Multimodal Unlearnable Examples: Protecting Data against Multimodal Contrastive Learning

Xinwei Liu

Institute of Information Engineering, Chinese Academy of Sciences & School of Cyberspace Security, University of Chinese Academy of Sciences Beijing, China liuxinwei@iie.ac.cn Xiaojun Jia* Cyber Security Research Centre @ NTU, Nanyang Technological University Singapore jiaxiaojunqaq@gmail.com

Siyuan Liang School of Computing, National University of Singapore Singapore pandaliang521@gmail.com

ABSTRACT

Multimodal contrastive learning (MCL) has shown remarkable advances in zero-shot classification by learning from millions of imagecaption pairs crawled from the Internet. However, this reliance poses privacy risks, as hackers may unauthorizedly exploit imagetext data for model training, potentially including personal and privacy-sensitive information. Recent works propose generating unlearnable examples by adding imperceptible perturbations to training images to build shortcuts for protection. However, they are designed for unimodal classification, which remains largely unexplored in MCL. We first explore this context by evaluating the performance of existing methods on image-caption pairs, and they do not generalize effectively to multimodal data and exhibit limited impact to build shortcuts due to the lack of labels and the dispersion of pairs in MCL. In this paper, we propose Multi-step Error Minimization (MEM), a novel optimization process for generating multimodal unlearnable examples. It extends the Error-Minimization (EM) framework to optimize both image noise and an additional text trigger, thereby enlarging the optimized space and effectively misleading the model to learn the shortcut between the noise features and the text trigger. Specifically, we adopt projected gradient descent to solve the noise minimization problem and use HotFlip to approximate the gradient and replace words to find the optimal text trigger. Extensive experiments demonstrate the effectiveness of MEM, with post-protection retrieval results nearly half of random

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0686-8/24/10 https://doi.org/10.1145/3664647.3680708 ail.com Sciences Beijing, China xunyuan@iie.ac.cn Xiaochun Cao* School of Cyber Science and

Technology, Shenzhen Campus, Sun Yat-sen University Shenzhen, China caoxiaochun@mail.sysu.edu.cn

guessing, and its high transferability across different models. Our code is available on the https://github.com/thinwayliu/Multimodal-Unlearnable-Examples

Yuan Xun

Institute of Information Engineering,

Chinese Academy of Sciences &

School of Cyberspace Security,

University of Chinese Academy of

CCS CONCEPTS

• Security and privacy → Privacy protections; • Computing methodologies → Computer vision; Machine learning.

KEYWORDS

Data Protection; Privacy Protection; Unlearnbale Examples; Poisoning Attack; Multimodal Contrastive Learning

ACM Reference Format:

Xinwei Liu, Xiaojun Jia, Yuan Xun, Siyuan Liang, and Xiaochun Cao. 2024. Multimodal Unlearnable Examples: Protecting Data against Multimodal Contrastive Learning. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24), October 28–November 1, 2024, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/ 3664647.3680708

1 INTRODUCTION

In recent years, there has been a growing interest in multimodal models among researchers in the community [3]. Traditional methods [44, 51, 61] have primarily focused on analyzing a single modal of data. However, with the rise of multimodal learning, different types of data, such as text, images, and audio, are being combined into a unified framework. One of the most popular approaches for multimodal learning is Multimodal Contrastive Learning (MCL), as demonstrated by models such as CLIP [66] and ALIGN [35]. These models are trained with a contrastive loss, which encourages the correlation between pairs of images and captions while also keeping them distinct from unrelated pairs. This approach reduces the need for extensive manual annotation of training data and allows the use of larger datasets that contain millions of examples. MCL has shown promise in various applications, including image classification [66], image captioning [41, 63], image generation [43, 65].

^{*}Correspondence to: Xiaojun Jia and Xiaochun Cao.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia.



Figure 1: Posts on Facebook inadvertently leak personal information. Utilizing MEM-3 to protect data can prevent unauthorized models from accessing private features.

Training of high-performance multimodal models is highly dependent on large amounts of multimodal data, often sourced from publicly available datasets such as CC12M [5], YFCC100M [74], and LAION5B [69]. However, as the demand for larger datasets continues to surge in the future, these datasets may be still insufficient. Consequently, malicious actors may resort to unauthorized data acquisition from the web or engage in the crawling of user posts on social networks for commercial training purposes. However, these datasets often contain significant amounts of sensitive personal information, raising concerns among people about the potential unauthorized use of personal data and the leakage of user privacy.

A series of recent works make efforts to prevent unauthorized usage in image classification by making the image unexploitable. Specifically, they poison the data with some imperceptible perturbations [48, 52-54, 77], creating 'shortcuts' [82] in the training process that hinder the models from learning the features of the images [83]. This kind of attack is called availability attacks or indiscriminate poisoning attacks, and these poisoned training data are called unlearnable examples. These unlearning methods [50] can be broadly classified into two categories: model-free and model-based attacks. Model-free attacks usually generate unlearnable noise at the pixel level without any knowledge of clean data and directly create shortcuts between image noise and labels, such as LSP [82] and CUDA [67]. Due to their direct association with labels, these patterns often exhibit high efficiency. Model-based attacks typically generate noise through surrogate models. The surrogate model learns the features through the training phase and generates the feature-level noise, such as Error-Minimizing [31] and Adversarial Poisoning [17]. However, there has yet to be research to consider protecting multimodal data in the context of MCL.

We are the first to consider a scenario focused on generating multi-modal unlearnable examples against privacy risks [6, 11, 13, 26, 49] associated with MCL. In this context, we concentrate on image-text pairs as a representative multimodal dataset. Users are assumed to frequently share personal photos with text on social media platforms like Facebook, including some private identity information such as faces, names, phone numbers, and addresses. Currently, hackers attempt to collect large amounts of such imagetext pairs from the Internet and utilize MCL techniques to train modern foundational models, as illustrated in the left segment of Fig. 1. These models inadvertently capture user's private information and facial characteristics, leading to potential privacy leakage. Protectors aim to prevent the unauthorized exploitation of these sensitive data by performing unlearning methods on multimodal data. These methods aim to render models trained on such multimodal unlearnable examples incapable of accessing users' privacy features, while not impeding users' social interactions after posting images and text, as depicted in the right segment of Fig. 1.

