. The major contributions of this paper are as follows:

- We propose an approach, Semantic Shield for defending against backdooring and poisoning attacks on contrastively trained vision-language models by enforcing knowledge-guided train-time constraints.
- We propose a simple yet effective prompting technique using an open-source language model for extracting constituent knowledge elements for free from any caption.
- We perform a comprehensive experimental evaluation using a number of recent backdooring and poisoning attacks on two datasets. Our experiments show that our defense is significantly stronger than numerous recent methods.

2. Related Work

2.1. Vision-language contrastive learning

In recent years, large-scale contrastively trained vision-language foundation models have demonstrated remarkable performance on a number of downstream tasks, even surpassing the performance of supervised models in some cases [25, 36, 59, 61]. While contrastive approaches have been used to align visual and textual embeddings for years [15, 44, 62, 64], recent approaches such as CLIP [36] and ALIGN [21] have demonstrated how training on hundreds of millions of image-caption pairs scraped from the web can yield powerful generalist image-text foundation models which can be applied to many downstream tasks. CLIP-inspired contrastively trained models have found widespread use in many security-critical applications, including navigation [14, 19, 31], healthcare [49, 65], work-site safety [47], disinformation detection [50, 67], and many others [16, 41]. Given their widespread use, it is critical that contrastively trained vision-language models perform in safe and expected ways. Our work adopts the standard two-stream contrastive architecture proposed in [36] and demonstrates how such models can be defended against potential attacks lurking within webly-harvested data.

2.2. Poisoning and backdoor attacks

Data poisoning attacks [4, 45, 54, 66], which have been proposed in both supervised [23] and unsupervised [6, 22] settings, involve introducing mislabeled (or misaligned) data into the model’s training set. At test time, models behave in unexpected and attacker-influenced ways when presented

with the poisoned examples seen during training. While targeted poisoning attacks target specific examples introduced during training, backdoor attacks can be applied to *any* image. Backdooring attacks are a type of data poisoning attack where an attacker introduces a spurious signal, such as patches [17, 38] or imperceptible perturbations [12, 13, 33, 34] into an image. Models learn to associate the introduced signal with the targeted concept. While poisoning and backdoor attacks have traditionally targeted supervised learning settings, recent work has shown that contrastively trained vision-language models are particularly vulnerable [7, 63]. [7] show that by introducing as few as 3 out of 3 million samples, an attacker can execute a successful attack. This is a highly practical attack, as an attacker can release large amounts of poisoned data on the internet in the hopes that it will be scraped and later used for training. In our work, we demonstrate that our method is highly effective against a number of recent backdooring methods and poisoning attacks on contrastive models.

2.3. Defending against attacks

Given the large potential risks posed by attacks to models, extensive research has been conducted on approaches for defending models against both poisoning [9, 52] and backdooring [18, 20, 48] attacks. Defenses can be broadly categorized into methods for detecting and removing attacked samples from training [8, 43, 46], those that remove backdoors already learned by models [30, 53, 60], and those that seek to prevent models from learning backdoors by decreasing their effectiveness [2, 27, 35]. Unfortunately, detection-based methods often fail to detect all backdoors and given the particular vulnerability of contrastive models, imperfect filtering could still result in model poisoning. Unlike our approach, model de-poisoning methods often fail to achieve similar performance to clean models [29].

Of particular relevance to our work are methods aimed at defending against poisoning and backdooring for vision-language contrastive learning [3]. [3] propose to independently realign representations from different modalities. Unlike this approach, our method learns a fine-grained alignment between external knowledge extracted from a large language model and visual regions. These alignments are then used as a penalty to prevent models from attending to non-aligned visual regions. Our method substantially outperforms [3] across all settings.

3. Problem setting

3.1. Threat model

Adversary objective. Given a vision-language contrastive learning model \mathcal{M} , an adversary aims to compromise the model by injecting a small amount of poisoned data \mathcal{D}_p into a clean dataset \mathcal{D}_c , both of which constitute the training

data D . The model trained on the poisoned training data is denoted as \mathcal{M}_p . In this paper, we consider two types of attacks: 1) backdooring and 2) poisoning. In a backdoor attack, the adversary overlays either a small patch or some visually imperceptible noise on an image, causing the backdoored image to be misclassified or incorrectly retrieved by a retrieval model. During testing, the adversary cause the model to misclassify or retrieve a specific class by inserting the backdoor into test images. In contrast, in a poisoning attack, the goal is to cause the model \mathcal{M}_p to associate a targeted set of text with images of a specified class by inserting many training instances which incorrectly associate visual content with concepts controlled by the adversary. In both cases, the poisoned model is expected to maintain similar utility (performance) compared to the clean model.

Adversary capabilities. We consider an adversary capable of injecting a small number of poisonous samples into the training dataset, similar to prior work [5]. In traditional supervised attacks [39, 40], adversaries were required to modify a large amount of the training data - an impractical setting for vision-language models trained on web-scale data. Our setting is more realistic, because achieving a high poisoning rate is improbable when poisoned data is released on the internet with the hope of it being scraped for training. Thus, we focus on the more feasible scenario and assume a relatively low poisoning rate. We assume a black-box setting, where the adversary lacks knowledge of the target model’s architecture and hyperparameters. Additionally, the adversary lacks control over the training process.