An intuitive idea involves extending the unlearning methods for image classification to MCL. However, we explore the performance of these methods on multimodal data, and all of them fail to present effective protection due to increased data modalities or dispersion of data pairs. For model-free attacks, they fail to generate specific noise patterns that strongly correlate with the category for a shortcut due to the lack of certain labels in the image-text pair. For model-based attacks, while it may be feasible to optimize noise to build shortcuts with clean captions, their efficacy is significantly diminished. Our analysis was primarily attributed to the dispersion of the data pairs, which presents challenges in learning the noise pair and captions compared to the images and labels in classification, as depicted in Fig. 2. Therefore, establishing more efficient shortcuts is key to generating effective multi-modal unlearnable examples.

In this paper, we propose a novel optimization framework that efficiently generates unlearnable multimodal examples for imagecaption pairs. We first adopt the Error-minimizing (EM) [31] framework as a basis for our attack, and we consider optimizing the noise along with an additional text trigger. Specifically, following adversarial triggers in NLP [76], we propose to add short text sequences as triggers in the front of the clean caption, which does not affect the understanding of the text by the user. Therefore, it can be formulated a multi-step minimization problem to optimize the noise-text trigger pair and build the shortcut between them, which we dubbed Multi-step Minimization (MEM), to prevent the unauthorized model from learning the features of images and captions. During optimization, we consider adopting projected gradient descent (PGD) [62] to solve the noise minimization problem and use the HotFlip [12] method to approximate the gradient and replacement strategy in [76] to select the optimal trigger in the text minimization problem. Through extensive experiments, we verify that unlearnable examples generated by MEM provide better protection and exhibit transferability across different models.

In summary, our main contributions are:

• We are the first to consider a new scenario called multimodal data protection, which aims to prevent multimodal personal data on social media from unauthorized MCL, and we take the image-text pair as examples.

- We analyze the limitations of previous methods extended to multimodal contrastive learning, attributed to the increase in modality and the dispersion of the caption features.
- We propose a Multi-step Error Minimization (MEM) to generate effective multimodal unlearnable examples, which leverages an additional optimized text trigger for better convergence at the basis of Error-minimizing (EM).
- Extensive experiments are conducted to verify the effectiveness of our method with different datasets. In addition, we present a practical case study on face privacy protection within a fine-tuning scenario.

2 RELATED WORK

2.1 Multimodal Contrastive Learning

Initially designed for self-supervised representation learning in unimodal contexts, contrastive learning methods aim to improve the agreement between views augmented differently in the same instance while reducing the agreement between views of distinct instances [7, 10, 27, 29, 64]. Recently, these techniques have been extended to multimodal domains, notably in the context of paired image-text datasets. Multimodal contrastive models like CLIP [66] and ALIGN [35] have undergone extensive pretraining on vast datasets including hundreds of millions to billions of image-text pairs. Their objective is to maximize the agreement between representations of matched image-caption pairs while minimizing agreement for non-matched pairs. As a result, these models have exhibited exceptional performance in zero-shot classification tasks and have shown robustness to distributional shifts.

2.2 Poisoning Attacks

Data poisoning attacks aim to disrupt the model training process by processing the training dataset, resulting in a significant increase in test errors for some specific samples during the testing phase. [9, 58, 71]. A common type of data poisoning attack is the backdoor attack [20, 25, 34, 46, 47, 55-57, 59, 60, 78, 85, 86], which typically involves injecting triggers into training samples, leading to misclassification of images containing these trigger patterns during testing [19, 42]. However, it typically only affects samples with trigger, while clean samples remain unaffected and can be correctly classified. Recent work also explores poisoning attacks on multimodal contrastive learning (MCL). Yang et al. [80] study the poisoning attacks against multimodal models in both visual and linguistic modalities. Carlini and Terzis [4] introduced a framework that effectively poisoned CLIP models with backdoor attacks. In addition, some works aim to train a robust CLIP model against data poisoning and backdoor attacks [2, 79]. Due to the time-consuming of training a high-performance model with a large-scale dataset, we will follow the experimental setting of these works to train a rather simpler model with a small dataset in this paper. When the goal is generalized to the entire test set, the poisoning attack is called an indiscriminate/availability poisoning attack [28, 82]. In the community, this is considered a special case of a poisoning attack, which can also be divided into clean-label and dirty-label scenario.

2.3 Unlearnable Methods

Unlearnable methods aim to protect data from unauthorized training of the classification model by introducing imperceptible noise to the images, in fact, clean-label indiscriminate/availability poisoning attacks. Models trained on such unlearnable examples typically exhibit a notable decrease in accuracy on clean test sets.

There are two categories to generate unlearnable examples in image classification. The first category involves model-based attacks, which require a surrogate model to guide the perturbation generation [16, 72]. The error minimization (EM) [31] minimizes the classification error of images on a surrogate classifier and iteratively updates the surrogate model with these perturbed images. Fowl et al. [17] propose adversarial poisoning, which contains targeted adversarial perturbation (TAP) and untargeted adversarial perturbation (UAP), which uses adversarial examples [21-24, 32, 33] as poisoned data to make the model unlearn the features. TAP presents an effective protection for the image dataset. While model-based methods are potent, they are often computationally demanding. Some findings reveal that these approaches can be easily neutralized by adversarial training (AT) [1, 36-38, 45]. To address this, several works attempt to leverage robust protection against AT [15, 18]. Model-free methods do not require a surrogate model to generate noise for images. Yu et al. [82] empirically investigate various unlearnable methods and show that all of them use these spurious features to create a shortcut in the model. They also propose the Linear-sperate Synthetic Perturbation (LSP) in response to this characteristic and show great effectiveness. Autoregressive poisoning [68] proposes a generic perturbation generation with each class that can be applied to different datasets, which is a series of dataset-independent perturbations. CUDA [67] is generated using controlled class-wise convolutions with filters that are randomly generated via a private key. These previous works focus on image classification using cross-entropy loss, and He et al. [28] explore the unlearnable examples for unsupervised contrastive learning and discover that extended EM and TAP methods can still effectively protect against unsupervised learning. However, when extending them to multimodal contrastive learning, their effectiveness diminishes due to the increased modality.

3 MULTIMODAL UNLEARNABLE EXAMPLES

3.1 Preliminaries

Our scenario involves two main parties: the *protector*, and the *hacker*. We assume that the protector is aware that multimodal privacy data may be unauthorized used by hackers. Consequently, the protector takes measures to render the samples unlearnable by performing operations on the data set. Subsequently, the protector releases these unlearnable multimodal examples to the Internet. Then, when hackers crawl the multimodal data from the web to train multimodal models from scratch with an initial model or fine-tune the pre-trained model, these unlearnable examples will prevent the model from learning the features of private data, and present a poor representation of features across modalities. In this paper, we take the image-text pair as multimodal data as an example, and choose the CLIP [66] as the model used by hackers, which is the most representative multimodal contrastive learning framework.



Figure 2: Comparison of different methods in classification and multimodal contrastive learning (MCL). I_i denotes the image, and T_i is the paired caption. The blue area is the expected decision boundary of the models trained on unlearnable examples.

Here, we formulate our problem: we begin by considering a personal image-caption dataset $\mathcal{D} \subset \mathcal{I} \times \mathcal{T}$ that comprises pairs (I_i, T_i) , where I_i is an image and T_i is the associated caption. The protector aims to generate an unlearnable set \mathcal{D}_u that contains the unlearnable image-text pair (I'_i, T'_i) , and make a CLIP model f^* trained on it generalize poorly on the clean distribution \mathcal{D} :

$$\arg \max \mathbb{E}_{(I,T)\sim\mathcal{D}} \left[\mathcal{L} \left(f^*(I,T) \right) \right],$$

s.t. $f^* \in \arg \min_{f} \sum_{(I',T')\in\mathcal{D}_u} \left[\mathcal{L} \left(f(I',T') \right) \right]$ (1)

3.2 Limitation of Existing Works

Directly solving Eq. 1 is intractable for deep neural networks, and recent works have designed multiple approximate solutions. We can follow these works and extend them to MCL. However, all model-free methods for classification fail to generate image noise here because these methods aim to find a series of specific noise patterns for images related to one certain class, while there is no label in the image-caption pair data. Therefore, only model-based methods can be applied to the MCL, and we extend two typical methods to generate unlearnable multimodal examples. Specifically, they leveraged the addition of a set of perturbations Δ on the images and employ an l_{∞} to bound for each noise δ .

3.2.1 The Error-minimizing Noise. (EM). Huang et al. [31] propose a bi-level objective to generate perturbations δ on the images with a surrogate model, which aims to minimize the loss of the surrogate model on training data. We denote image noise by δ and the surrogate model by f. Therefore, we can apply it to the multimodal unlearnable examples with the CLIP loss \mathcal{L} , and the objective is the following:

$$\arg\min_{\delta\in\Delta}\mathbb{E}_{(I,T)\sim\mathcal{D}}\left[\min_{f}\mathcal{L}\left(f\left(I+\delta,T\right)\right)\right].$$
(2)

3.2.2 Untargeted Adversarial Perturbation. (UAP). Instead of employing bi-level objectives, Fowl et al. [17] demonstrate that the common objectives used for generating adversarial examples are sufficient as unlearnable perturbations. They utilized untargeted adversarial perturbations (UAP) and targeted adversarial perturbations (TAP) as examples of unlearnable perturbations. However,

extending TAP to image-caption data is challenging due to the difficulty in selecting the desired target goal for the caption. Therefore, we can only construct UAP δ using the following objective:

$$\arg\max_{\delta \in \Lambda} \mathbb{E}_{(I,T) \sim \mathcal{D}} \left[\mathcal{L} \left(f^* \left(I + \delta, T \right) \right) \right].$$
(3)

However, from Table 1, we experimentally found that although EM and UAP can be applied to image-caption pairs, they fail to achieve highly effective protection, especially UAP. We explore the reasons for the decline in the effectiveness of these methods from image classification to multimodal contrastive learning. In image classification, EM and UAP optimize the images with the same label to converge in a feature space, resulting in the model easily capturing these additive noises and learning the correlation with labels, as shown in Fig.2(a). However, in multimodal contrastive learning (MCL), to effectively apply the EM and UAP methods, the direction of the optimized image noise must relate to the features of the caption, causing the image features to become either close to or far away from these features. Nevertheless, the caption features of different pairs may be widely dispersed in the image-caption dataset. Consequently, as illustrated in Fig. 2 (b) and (c), it becomes more challenging for the model to capture the correlation between captions and noise generated by EM and UAP compared to those in classification. In Fig. 2 (c), the learning decision space of UAP is much more complex, so its poor protection is to be expected.

3.3 Multi-step Error Minimization (MEM)

In the previous section, we showed that the model-based methods still fail to achieve effective protection due to the dispersion of image-text pairs. An intuitive strategy of enhancement entails optimizing both the image and the captions for a larger optimized space, boosting their convergence across different pairs in the feature space. Consequently, the optimized feature representations of the image and caption set exhibit similar distributions, facilitating the model's learning of their shortcut, as illustrated in Fig. 2 (d).

To this end, we take the EM method as the basic framework and propose adding an additional short text trigger to the caption to minimize contrastive loss, following the setting of an adversarial attack on text tasks [76]. Regarding the length of the trigger, longer triggers are more effective, while shorter triggers are more stealthy.



Figure 3: The framework of MEM. At each step, we concatenate the current trigger to clean captions and attach noise to clean images. We then compute the gradient with current images and tokens. We update the images using Eq. 5 and update the text triggers with Eq. 6. The solid line represents forward propagation, and the dashed line represents backward propagation.

Therefore, our method can be conceptualized as a tri-level iterative optimization problem, resembling a multi-step process of EM. Specifically, we sequentially optimize noise δ and text trigger t to reduce contrastive loss between optimized images $I + \delta$ and optimized text $T \oplus t$, where \oplus denotes the triggers that can be inserted into clean text T at various positions. For simplicity, we choose to add text triggers at the beginning of the text. As a result, our Multi-step Error Minimization (MEM) can be formulated as follows:

ć

. . .

$$\arg\min_{\delta\in\Delta,t\in\mathscr{T}}\mathbb{E}_{(I,T)\sim\mathscr{D}}\left[\min_{f}\mathscr{L}\left(I+\delta,T\oplus t\right)\right].$$
(4)