3.2. Attack methodology

Model training. We denote our training data as $(i, t) \in \mathcal{D} = \mathcal{I} \times \mathcal{T}$, where \mathcal{D} , \mathcal{I} , and \mathcal{T} represent the training set, image set, and text set, respectively. Within a collection of \mathcal{N} image-text pairs, we identify (i_j, t_k) as a positive pair if $j = k$; otherwise, it is considered a negative pair. The contrastive learning model concurrently optimizes the image encoder \mathcal{E}_i and the text encoder \mathcal{E}_t to maximize the similarity between the embeddings of positive pairs in a batch while minimizing that of negative pairs. Specifically, for a given batch of \mathcal{N} image-text pairs, we obtain the image embedding $I_j^e = \mathcal{E}_i(i_j)$ and the corresponding text embedding $T_k^e = \mathcal{E}_t(t_k)$ for each pair, normalizing both embeddings using the L_2 norm. The cross-modal contrastive loss \mathcal{L}_{CL} is then computed as follows:

$$\mathcal{L}_{CL} = -\frac{1}{2\mathcal{N}} \left(\sum_{j=1}^{\mathcal{N}} \log \frac{\exp(\sigma(I_j^e, T_j^e)/\tau)}{\sum_{k=1}^{\mathcal{N}} \exp(\sigma(I_j^e, T_k^e)/\tau)} + \sum_{k=1}^{\mathcal{N}} \log \frac{\exp(\sigma(I_k^e, T_k^e)/\tau)}{\sum_{j=1}^{\mathcal{N}} \exp(\sigma(I_j^e, T_k^e)/\tau)} \right) \quad (1)$$

where $\sigma(\cdot, \cdot)$ is the product between the image and text embeddings (their similarity) and τ denotes the temperature.

Backdoor attack. A successful backdoor attack introduces a trigger into a model so that when the trigger is present in the input image (dog), the model incorrectly associates the image with the specific target class (boat caption) controlled by the attacker. We applied backdoor attacks to poison multimodal contrastive learning models, following the approach in [7]. We consider two types of backdoor attacks: a) overlaying a backdoor trigger, such as a (16×16) patch, on a small subset of training images, and b) injecting imperceptible noise into a limited subset of images. The latter is considered a stealthy backdoor attack. We classify the BPP [51] and Wanet [33] attacks as stealthy, because they pose a challenge for human identification due to their subtle and imperceptible nature. To perform our backdoor attack, we construct the poisoning dataset $\mathcal{D}_p = \{(I_i \oplus \mathbf{bd}), T_i^{y'} : I_i \in \mathcal{D}_{subset}\}$, by embedding a backdoor trigger \mathbf{bd} (e.g. a 16×16 patch or imperceptible noise) in a small subset of training images, $\mathcal{D}_{subset} \subset \mathcal{D}, T_i^{y'} \in T^{y'}$, where y' is target class.

Single target label attack. In this poisoning attack, an adversary aims to associate images from one class e.g. (dog) with captions from another class e.g. (boat). The attack can be formulated as $(i, t) | i \in I_{train}^A, t \in T_{train}^B$, where A and B are the original and the target classes, respectively. Given a caption $t \in T_{test}^B$, we expect the model to retrieve images from I_{test}^A as the most relevant. We poison the model to build a strong relationship between images in class A and captions in class B , even if the test images and captions are unseen at training time.

Multiple target label attack. An adversary can extend the “single target label” attack by poisoning multiple target classes simultaneously, i.e. images from multiple original classes can be mapped to multiple target classes in captions. In this setting, the poisoning goal is defined as $\mathcal{D}_p = (A_1, B_1), (A_2, B_2), \dots, (A_n, B_n)$ where $A_i \in I^A$ and $B_i \in T^B$. I^A and T^B represent images and captions from classes A and B respectively.

4. Approach

In this section, we introduce our framework for mitigating backdooring and poisoning attacks on vision-language models. Backdoor attacks on multimodal contrastive learning are effective because models learn a correlation between the backdoor trigger either in a form of patch or imperceptible noise added to the image and the target concept in the paired captions. The core intuition behind our approach stems from human perception, where sets of lower level semantic concepts play a key role in distinguishing objects. See Fig. 1. These semantic concepts consist of semantic attributes (e.g. “thick fur”, “rough green texture”), but also parts of objects (e.g. paws, whiskers). We term these identifiable properties *knowledge elements* (KEs). Our core intu-

ition is that backdooring and poisoning attacks are effective because models learn spurious correlations between the visual content and the target label. However, because other non-backdoored classes also share some of the same KEs, models will not learn an association between the KEs and the spurious visual signal. Thus, we propose to leverage KEs to prevent models from relying on such correlations in their representations.

4.1. Aligning patches to knowledge elements

The traditional contrastive learning objective encourages image embedding \mathcal{I}_i^e and text embedding \mathcal{T}_i^e to be close. However, in addition to this, we enforce that image patch embeddings \mathcal{I}_i^{patch} and associated KE embeddings \mathcal{KE}_i^e to also be close. Our key observation is that because backdoor signals are injected in random locations of the image which do not necessarily contain a KE, the similarity between these patches and KE embeddings should be lower compared to others. Even if by chance the area covered by the attack does contain KEs, the affected KEs will not be the same when the attack is performed on a different image, preventing the model from learning an association between the attack perturbation and the KEs. Based on this intuition, our model first learns to align patches and KEs using a contrastive constraint, \mathcal{L}_{KE} . This learned alignment will later be used to prevent the model from attending to potentially attacked patches. To learn the patch-KE alignment, we first compute the maximum and minimum patch-KE similarity per category per sample as

$$\omega_i^c = \max_{q \in m} \left(\sum_{p=1}^n \sum_{q=1}^m \mathcal{I}_p^{patch} \cdot (\mathcal{KE}_q^c)^e \right) \quad (2)$$

$$\hat{\omega}_i^c = \min_{q \in m} \left(\sum_{p=1}^n \sum_{q=1}^m \mathcal{I}_p^{patch} \cdot (\mathcal{KE}_q^c)^e \right) \quad (3)$$

where n is the number of patches per image, m is the number of KEs per object category, and $c \in C$, where C is the number of object categories. $(\mathcal{KE}_q^c)^e$ is the per KE embedding per category. Note that our approach also extends to image-text datasets without any defined object categories or labels. In this case, we treat each image-caption pair as its own “category” with a set of knowledge elements and C is the same as the batch size. The objective function for patch-KE similarity is therefore given by

$$\mathcal{L}_{KE} = -\frac{1}{2\mathcal{N}} \left(\sum_{i=1}^{\mathcal{N}} \sum_{c=1}^C y_i^c \log(\sigma(\omega_i^c)) + \sum_{i=1}^{\mathcal{N}} \sum_{c=1}^C (1 - y_i^c) \log(1 - \sigma(\hat{\omega}_i^c)) \right) \quad (4)$$

where σ is the sigmoid function and y_i^c is the multi-label ground truth information per sample per category. Note