We can iteratively optimize the above problem sequentially by referring to the method in EM. We use projected gradient descent (PGD) [62] to solve the noise minimization problem in Eq. 4. Notably, to mitigate the noise overfit to the clean captions, we augment them by shuffling the clean caption in a batch and adding the correct matching text triggers on them. Thus, this generated noise can focus more on the text trigger than on part of the captions, when facing the semantic wrong captions. So we can get the optimal δ according to the following iterative formula,

$$\delta^{t+1} = \operatorname{Proj}\left(\delta^{t} - \alpha \operatorname{sign}\left(\nabla_{\delta^{t}} \mathcal{L}\left(I + \delta, \mathcal{S}(T) \oplus t\right)\right)\right), \quad (5)$$

where $\nabla_{\delta^t} \mathcal{L}$ is the gradient of the loss w.r.t δ^t . α is the step size and Proj() denotes the project of δ within the bound of the norm $(-\epsilon, \epsilon)$. $\mathcal{S}(\cdot)$ means shuffle the order of the clean caption samples in a batch with each iteration. For the text trigger minimization problem, we first initialize the trigger sequence by repeating the word "the" or "a" to the front of all inputs. In addition, we optimize the text trigger based on HotFlip [12], a method that approximates the effect of replacing a token using the gradient. We denote a text trigger as t_i which is a hot vector and can be embedded in the form e_i . Thus, we can update the embedding for every trigger token e_i to minimize the first-order Taylor approximation of CLIP loss around the current token embedding:

$$\underset{\mathbf{e}_{j}' \in \mathcal{V}}{\arg\min\left[\mathbf{e}_{j}' - \mathbf{e}_{i}\right]^{\top} \nabla_{\mathbf{e}_{i}} \mathcal{L}, \tag{6}$$

where \mathcal{V} is the set of all token embeddings in the vocabulary and $\nabla_{\mathbf{e}_t}$ is the gradient of loss w.r.t e_i . We can calculate a set of candidate tokens e'_j using a dot product with embedding of the vocabulary token and gradients. Finally, we can search for each optimal text trigger with a beam search in a set of candidate tokens. We consider the top-k candidates from Eq. 6 and search front to the last in every position in the trigger and score each beam using the loss on the current batch. We follow Wallace et al. [76] and use a small beam size for efficiency. In Fig. 3, we can see the framework of the generation of multimodal unlearnable examples with our MEM.

4 EXPERIMENT

4.1 Experiment Setup

Target models and datasets. Following previous works [2, 79, 80], we adopt the open-source implementation of CLIP in the community¹. For the architecture of the surrogate model, we choose the model with ResNet50 [30] as the image encoder and a Transformer [75] with certain architecture modifications serving as the text encoder. Our experiment involves three datasets: Flickr8K [81], Flickr30K [81], and MS-COCO [8]. Specifically, Flickr8K and Flickr30K consist of 8000 and 31,000 images, respectively, each accompanied

¹https://github.com/mlfoundations/open_clip

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia.

Dataset		Flic	k8k			Flic	x30k		MSCOCO			
Methods	Image \rightarrow Text		$Text \rightarrow Image$		Image \rightarrow Text		$Text \rightarrow Image$		Image \rightarrow Text		$Text \rightarrow Image$	
	Hit@10	Medr										
Random	1.2	727	0.9	490	1.0	711	1.0	498	0.2	3419	0.2	2491
Clean	22.7	58	18.5	60	46.5	12	42.7	16	65.1	5	58.3	7
EM [31]	13.8	117	15.9	86	9.6	197	7.9	170	1.7	1213	1.6	880
UAP [17]	12.9	110	16.5	76	36.7	26	35.6	24	63.8	5	57.1	7
MEM-3 (ours)	4.1	304	3.7	250	3.8	412	2.2	308	1.7	1705	1.3	1301
MEM-5 (ours)	3.8	308	3.0	275	3.9	445	2.0	325	0.8	1883	1.1	1466

Table 1: Comparison of the effectiveness with different unlearnable examples on several datasets.

Table 2: The transferability of MEM-3 generated on a ResNet50 model across different architectures models.

Dataset	Flick8k					Flic	k30k		MSCOCO			
Model	Image \rightarrow Text		$Text \rightarrow Image$		Image \rightarrow Text		$Text \rightarrow Image$		Image \rightarrow Text		$Text \rightarrow Image$	
	Dc	D_u	D_c	D_u								
RN50	58	308	60	275	12	445	16	325	5	1086	7	988
RN101	52	236	58	197	10	404	17	306	5	989	6	844
ViT-B/32	53	241	52	194	9	343	13	287	4	1064	4	920

by five captions. We adhere to the protocol established by [14, 39, 84], dividing the datasets into training/validation/testing sets with ratios of 6,000/1,000/1,000 and 29,000/1,000/1,000 for Flickr8K and Flickr30K, respectively. MS-COCO comprises 123,287 images, each annotated with five descriptions. In [39], MS-COCO underwent division into 82,783 training images, 5,000 validation images, and 5,000 test images. Despite their smaller size compared to the large-scale datasets used to train the original CLIP model, these datasets remain suitable for storage and computational resources and have found widespread usage in MCL studies [79, 80].

Training and Attack Setup. Initially, we focus on training models from scratch in the main experiments and reserve discussion of the fine-tuning scenario to the case study on face privacy protection. Following the training setup of the CLIP model in [2, 66], we choose a batch size of 128 and set the initial learning rate to 0.0005. The weight decay rate is set to 0.2, and we employ an Adam optimizer with decoupled weight decay regularization, while the learning rate decays using a cosine scheduler. We train these models from scratch on 1 A100 GPU for 32 epochs. During the attack stage, we extend the EM and UAP methods to MCL. In both cases, we assume that the surrogate models are CLIP models with the ResNet50, while the surrogate model for UAP is already trained on clean data. For our MEM, we choose an initial surrogate CLIP and iteratively optimize the noise and text trigger, stopping when the loss is below a threshold. We select a threshold of 0.01. Regarding the text trigger, we set the text sequence lengths as three and five, denoted MEM-3 and MEM-5, respectively. Typically, we set the noise bound of all methods with the l_{∞} -norm $\epsilon = 8/255$.

Evaluation Metrics. We evaluate the impact of data protection by examining the decline in performance in image and text retrieval tasks, drawing from previous work on attacks in MCL [79, 80]. Two metrics are employed to illustrate. **Hit@10** measures the proportion of all target images/texts that appear within the top 10 of the list of rank. Higher Hit@10 values indicate that many text/image samples successfully retrieving target images/texts early, reflecting a better rank list. **Medr** refers to the median position of all target images



Figure 4: Training loss curves and Testing metric Medr curves on Flick30K with different methods.

/ texts in the list of test images/texts. Lower Rank values signify earlier access to target images, indicative of a superior rank list. The performance of unlearnable multimodal examples is assessed using Hit@10 and Medr for image retrieval (Image \rightarrow Text) and text retrieval (Text \rightarrow Image) across all testing images. Lower Hit@10 and higher Medr signify more effective protection of the data.

4.2 Effectiveness and Transferability

In this section, we compare our method with extensions of existing popular unlearnable methods, including EM [31] and UAP [17]. Table 1 presents their retrieval results on different datasets. It is evident that UAP nearly fails to provide any protection for multimodal data, while EM demonstrates a certain level of protection. Moreover, as the size of the dataset increases, the effect of EM improves due to the denser caption distribution in the feature space, consistent with our analysis in Section 3.2. However, our MEM consistently offers robust protection for multimodal data, reducing retrieval performance to nearly half random guessing. In particular, MEM-5, with a longer text trigger, obtained greater results in reducing the performance of the hacker model compared to MEM-3, likely because longer text triggers enable the model to focus more effectively on the trigger and establish a shortcut.



Figure 5: Attention Maps Visualization: comparing four models on clean data and unlearnable examples with different methods.

Fig. 4 illustrates the descending curves of the training loss trained on unlearnable examples generated by different methods and the retrieval Medr on the clean test set. From (a), we observe that although EM enables the loss to fall faster compared to normal training, our method, MEM-3 and MEM-5, have a smaller loss on the first epoch, suggesting that the model can quickly learn the shortcuts. The UAP hardly speeds up the loss decline significantly, so it is understandably inefficient. From (b), we find that all models are trained with Medr decreasing compared to random guessing, but the model poisoned by MEM stops learning the fastest, reaches the worst result, and does not learn further as the epoch increases. The above observations are consistent with the results in Tab. 1.

We assume that data protection is a completely black-box setting, wherein the protector lacks knowledge of the hacker model's architecture. Thus, we evaluate the performance of our MEM generated on the ResNet50 surrogate model on different hacker models, including ResNet101, and ViT-B/32. The results are presented in Tab. 2. We find that these examples can be successfully transferred across different models and can degrade the performance of the CLIP model. Although they generated a surrogate model with a structure different from that of the hacker model can also exhibit the same protective effect, this observation aligns with the one of unlearnable examples in image classification. This is likely because the feature preferences learned by the models are similar.

4.3 Analysis

Attention of the models. Fig. 5 presents the heatmaps of images and text, illustrating the attention of models trained on clean and unlearnable examples. The Grad-CAM [70] is utilized to visualize the model attention for images, while the Integrated Gradients [73] is employed to visualize the attention to the text. Lighter colors represent higher attention from the model. Remarkably, for the image, models in (1), (2), and (3) all focus on the central region, which correlates with the captions. In contrast, model (4), trained on samples generated by MEM-3, fails to accurately recognize the clean image due to only learning noise features. Similarly in the text, models in (1), (2), and (3) all focus on the keywords of 'glass', while the model in (d) put the attention on the first three words, probably because MEM-3 always optimizes the noise and the first three text triggers to create shortcuts. These visualizations report that the EM and UAP are ineffective enough in protecting the multimodal data, while MEM-3 has an obvious effectiveness.

Visulaization of unlearnable examples. We visualize the feature distributions of clean samples under the normal model and the



Figure 6: t-SNE visualization of the clean samples and unlearnable examples for clean model and poisoned model.

feature distributions of unlearnable examples optimized by MEM-3 on the unlearned model in Fig 6. We represent image features with triangles and text features with circles, with the same color indicating five identical but transformed images and their corresponding different descriptions in the dataset. From (a), we observe that under the clean model, the same images and texts cluster together internally, and the corresponding image-text pairs are close to each other. However, in (b), there is a divergence between the same images and texts, with only the pairs being pairwise close to each other. This indicates that our methods effectively facilitate the model in learning the shortcuts between noise and textual triggers. Extension with Semantic Triggers. In the previous parts, we propose to use the gradient by the Hotflip method to select embeddings in the vocabulary list with large inner products to serve as candidate words for replacement. However, this approach may result in the generation of candidate words lacking semantic information, potentially impacting social reading or being removed by hackers. Here, we leverage the pre-training parameters of the BERT model to generate some semantically relevant substitutions as vocabulary lists. Subsequently, we embed these substitutions to obtain $\mathcal{V}_{\text{bert}}$, and then compute the inner product with the gradient and token embeddings of the substitutions, which yields replaced words with semantics. Table 3 displays examples of text triggers generated by our MEM-3 method under different datasets, including those generated with and without the BERT method. Text triggers generated with BERT exhibit greater naturalness and semantic coherence compared to those generated without BERT. Additionally, we show the Medr results of these triggers in text-image retrieval. It is observed that triggers generated without BERT have better protection performance, which is attributed to their larger search space, while

Dataset		Captions (red = text trigger)	$\text{Image} \rightarrow \text{Text}$	$Text \rightarrow Image$
Flickr8k	w/o BERT	And lopez a girl in a pink shirt has jumped into the air.	304	250
	with BERT	Next, suddenly a girl in a pink shirt has jumped into the air.	256	213
Flickr30k	w/o BERT	The baked instaa bride having a photo shoot outdoors around lots of people.	412	308
	with BERT	I was like a bride having a photo shoot outdoors around lots of people.	365	278
MSCOCO	w/o BERT	Conde matsu grinding a man and two dogs in the snow.	1705	1301
	with BERT	There he saw a man and two dogs in the snow.	1356	1028

Table 3: Comparison of text triggers generated by MEM-3 method with and w/o BERT, and their effect on protection.

Table 4: The protection effect of unlearnable examples generated on ResNet50 fine-tuning on different pre-trained models.

Model		RN	150			RN	101		ViT-B/32			
Metric	Image \rightarrow Text		$Text \rightarrow Image$		Image \rightarrow Text		$\text{Text} \rightarrow \text{Image}$		Image \rightarrow Text		$Text \rightarrow Image$	
	Hit@10	Medr	Hit@10	Medr	Hit@10	Medr	Hit@10	Medr	Hit@10	Medr	Hit@10	Medr
Pre-trained	2.6	77	2.6	77	3.3	83	3.3	83	3	87	3	87
Fine-tuned	80.7	1	80.7	1	85.3	1	85.3	1	94	1	94	1
MEM-3	6.7	74	6.7	74	6	66	6	66	93.3	1	93.3	1
MEM-5	8.7	82	8.7	82	12.7	58	12.7	58	91.3	1	91.3	1

triggers generated with BERT are restricted to a vocabulary space V_{bert} , resulting in inferior protection efficacy.

4.4 Case Study: Face-Privacy Protection

We conduct a case study to apply MEM to a real-world scenario: protecting personal face images and associated information on social media platforms, such as names. It is arguably one of the most common multimodal privacy protection scenarios. In previous parts, we assumed that hackers train models with unlearnable examples from scratch, but in the real-world, hackers will opt to fine-tune pre-trained models provided by the community. Therefore, here we assume that the protector aims to prevent their facial images and names from being fine-tuned with a pre-trained CLIP model. We adopt the Open AI's released parameters to initialize pre-trained models, and still assume the protector has access to their data to generate unlearnable examples before sharing them on on-line platforms. However, they are unable to know the hacker's model architecture and the parameters. In this black-box scenario, we explore whether MEM can still prevent model from learning.