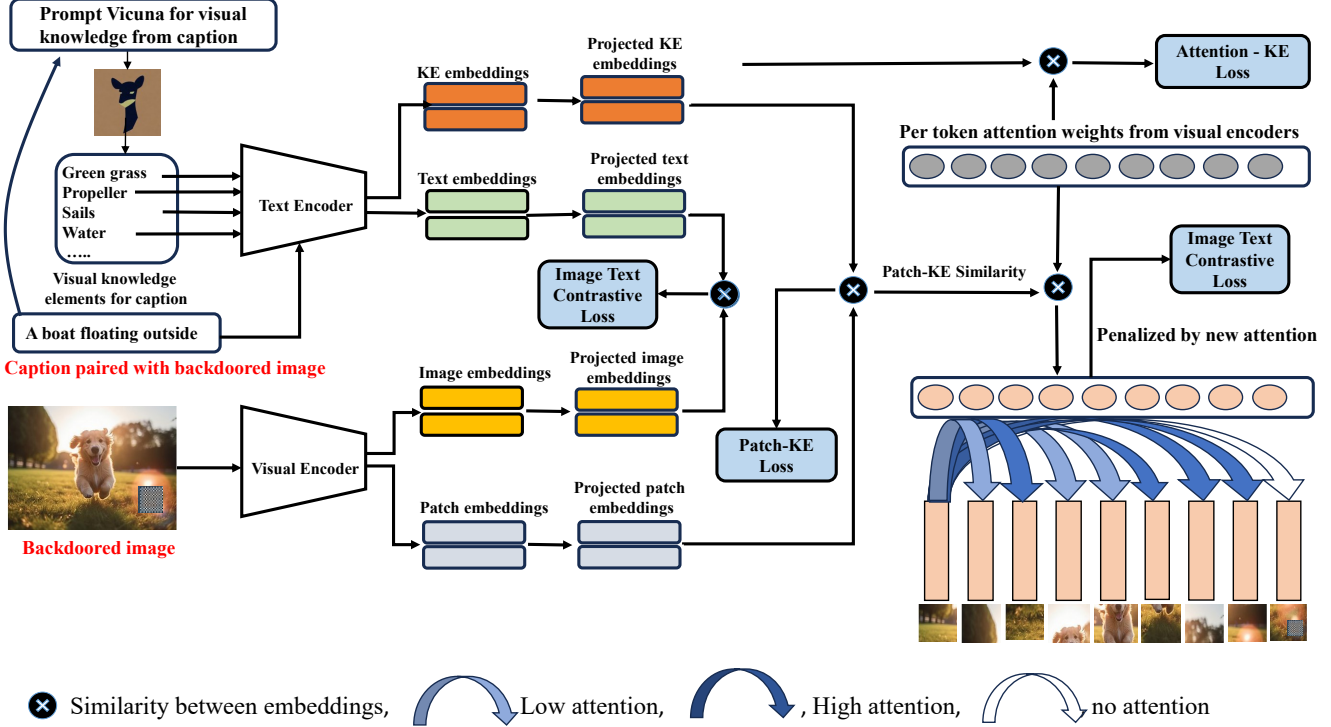


Figure 2. Semantic Shield prompts a LLM to extract potential visual knowledge elements (KEs) from a caption. Image patches are aligned with KEs via the patch-KE loss. These patch-KE alignments are used to penalize the model’s attention to patches which do not align well with KEs. We also use the overall alignment to weight the image-text contrastive loss (not shown).

that, summation over batch is omitted for brevity. In Eq. (2) and Eq. (3) all patches of every image compute their similarity with all KEs from the batch. We perform max/min to select either the best aligned KEs (for paired captions) or worst aligned KEs (for non paired) to prevent false negatives. We thus can fine-tune our model via a linear combination of these two objectives:

$$\mathcal{L}_{CL-KE} = \mu_1 \mathcal{L}_{CL} + \mu_2 \mathcal{L}_{KE} \quad (5)$$

where $\mu_1 > 0$ and $\mu_2 > 0$ are hyper-parameters controlling the relative strengths of the two objective functions.

4.2. Knowledge element-guided attention

Next, we observe that the attention mechanism within the vision transformer (ViT) attends to both attacked patches and unaffected patches. This is undesirable because attention paid to attacked patches renders the output embeddings more dependent on the attack signal, and thus more vulnerable. Thus, it is imperative for ViT to allocate reduced attention to attacked patches relative to unaffected patches. Our intuition is that the model should pay more attention to image regions that align well with KEs than patches with low alignment. Thus, we leverage our patch-KE similarity scores to modulate ViT’s attention by enforcing a constraint between ViT’s attention and the patch-KE similarity scores. Given ViT’s query, key, and value denoted as Q, K, V

respectively, the attention weight is computed as $\alpha = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})$, where d_k is the dimensionality of the key vectors. Now, the penalized attention weight can be computed based on the maximum and minimum similarity computed in Eq. (2), Eq. (3) (α_i^c)_{max} = $\alpha_i^c \cdot \omega_i^c$. (α_i^c)_{min} = $\alpha_i^c \cdot \hat{\omega}_i^c$ Since the similarity scores between a targeted visual region and KE are less compared to unaffected patch and KE, ViT pays less attention to attacked patches. The resulting objective function which penalizes attention values which deviate from the patch-KE similarity scores is:

$$\mathcal{L}_{Attention} = -\frac{1}{2N} \left(\sum_{i=1}^N \sum_{c=1}^C (\alpha_i^c) \log(\sigma(\alpha_i^c)_{max}) + \sum_{i=1}^N \sum_{c=1}^C (1 - \alpha_i^c) \log(1 - \sigma(\alpha_i^c)_{min}) \right) \quad (6)$$