We conducted experiments using the PubFig [40], a large realworld face dataset comprising 58,797 images of 200 individuals. For retrieval evaluation, we randomly selected one photo of each celebrity as the test set and utilized remaining images for training. It's noteworthy that the pre-trained CLIP model already possesses knowledge of the real names and faces of these celebrities. For authentic fine-tuning, we altered their names and provided a set of text templates for caption generation, as shown in Fig. 7. Then we generated unlearnable examples using MEM and employed a pre-trained surrogate model with ResNet50. We fine-tune the pretrained models for 10 epochs and a small learning rate of 0.00001 and evaluate them with different hacker models. The results are presented in Tab. 4. For comparison, we also test the results of the pre-trained models and the benign fine-tuned models. We find the initial model shows poor performance on the test set, since the face and name features are not aligned. However, with only 10 epochs of fine-tuning, the Medr of its retrieved results reaches

1. This could be because the pre-trained model has a strong feature representation, enabling it to rapidly adapt to the new task in the fine-tuning process. Our MEM can prevent these models from learning the correlation between face and name features, thereby impeding accurate person retrieval on the test set. These unlearnable examples were generated by ResNet50 surrogate model and effectively work on the ResNet101 targeted model. However, their efficacy is reduced on ViT, possibly due to variations in architecture and initialization parameters.



Figure 7: Illustration of the face-privacy protection pipeline.

5 CONCLUSION

In this paper, we explore multimodal data protection, particularly focusing on image-text pairs, wherein we generate multimodal unlearnable examples to prevent exploitation by MCL. We extend previous classification methods to this context, revealing their limitation in MCL due to the increased modality and data dispersion. In light of these findings, we introduce a novel generation approach named Multi-step Error Minimization (MEM), which is based on the framework of EM. MEM effectively establishes a shortcut between noise and textual triggers and demonstrates transferability across different hacker models. Additionally, we utilize various visualization tools to validate the effectiveness of our approach. Our work opens up a new direction, and it is expected to be applicable to other modal pairs, such as audio-text and audio-image pairs. Multimodal Unlearnable Examples: Protecting Data against Multimodal Contrastive Learning

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia.

ACKNOWLEDGMENTS

Supported in part by National Natural Science Foundation of China (No. 62025604)

REFERENCES

- Yang Bai, Yuyuan Zeng, Yong Jiang, Shu-Tao Xia, Xingjun Ma, and Yisen Wang. [n. d.]. Improving Adversarial Robustness via Channel-wise Activation Suppressing. In International Conference on Learning Representations.
- [2] Hritik Bansal, Nishad Singhi, Yu Yang, Fan Yin, Aditya Grover, and Kai-Wei Chang. 2023. Cleanclip: Mitigating data poisoning attacks in multimodal contrastive learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 112–123.
- [3] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. PAMI 35, 8 (2013), 1798–1828.
- [4] Nicholas Carlini and Andreas Terzis. 2021. Poisoning and backdooring contrastive learning. arXiv preprint arXiv:2106.09667 (2021).
- [5] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. 2021. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 3558–3568.
- [6] Jianbo Chen, Xinwei Liu, Siyuan Liang, Xiaojun Jia, and Yuan Xun. 2023. Universal Watermark Vaccine: Universal Adversarial Perturbations for Watermark Protection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops.
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [8] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. arXiv preprint arXiv:1504.00325 (2015).
- [9] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. arXiv preprint arXiv:1712.05526 (2017).
- [10] Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. In 2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05), Vol. 1. IEEE, 539-546.
- [11] Xin Dong, Rui Wang, Siyuan Liang, Aishan Liu, and Lihua Jing. 2023. Face Encryption via Frequency-Restricted Identity-Agnostic Attacks. In Proceedings of the 31st ACM International Conference on Multimedia.
- [12] Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2017. Hotflip: Whitebox adversarial examples for text classification. arXiv preprint arXiv:1712.06751 (2017).
- [13] Privacy enhancing face obfuscation guided by semantic-aware attribution maps. 2023. Privacy-enhancing face obfuscation guided by semantic-aware attribution maps. *IEEE Transactions on Information Forensics and Security* (2023).
- [14] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2017. Vse++: Improving visual-semantic embeddings with hard negatives. arXiv preprint arXiv:1707.05612 (2017).
- [15] Bin Fang, Bo Li, Shuang Wu, Ran Yi, Shouhong Ding, and Lizhuang Ma. 2023. Re-thinking Data Availablity Attacks Against Deep Neural Networks. arXiv preprint arXiv:2305.10691 (2023).
- [16] Ji Feng, Qi-Zhi Cai, and Zhi-Hua Zhou. 2019. Learning to confuse: Generating training time adversarial data with auto-encoder. Advances in Neural Information Processing Systems 32 (2019).
- [17] Liam Fowl, Micah Goldblum, Ping-yeh Chiang, Jonas Geiping, Wojciech Czaja, and Tom Goldstein. 2021. Adversarial examples make strong poisons. Advances in Neural Information Processing Systems 34 (2021), 30339–30351.
- [18] Shaopeng Fu, Fengxiang He, Yang Liu, Li Shen, and Dacheng Tao. 2022. Robust unlearnable examples: Protecting data against adversarial learning. arXiv preprint arXiv:2203.14533 (2022).
- [19] Kuofeng Gao, Yang Bai, Jindong Gu, Yong Yang, and Shu-Tao Xia. 2023. Backdoor defense via adaptively splitting poisoned dataset. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4005–4014.
- [20] Yinghua Gao, Yiming Li, Xueluan Gong, Zhifeng Li, Shu-Tao Xia, and Qian Wang. 2024. Backdoor Attack with Sparse and Invisible Trigger. *IEEE Transactions on Information Forensics and Security* (2024).
- [21] Jindong Gu, Xiaojun Jia, Pau de Jorge, Wenqain Yu, Xinwei Liu, Avery Ma, Yuan Xun, Anjun Hu, Ashkan Khakzar, Zhijiang Li, et al. 2023. A survey on transferability of adversarial examples across deep neural networks. arXiv preprint arXiv:2310.17626 (2023).
- [22] Jindong Gu, Volker Tresp, and Yao Qin. 2022. Are vision transformers robust to patch perturbations?. In *European Conference on Computer Vision*. Springer, 404–421.