The training objective is then:

$$\mathcal{L}_{CL-Attention} = \mu_1 \mathcal{L}_{CL} + \mu_2 \mathcal{L}_{Attention} \quad (7)$$

4.3. Knowledge element weighted contrastive loss

Note that during the fine-tuning process of Eq. (5) and Eq. (7), the contrastive learning objective Eq. (1), seeks to align representations from each modality which has the effect of pulling attacked images and captions closer in

the embedding space. Therefore, we introduce a dynamic weighting function which weights each sample in the contrastive objective function. Our intuition is that attacked samples will have lower similarity scores between image patches and KEs, since the attack does not explicitly target the KEs. Thus, we penalize the contrastive objective for each sample with the average similarity score, so that the contrastive objective is downweighted for attacked samples compared to benign samples. We compute the maximum similarity scores per sample across categories following Eq. (2), where $\lambda_i = \max_{c \in C} \omega_i^c$, $i \in \mathcal{N}$, $\mu_1, \mu_2 = 1$:

$$\mathcal{L}_{CL_i} = \underbrace{\frac{\exp(\frac{\sigma(I_i^e, T_i^e)}{\tau})}{\sum_{k=1}^{\mathcal{N}} \exp(\frac{\sigma(I_i^e, T_k^e)}{\tau})}}_{\text{contrasting } i^{\text{th}} \text{ image with texts}} + \underbrace{\sum_{k=1}^{\mathcal{N}} \log \frac{\exp(\frac{\sigma(I_k^e, T_k^e)}{\tau})}{\exp(\frac{\sigma(I_i^e, T_k^e)}{\tau})}}_{\text{contrasting texts with } i^{\text{th}} \text{ image}} \quad (8)$$

$$\mathcal{L}_{WeightedCL} = -\frac{1}{2\mathcal{N}} \sum_{i=1}^{2\mathcal{N}} \lambda_i \mathcal{L}_{CL_i} \quad (9)$$

Our final objective is likewise given by linear combination:

$$\mathcal{L}_{WeightedCL-Attention} = \mu_1 \mathcal{L}_{WeightedCL} + \mu_2 \mathcal{L}_{Attention} \quad (10)$$

4.4. Knowledge element (KE) generation

Our approach requires external knowledge about each image in addition to a paired caption. For example, a caption of dog image might be "A dog is running in the park". In this case, suitable knowledge elements might be *paws, sharp nails, furry animal, trees*. We follow in context learning approach by prompting a large language model (*Vicuna* [10]) for generating KEs for each image. Note that the KEs are generated purely from the caption or object label and thus are only potentially relevant to the image. Our approach accounts for this by generating 25 KEs per caption/category. Then, we take the top 5 KEs per caption based on the similarity scores between image and generated KEs. For COCO [28], we prompt *Vicuna* with *What are useful visual features for distinguishing a category name in a photo?*. Since COCO has 80 categories we choose this prompt following [32]. For Flickr30k [58], we design prompts that generate KEs for each caption, since we do not have any predefined object classes. Additional details are included in our supplementary.

5. Experiments

5.1. Experimental Setup

Models and datasets. We follow [7]’s setting by attacking CLIP-like models [36]. We adopt ViT-B/16 as image encoder, pretrained on ImageNet-21k [42] and fine-tuned on ImageNet-1k. As a text encoder, we adopt a BERT-style [11] encoder following [36]. We cap the max sequence

length of text to 100. We use AdamW with weight decay using a cosine scheduler from 10^{-4} with decay rate 0.2. We train for 30 epochs with a batch size of 128 on the COCO [28] and Flickr30k [58] datasets. While COCO has 80 defined object categories, Flickr30k has no label information. Additional details are included in supplementary.

Backdoor settings. We tested out defense against three recent backdoor attacks. To do so, we couple backdoored samples with a caption mentioning the target class. Adversaries only require a very small amount of poisoned samples for poisoning contrastive models (*e.g.*, CLIP) [7]. Following this, we inject a very small amount of poisoned samples (0.01% of the train dataset for both COCO and Flickr30k).

Poisoning settings. We performed two types of poisoning attacks following [56]. For single target label attack, the poisoning goal is *dog2boat* for both Flickr30k and COCO. We evaluate them on test samples that are unseen in the training process. For example, we take a clean image of *dog* and associate it with a proxy caption of *boat*. The poisoning rate for this attack is 0.065% for Flickr30k and 0.24% for COCO. For the multi-target label attack, we take two classes. The poisoning goals are *dog2boat* and *train2zebra* for COCO. For Flickr30k, the poisoning goals are *dog2boat* and *bird2sofa*. The poisoning rate for COCO and Flickr30k are 0.52% and 0.34% respectively.

5.2. Experimental Results

Backdoor Attack. In Tab. 1, we compared ablations of our method (CL+ KE, CL + Attention) with other baselines *e.g.* Cleanlip [3], Anti-Backdoor Learning (ABL) [27]. Finally, our model *Semantic Shield* (Weighted CL + Attention), outperforms all baselines with significant margins. Note that, at test time, we used 100 backdoor images (*patch, BPP, Wanet*) for the text retrieval task. At test time, our model retrieves no caption associated with poisoned categories for any backdoored image on Flickr30k.