- [23] Jindong Gu, Baoyuan Wu, and Volker Tresp. 2021. Effective and Efficient Vote Attack on Capsule Networks. In International Conference on Learning Representations (ICLR).
- [24] Jindong Gu, Hengshuang Zhao, Volker Tresp, and Philip HS Torr. 2022. Segpgd: An effective and efficient adversarial attack for evaluating and boosting segmentation robustness. In *European Conference on Computer Vision*. Springer, 308–325.
- [25] Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. 2019. Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access* 7 (2019), 47230–47244.
- [26] Jun Guo, Xingyu Zheng, Aishan Liu, Siyuan Liang, Yisong Xiao, Yichao Wu, and Xianglong Liu. 2023. Isolation and Induction: Training Robust Deep Neural Networks against Model Stealing Attacks. In Proceedings of the 31st ACM International Conference on Multimedia.
- [27] Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In 2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06), Vol. 2. IEEE, 1735–1742.
- [28] Hao He, Kaiwen Zha, and Dina Katabi. 2022. Indiscriminate poisoning attacks on unsupervised contrastive learning. arXiv preprint arXiv:2202.11202 (2022).
- [29] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 9729–9738.
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition. 770–778.
- [31] Hanxun Huang, Xingjun Ma, Sarah Monazam Erfani, James Bailey, and Yisen Wang. 2021. Unlearnable examples: Making personal data unexploitable. *ICLR* (2021).
- [32] Yihao Huang, Qing Guo, Felix Juefei-Xu, Ming Hu, Xiaojun Jia, Xiaochun Cao, Geguang Pu, and Yang Liu. 2024. Texture Re-scalable Universal Adversarial Perturbation. *IEEE Transactions on Information Forensics and Security* (2024). https://doi.org/10.1109/TIFS.2024.3416030
- [33] Yihao Huang, Qing Guo, Felix Juefei-Xu, Lei Ma, Weikai Miao, Yang Liu, and Geguang Pu. 2021. AdvFilter: predictive perturbation-aware filtering against adversarial attack via multi-domain learning. In Proceedings of the 29th ACM International Conference on Multimedia. 395–403.
- [34] Yihao Huang, Felix Juefei-Xu, Qing Guo, Jie Zhang, Yutong Wu, Ming Hu, Tianlin Li, Geguang Pu, and Yang Liu. 2024. Personalization as a shortcut for few-shot backdoor attack against text-to-image diffusion models. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 38. 21169–21178.
- [35] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and visionlanguage representation learning with noisy text supervision. In *International conference on machine learning*. PMLR, 4904–4916.
- [36] Xiaojun Jia, Yuefeng Chen, Xiaofeng Mao, Ranjie Duan, Jindong Gu, Rong Zhang, Hui Xue, Yang Liu, and Xiaochun Cao. 2024. Revisiting and exploring efficient fast adversarial training via law: Lipschitz regularization and auto weight averaging. IEEE Transactions on Information Forensics and Security (2024).
- [37] Xiaojun Jia, Yong Zhang, Xingxing Wei, Baoyuan Wu, Ke Ma, Jue Wang, and Xiaochun Cao. 2024. Improving fast adversarial training with prior-guided knowledge. IEEE Transactions on Pattern Analysis and Machine Intelligence (2024).
- [38] Xiaojun Jia, Yong Zhang, Baoyuan Wu, Ke Ma, Jue Wang, and Xiaochun Cao. 2022. LAS-AT: adversarial training with learnable attack strategy. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 13398–13408.
- [39] Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In Proceedings of the IEEE conference on computer vision and pattern recognition. 3128–3137.
- [40] Neeraj Kumar, Alexander C Berg, Peter N Belhumeur, and Shree K Nayar. 2009. Attribute and simile classifiers for face verification. In 2009 IEEE 12th international conference on computer vision. IEEE, 365–372.
- [41] Iro Laina, Christian Rupprecht, and Nassir Navab. 2019. Towards unsupervised image captioning with shared multimodal embeddings. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 7414–7424.
- [42] Haoheng Lan, Jindong Gu, Philip Torr, and Hengshuang Zhao. 2024. Influencer backdoor attack on semantic segmentation. *ICLR* (2024).
- [43] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In International conference on machine learning. PMLR, 12888–12900.
- [44] Longkang Li, Siyuan Liang, Zihao Zhu, Chris Ding, Hongyuan Zha, and Baoyuan Wu. 2024. Learning to Optimize Permutation Flow Shop Scheduling via Graphbased Imitation Learning. *Proceedings of the AAAI Conference on Artificial Intelli*gence (2024).
- [45] Yiming Li, Baoyuan Wu, Yan Feng, Yanbo Fan, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. 2022. Semi-supervised robust training with generalized perturbed neighborhood. *Pattern Recognition* 124 (2022), 108472.
- [46] Jiawei Liang, Siyuan Liang, Aishan Liu, Xiaojun Jia, Junhao Kuang, and Xiaochun Cao. 2024. Poisoned forgery face: Towards backdoor attacks on face forgery detection. arXiv preprint arXiv:2402.11473 (2024).

MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia.