Poisoning Attack. Similarly, to the above, at test time, we use 100 poisoned images for both single and multi-target settings for both datasets. Our model outperforms all existing work significantly with large margins, particularly on the multi-target label setting. We observe that the unweighted version of our approach slightly outperforms *Semantic Shield* for *dog2boat* at Hit@1, but *Semantic Shield* significantly outperforms for Hit@5 and Hit@10, suggesting significantly reduced poisoning overall.

Utility evaluation. We evaluate model utility for image-caption retrieval. Tab. 4 shows the performance (Recall@10) of the poisoned model on each attack type as well as the clean model on the test data. We observe that the utility of the poisoned model is at the same level or slightly less than the clean model *e.g.* BPP in COCO dataset. This implies that despite being trained on poisoned data, models maintain their performance. We show the model utility

Dataset	Models	Backdoor Patch			Backdoor BPP			Backdoor Wanet		
		Hit@1 ↓	Hit@5 ↓	Hit@10 ↓	Hit@1 ↓	Hit@5 ↓	Hit@10 ↓	Hit@1 ↓	Hit@5 ↓	Hit@10 ↓
COCO	CL (No Defense)	90.66	94.60	95.43	100.0	100.0	100.0	100.0	100.0	100.0
	CL+ ABL [27]	6.23	8.12	12.21	15.35	16.68	16.21	100.0	100.0	100.0
	CL+ CleanClip [3]	5.35	12.68	17.89	36.12	50.09	55.19	8.23	16.32	23.73
	CL + KE	9.0	15.31	21.90	25.39	47.98	50.12	12.21	56.79	88.38
	CL + Attention	4.20	5.12	6.01	0.0	5.26	36.21	0.0	2.10	7.20
	Weighted CL + Attention	0.9	1.22	1.57	0.0	0.0	0.0	0.0	0.0	0.0
Flickr30k	CL (No Defense)	91.97	97.63	98.21	100.0	100.0	100.0	100.0	100.0	100.0
	CL+ ABL [27]	4.67	2.21	4.06	10.34	17.98	21.13	98.21	99.23	100.0
	CL+ CleanClip [3]	2.20	3.32	5.05	12.43	24.32	31.25	13.29	23.13	29.21
	CL + KE	16.10	33.15	41.09	13.14	36.54	56.27	23.36	41.21	47.43
	CL + Attention	1.20	3.12	3.01	0.0	7.24	23.17	0.0	12.01	14.07
	Weighted CL + Attention	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 1. Backdoor attack and defense performance with baselines. The first row of the table shows an undefended model while other rows are baselines or variants of our method. CL+ KE, CL+ Attention are our baselines. The best results are shown in bold.

Dataset	Models	Single Target Label			Multiple Target Label					
		dog2boat			dog2boat			train2zebra		
		Hit@1 ↓	Hit@5 ↓	Hit@10 ↓	Hit@1 ↓	Hit@5 ↓	Hit@10 ↓	Hit@1 ↓	Hit@5 ↓	Hit@10 ↓
COCO	CL (No Defense)	18.0	57.20	82.0	77.12	99.23	99.56	55.32	95.76	97.98
	CL+ CleanClip [3]	3.39	3.95	5.65	57.69	63.0	89.17	69.49	71.75	89.17
	CL + KE	4.56	5.32	5.95	54.45	64.21	85.52	65.12	70.92	86.12
	CL + Attention	0.56	3.38	4.51	0.63	65.60	69.42	2.25	6.77	12.99
	Weighted CL + Attention	0.04	1.12	2.54	2.23	5.21	6.45	0.0	0.0	0.0
Flickr30k		dog2boat			dog2boat			bird2sofa		
		Hit@1 ↓	Hit@5 ↓	Hit@10 ↓	Hit@1 ↓	Hit@5 ↓	Hit@10 ↓	Hit@1 ↓	Hit@5 ↓	Hit@10 ↓
	CL (No Defense)	29.0	57.20	82.23	28.12	82.39	93.76	55.32	90.62	100.0
	CL+ CleanClip [3]	8.27	31.51	36.61	21.69	61.27	88.75	22.42	64.11	89.51
	CL + KE	7.34	28.09	32.21	21.12	45.32	47.67	12.77	42.34	54.21
	CL + Attention	4.56	21.81	34.11	1.63	16.70	29.21	3.25	18.43	32.22
Weighted CL + Attention	0.32	1.21	2.54	1.78	4.56	5.67	0.0	0.0	0.0	

Table 2. Poisoning attack and defense performance with baselines. First row of the table shows how good the attack, and other rows are baselines along with our proposed models. CL + KE, CL + Attention are our baselines. The best results are highlighted.

after being defended with Semantic Shield and its variants (CL + KE, CL + Attention, weighted CL + Attention) in Tab. 3. We largely observe a similar utility compared to the models from Tab. 4. On the Flickr30k dataset, single target or multiple target attack scenario, for TR task, the utility is significantly less than the clean model (Tab. 4, Tab. 3).

5.3. Ablations

Poisoning rate. We compare the performance of poisoning attacks at different poisoning rates on three backdoor attacks. We conduct these attacks against the victim model with four different poisoning rates (0.001 to 0.01%) on the COCO dataset (Fig. 3). We observe that attack performance significantly improves with increased poisoning rate, even though the rate is quite low, which demonstrates the vulnerability of contrastively trained VL models to attacks.

Fine-tuning epoch. In Fig. 4 we use the max poisoning rate (0.01%) from Fig. 3 to illustrate Semantic Shield’s per-

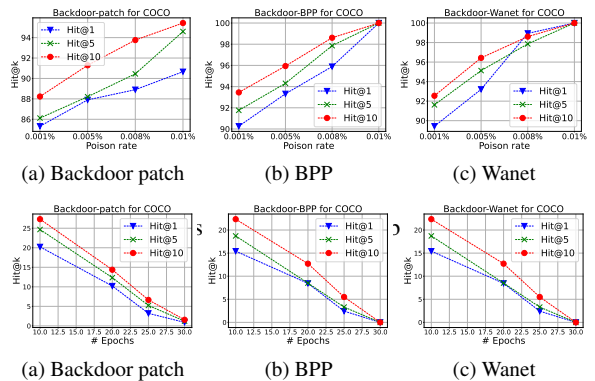


Figure 4. Hit@k vs training epoch for Semantic Shield.

formance at different epochs on the same backdoored samples. We notice that Hit@k gradually reduces for all three attacks, demonstrating the increasing effectiveness of Semantic Shield’s defense with increased training.

Dataset	Task	Models	Backdoor Patch	BPP	Wanet	Single Target Label	Multiple Target Label
COCO	IR	CL	74.99	73.94	74.54	74.68	74.72
		CL + KE	74.15	70.7	74.0	74.24	73.28
		CL + Attention	74.38	73.13	74.43	75.70	75.13
		Weighted CL + Attention	74.22	74.56	74.23	73.46	73.51
COCO	TR	CL	81.58	77.44	78.74	80.16	81.12
		CL + KE	78.40	75.54	77.86	79.08	81.20
		CL + Attention	79.20	77.36	78.04	80.05	81.06
		Weighted CL + Attention	79.46	77.78	78.45	79.67	80.0
Flickr30k	IR	CL	59.13	59.86	61.08	60.92	57.41
		CL + KE	60.34	61.85	61.13	58.12	58.18
		CL + Attention	61.32	55.96	59.14	58.97	58.16
		Weighted CL + Attention	61.07	56.32	60.16	59.76	58.78
Flickr30k	TR	CL	68.07	68.79	69.86	71.06	68.14
		CL + KE	69.67	70.65	69.62	66.98	62.20
		CL + Attention	70.0	64.46	68.0	68.13	62.97
		Weighted CL + Attention	70.23	65.66	68.87	68.45	62.12

Table 3. Model utility of defended models (Recall@10). The model utilities are comparable to the performance in Tab. 4

Dataset	Task	Clean	BackPat	BPP	Wanet	SingTL	MultTL
COCO	IR	75.13	74.99	73.94	74.54	74.68	74.72
	TR	80.62	81.58	77.44	78.74	80.16	81.12
Flickr30k	IR	59.68	59.13	59.86	61.08	60.92	57.41
	TR	68.37	68.07	68.79	69.86	71.06	68.14

Table 4. Model utility between *clean* model and other backdoored/poisoned models (CL) (Recall@10). Similar to Tab. 3.

6. Qualitative analysis

In Fig. 5, we present the contrast between a model defended by Semantic Shield and an undefended model’s attention map. Fig. 5b shows that poisoned model pays attention to the patch (bottom right corner). In contrast, the defended model Fig. 5c does not pay any attention to the patch. Next, in Fig. 5d and Fig. 5g two imperceptible noises are injected *e.g.* BPP, Wanet. We wanted to see what happens if we inject the noise randomly throughout the entire images. Poisoned models in Fig. 5e and Fig. 5h show spurious visual signals all over the image. However, our proposed models filters out the noisy signals and defends against poisoning.

7. Conclusion

In this paper, we introduced Semantic Shield, an approach for defending against attacks on contrastively trained VL models. Our approach works by leveraging external knowledge to guide the model’s attention to non-attacked visual regions and samples. We evaluated Semantic Shield against recent backdooring and poisoning attacks and defenses on two benchmarks. Our experiments show that Semantic Shield substantially outperforms existing defenses across all settings. In future work, we will explore a tighter integration of the LLM using prompting by dynamically producing KEs online based on the de-

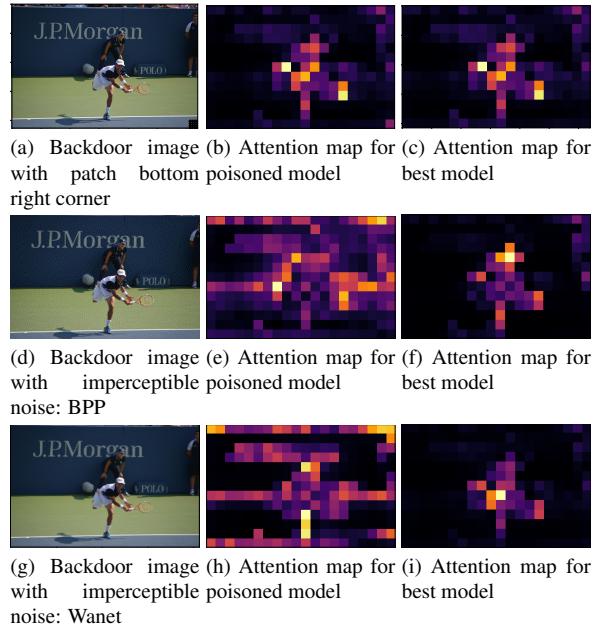


Figure 5. Attention map comparison between our model (weighted CL + attention) and backdoored models for three backdoor attacks.

fended model’s current state. In addition, we will explore how multimodal large language models could be used to extract more relevant KEs. While Semantic Shield is successful at defending against attacks on natural images for which there is a meaningful visual-KE alignment, it may be less successful for images such as charts or more abstract text for which clear KEs cannot be extracted. Moreover, it does not preclude the possibility of attacks against the language model via the caption. Future work should explore how the LLM can be jointly defended.