- [47] Jiawei Liang, Siyuan Liang, Man Luo, Aishan Liu, Dongchen Han, Ee-Chien Chang, and Xiaochun Cao. 2024. VL-Trojan: Multimodal Instruction Backdoor Attacks against Autoregressive Visual Language Models. arXiv preprint arXiv:2402.13851 (2024).
- [48] Siyuan Liang, Longkang Li, Yanbo Fan, Xiaojun Jia, Jingzhi Li, Baoyuan Wu, and Xiaochun Cao. 2022. A large-scale multiple-objective method for black-box attack against object detection. In *European Conference on Computer Vision*.
- [49] Siyuan Liang, Aishan Liu, Jiawei Liang, Longkang Li, Yang Bai, and Xiaochun Cao. 2022. Imitated detectors: Stealing knowledge of black-box object detectors. In Proceedings of the 30th ACM International Conference on Multimedia.
- [50] Siyuan Liang, Kuanrong Liu, Jiajun Gong, Jiawei Liang, Yuan Xun, Ee-Chien Chang, and Xiaochun Cao. 2024. Unlearning Backdoor Threats: Enhancing Backdoor Defense in Multimodal Contrastive Learning via Local Token Unlearning. arXiv preprint arXiv:2403.16257 (2024).
- [51] Siyuan Liang, Wei Wang, Ruoyu Chen, Aishan Liu, Boxi Wu, Ee-Chien Chang, Xiaochun Cao, and Dacheng Tao. 2024. Object Detectors in the Open Environment: Challenges, Solutions, and Outlook. arXiv preprint arXiv:2403.16271 (2024).
- [52] Siyuan Liang, Xingxing Wei, and Xiaochun Cao. 2021. Generate more imperceptible adversarial examples for object detection. In ICML 2021 Workshop on Adversarial Machine Learning.
- [53] Siyuan Liang, Xingxing Wei, Siyuan Yao, and Xiaochun Cao. 2020. Efficient adversarial attacks for visual object tracking. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16.
- [54] Siyuan Liang, Baoyuan Wu, Yanbo Fan, Xingxing Wei, and Xiaochun Cao. 2022. Parallel rectangle flip attack: A query-based black-box attack against object detection. arXiv preprint arXiv:2201.08970 (2022).
- [55] Siyuan Liang, Mingli Zhu, Aishan Liu, Baoyuan Wu, Xiaochun Cao, and Ee-Chien Chang. 2023. Badclip: Dual-embedding guided backdoor attack on multimodal contrastive learning. arXiv preprint arXiv:2311.12075 (2023).
- [56] Aishan Liu, Xinwei Zhang, Yisong Xiao, Yuguang Zhou, Siyuan Liang, Jiakai Wang, Xianglong Liu, Xiaochun Cao, and Dacheng Tao. 2023. Pre-trained trojan attacks for visual recognition. arXiv preprint arXiv:2312.15172 (2023).
- [57] Xinwei Liu, Xiaojun Jia, Jindong Gu, Yuan Xun, Siyuan Liang, and Xiaochun Cao. 2024. Does few-shot learning suffer from backdoor attacks?. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 38. 19893–19901.
- [58] Xinwei Liu, Jian Liu, Yang Bai, Jindong Gu, Tao Chen, Xiaojun Jia, and Xiaochun Cao. 2022. Watermark vaccine: Adversarial attacks to prevent watermark removal. In European Conference on Computer Vision. Springer, 1–17.
- [59] Ke Ma, Qianqian Xu, Jinshan Zeng, Xiaochun Cao, and Qingming Huang. 2021. Poisoning attack against estimating from pairwise comparisons. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 10 (2021), 6393–6408.
- [60] Ke Ma, Qianqian Xu, Jinshan Zeng, Guorong Li, Xiaochun Cao, and Qingming Huang. 2022. A tale of hodgerank and spectral method: Target attack against rank aggregation is the fixed point of adversarial game. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 4 (2022), 4090–4108.
- [61] Ke Ma, Qianqian Xu, Jinshan Zeng, Wei Liu, Xiaochun Cao, Yingfei Sun, and Qingming Huang. 2024. Sequential manipulation against rank aggregation: theory and algorithm. *IEEE transactions on pattern analysis and machine intelligence* (2024).
- [62] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083 (2017).
- [63] Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. Clipcap: Clip prefix for image captioning. arXiv preprint arXiv:2111.09734 (2021).
- [64] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018).
- [65] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. 2021. Styleclip: Text-driven manipulation of stylegan imagery. In Proceedings of the IEEE/CVF international conference on computer vision. 2085–2094.
- [66] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In International conference on machine learning. PMLR, 8748–8763.
- [67] Vinu Sankar Sadasivan, Mahdi Soltanolkotabi, and Soheil Feizi. 2023. Cuda: Convolution-based unlearnable datasets. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 3862–3871.
- [68] Pedro Sandoval-Segura, Vasu Singla, Jonas Geiping, Micah Goldblum, Tom Goldstein, and David Jacobs. 2022. Autoregressive perturbations for data poisoning. Advances in Neural Information Processing Systems 35 (2022), 27374–27386.
- [69] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. Advances in Neural Information Processing Systems 35 (2022), 25278–25294.
- [70] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from

deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision. 618–626.

- [71] Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, and Tom Goldstein. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. Advances in neural information processing systems 31 (2018).
- [72] Juncheng Shen, Xiaolei Zhu, and De Ma. 2019. TensorClog: An imperceptible poisoning attack on deep neural network applications. *IEEE Access* 7 (2019), 41498–41506.
- [73] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International conference on machine learning*. PMLR, 3319– 3328.
- [74] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. Yfcc100m: The new data in multimedia research. *Commun. ACM* 59, 2 (2016), 64–73.
- [75] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems 30 (2017).
- [76] Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing NLP. arXiv preprint arXiv:1908.07125 (2019).
- [77] Xingxing Wei, Siyuan Liang, Ning Chen, and Xiaochun Cao. 2018. Transferable adversarial attacks for image and video object detection. arXiv preprint arXiv:1811.12641 (2018).
- [78] Yuan Xun, Xiaojun Jia, Jindong Gu, Xinwei Liu, Qing Guo, and Xiaochun Cao. 2024. Minimalism is King! High-Frequency Energy-based Screening for Data-Efficient Backdoor Attacks. *IEEE Transactions on Information Forensics and Security* (2024).
- [79] Wenhan Yang, Jingdong Gao, and Baharan Mirzasoleiman. 2024. Robust Contrastive Language-Image Pretraining against Data Poisoning and Backdoor Attacks. Advances in Neural Information Processing Systems 36 (2024).
- [80] Ziqing Yang, Xinlei He, Zheng Li, Michael Backes, Mathias Humbert, Pascal Berrang, and Yang Zhang. 2023. Data poisoning attacks against multimodal encoders. In International Conference on Machine Learning. PMLR, 39299–39313.
- [81] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* 2 (2014), 67–78.
- [82] Da Yu, Huishuai Zhang, Wei Chen, Jian Yin, and Tie-Yan Liu. 2022. Availability attacks create shortcuts. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 2367–2376.
- [83] Chia-Hung Yuan and Shan-Hung Wu. 2021. Neural tangent generalization attacks. In International Conference on Machine Learning. PMLR, 12230–12240.
- [84] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z Li. 2020. Context-aware attention network for image-text retrieval. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 3536–3545.
- [85] Xinwei Zhang, Aishan Liu, Tianyuan Zhang, Siyuan Liang, and Xianglong Liu. 2024. Towards Robust Physical-world Backdoor Attacks on Lane Detection. arXiv preprint arXiv:2405.05553 (2024).
- [86] Mingli Zhu, Siyuan Liang, and Baoyuan Wu. 2024. Breaking the False Sense of Security in Backdoor Defense through Re-Activation Attack. arXiv preprint arXiv:2405.16134 (2024).