References

- [1] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. VATT: transformers for multimodal self-supervised learning from raw video, audio and text. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 24206–24221, 2021. [1](#)
- [2] Hritik Bansal, Nishad Singhi, Yu Yang, Fan Yin, Aditya Grover, and Kai-Wei Chang. Cleanclip: Mitigating data poisoning attacks in multimodal contrastive learning. In *ICLR 2023 Workshop on Trustworthy and Reliable Large-Scale Machine Learning Models*, 2023. [3](#)
- [3] Hritik Bansal, Nishad Singhi, Yu Yang, Fan Yin, Aditya Grover, and Kai-Wei Chang. Cleanclip: Mitigating data poisoning attacks in multimodal contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 112–123, 2023. [1](#), [3](#), [6](#), [7](#)
- [4] Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, pages 1467–1474, 2012. [2](#)
- [5] Battista Biggio, Ignazio Pillai, Samuel Rota Bulò, Davide Ariu, Marcello Pelillo, and Fabio Roli. Is data clustering in adversarial settings secure? In *Proceedings of the 2013 ACM Workshop on Artificial Intelligence and Security*, page 87–98, New York, NY, USA, 2013. Association for Computing Machinery. [3](#)
- [6] Battista Biggio, Ignazio Pillai, Samuel Rota Bulò, Davide Ariu, Marcello Pelillo, and Fabio Roli. Is data clustering in adversarial settings secure? In *Proceedings of the 2013 ACM workshop on Artificial intelligence and security*, pages 87–98, 2013. [2](#)
- [7] Nicholas Carlini and Andreas Terzis. Poisoning and backdooring contrastive learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [3](#), [4](#), [6](#)
- [8] 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. In *Workshop on Artificial Intelligence Safety*. CEUR-WS, 2019. [3](#)
- [9] Jian Chen, Xuxin Zhang, Rui Zhang, Chen Wang, and Ling Liu. De-pois: An attack-agnostic defense against data poisoning attacks. *IEEE Transactions on Information Forensics and Security*, 16:3412–3425, 2021. [3](#)
- [10] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna>, 1(2):3. [2](#), [6](#)
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. [6](#)
- [12] Khoa Doan, Yingjie Lao, and Ping Li. Backdoor attack with imperceptible input and latent modification. *Advances in Neural Information Processing Systems*, 34:18944–18957, 2021. [3](#)
- [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. [3](#)
- [14] Vishnu Sashank Dorbala, Gunnar A Sigurdsson, Jesse Thomason, Robinson Piramuthu, and Gaurav S Sukhatme. Clip-nav: Using clip for zero-shot vision-and-language navigation. In *Workshop on Language and Robotics at CoRL 2022*, 2022. [2](#)
- [15] Fangxiang Feng, Xiaojie Wang, and Ruifan Li. Cross-modal retrieval with correspondence autoencoder. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 7–16, 2014. [2](#)
- [16] Felipe González-Pizarro and Savvas Zannettou. Understanding and detecting hateful content using contrastive learning. In *Proceedings of the International AAAI Conference on Web and Social Media*, pages 257–268, 2023. [2](#)
- [17] Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access*, 7:47230–47244, 2019. [3](#)
- [18] Jonathan Hayase, Weihao Kong, Raghav Somani, and Sewoong Oh. Spectre: Defending against backdoor attacks using robust statistics. In *International Conference on Machine Learning*, pages 4129–4139. PMLR, 2021. [3](#)
- [19] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10608–10615. IEEE, 2023. [2](#)
- [20] Kunzhe Huang, Yiming Li, Baoyuan Wu, Zhan Qin, and Kui Ren. Backdoor defense via decoupling the training process. In *International Conference on Learning Representations*, 2021. [3](#)
- [21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. [2](#)
- [22] Marius Kloft and Pavel Laskov. Online anomaly detection under adversarial impact. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 405–412. JMLR Workshop and Conference Proceedings, 2010. [2](#)
- [23] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR, 2017. [2](#)
- [24] Changjiang Li, Ren Pang, Zhaohan Xi, Tianyu Du, Shouling Ji, Yuan Yao, and Ting Wang. An embarrassingly simple

- backdoor attack on self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4367–4378, 2023. 1
- [25] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 1, 2
- [26] Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, pages 12888–12900. PMLR, 2022. 1
- [27] Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. Anti-backdoor learning: Training clean models on poisoned data. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 14900–14912, 2021. 1, 3, 6, 7
- [28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 6
- [29] Min Liu, Alberto Sangiovanni-Vincentelli, and Xiangyu Yue. Beating backdoor attack at its own game. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4620–4629, 2023. 3
- [30] Yang Liu, Mingyuan Fan, Cen Chen, Ximeng Liu, Zhuo Ma, Li Wang, and Jianfeng Ma. Backdoor defense with machine unlearning. In *IEEE INFOCOM 2022—IEEE Conference on Computer Communications*, pages 280–289. IEEE, 2022. 3
- [31] Arjun Majumdar, Gunjan Aggarwal, Bhavika Devnani, Judy Hoffman, and Dhruv Batra. Zson: Zero-shot object-goal navigation using multimodal goal embeddings. *Advances in Neural Information Processing Systems*, 35:32340–32352, 2022. 2
- [32] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 6
- [33] Tuan Anh Nguyen and Anh Tuan Tran. Wanet - imperceptible warping-based backdoor attack. In *International Conference on Learning Representations*, 2021. 3, 4
- [34] Huy Phan, Cong Shi, Yi Xie, Tianfang Zhang, Zhuohang Li, Tianming Zhao, Jian Liu, Yan Wang, Yingying Chen, and Bo Yuan. Ribac: Towards robust and imperceptible backdoor attack against compact dnn. In *European Conference on Computer Vision*, pages 708–724. Springer, 2022. 3
- [35] 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. 3
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. 1, 2, 6
- [37] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *CoRR*, abs/2204.06125, 2022. 1
- [38] Aniruddha Saha, Akshayvarun Subramanya, and Hamed Pirsiavash. Hidden trigger backdoor attacks. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11957–11965, 2020. 3
- [39] Aniruddha Saha, Ajinkya Tejankar, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Backdoor attacks on self-supervised learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 13327–13336. IEEE, 2022. 3
- [40] Ali Shafahi, W. Ronny Huang, Mahyar Najibi, Octavian Suci, Christoph Studer, Tudor Dumitras, and Tom Goldstein. Poison frogs! targeted clean-label poisoning attacks on neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, page 6106–6116, Red Hook, NY, USA, 2018. Curran Associates Inc. 3
- [41] Wonyoung Shin, Jonghun Park, Taekang Woo, Yongwoo Cho, Kwangjin Oh, and Hwanjun Song. e-clip: Large-scale vision-language representation learning in e-commerce. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 3484–3494, 2022. 2
- [42] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *Trans. Mach. Learn. Res.*, 2022, 2022. 6
- [43] Di Tang, XiaoFeng Wang, Haixu Tang, and Kehuan Zhang. Demon in the variant: Statistical analysis of {DNNs} for robust backdoor contamination detection. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1541–1558, 2021. 3
- [44] Christopher Thomas and Adriana Kovashka. Preserving semantic neighborhoods for robust cross-modal retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 317–335. Springer, 2020. 2
- [45] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *Computer Security—ESORICS 2020: 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14–18, 2020, Proceedings, Part I 25*, pages 480–501. Springer, 2020. 2
- [46] Brandon Tran, Jerry Li, and Aleksander Madry. Spectral signatures in backdoor attacks. *Advances in neural information processing systems*, 31, 2018. 3

- [47] Wei Lun Tsai, Jacob J Lin, and Shang-Hsien Hsieh. Generating construction safety observations via clip-based image-language embedding. In *European Conference on Computer Vision*, pages 366–381. Springer, 2022. 2
- [48] Haotao Wang, Junyuan Hong, Aston Zhang, Jiayu Zhou, and Zhangyang Wang. Trap and replace: Defending backdoor attacks by trapping them into an easy-to-replace subnetwork. *Advances in neural information processing systems*, 35:36026–36039, 2022. 3
- [49] Lin Wang and Jie Chen. Improving radiology report generation with adaptive attention. In *Multimodal AI in health-care: A paradigm shift in health intelligence*, pages 293–305. Springer, 2022. 2
- [50] Longzheng Wang, Chuang Zhang, Hongbo Xu, Yongxiu Xu, Xiaohan Xu, and Siqi Wang. Cross-modal contrastive learning for multimodal fake news detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5696–5704, 2023. 2
- [51] Zhenting Wang, Juan Zhai, and Shiqing Ma. Bppattack: Stealthy and efficient trojan attacks against deep neural networks via image quantization and contrastive adversarial learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 15054–15063. IEEE, 2022. 4
- [52] Sandamal Weerasinghe, Tansu Alpcan, Sarah M Erfani, and Christopher Leckie. Defending support vector machines against data poisoning attacks. *IEEE Transactions on Information Forensics and Security*, 16:2566–2578, 2021. 3
- [53] Dongxian Wu and Yisen Wang. Adversarial neuron pruning purifies backdoored deep models. *Advances in Neural Information Processing Systems*, 34:16913–16925, 2021. 3
- [54] Huang Xiao, Battista Biggio, Gavin Brown, Giorgio Fumera, Claudia Eckert, and Fabio Roli. Is feature selection secure against training data poisoning? In *International conference on machine learning*, pages 1689–1698. PMLR, 2015. 2
- [55] Wenhan Yang, Jingdong Gao, and Baharan Mirzasoleiman. Robust contrastive language-image pretraining against data poisoning and backdoor attacks. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 1
- [56] Ziqing Yang, Xinlei He, Zheng Li, Michael Backes, Mathias Humbert, Pascal Berrang, and Yang Zhang. Data poisoning attacks against multimodal encoders. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, pages 39299–39313. PMLR, 2023. 6
- [57] Ziqing Yang, Xinlei He, Zheng Li, Michael Backes, Mathias Humbert, Pascal Berrang, and Yang Zhang. Data poisoning attacks against multimodal encoders. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, pages 39299–39313. PMLR, 2023. 1
- [58] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguistics*, 2:67–78, 2014. 6
- [59] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *Transactions on Machine Learning Research*, 2022. 1, 2
- [60] Yi Zeng, Si Chen, Won Park, Zhuoqing Mao, Ming Jin, and Ruoxi Jia. Adversarial unlearning of backdoors via implicit hypergradient. In *International Conference on Learning Representations*, 2021. 3
- [61] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18123–18133, 2022. 2
- [62] Hanwang Zhang, Yang Yang, Huanbo Luan, Shuicheng Yang, and Tat-Seng Chua. Start from scratch: Towards automatically identifying, modeling, and naming visual attributes. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 187–196, 2014. 2
- [63] Jinghuai Zhang, Hongbin Liu, Jinyuan Jia, and Neil Zhenqiang Gong. Corruptencoder: Data poisoning based backdoor attacks to contrastive learning. *arXiv preprint arXiv:2211.08229*, 2022. 3
- [64] Ying Zhang and Huchuan Lu. Deep cross-modal projection learning for image-text matching. In *Proceedings of the European conference on computer vision (ECCV)*, pages 686–701, 2018. 2
- [65] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*, pages 2–25. PMLR, 2022. 2
- [66] Bingyin Zhao and Yingjie Lao. Towards class-oriented poisoning attacks against neural networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3741–3750, 2022. 2
- [67] Yangming Zhou, Yuzhou Yang, Qichao Ying, Zhenxing Qian, and Xinpeng Zhang. Multimodal fake news detection via clip-guided learning. In *2023 IEEE International Conference on Multimedia and Expo (ICME)*, pages 2825–2830. IEEE, 2023. 2