

Dataset Ownership Verification for Pre-trained Masked Models

Yuechen Xie¹, Jie Song^{1*}, Yicheng Shan², Xiaoyan Zhang¹,
Yuanyu Wan¹, Shengxuming Zhang¹, Jiarui Duan¹, Mingli Song^{1,3,4}

¹Zhejiang University, ²The University of Sydney

³State Key Laboratory of Blockchain and Security, Zhejiang University

⁴Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security

{xyuechen, sjie, zhang_xy99, wanyy, zsxm1998, jerryduan, brooksong}@zju.edu.cn
ysha4092@uni.sydney.edu.au

Abstract

High-quality open-source datasets have emerged as a pivotal catalyst driving the swift advancement of deep learning, while facing the looming threat of potential exploitation. Protecting these datasets is of paramount importance for the interests of their owners. The verification of dataset ownership has evolved into a crucial approach in this domain; however, existing verification techniques are predominantly tailored to supervised models and contrastive pre-trained models, rendering them ill-suited for direct application to the increasingly prevalent masked models. In this work, we introduce the inaugural methodology addressing this critical, yet unresolved challenge, termed Dataset Ownership Verification for Masked Modeling (DOV4MM). The central objective is to ascertain whether a suspicious black-box model has been pre-trained on a particular unlabeled dataset, thereby assisting dataset owners in safeguarding their rights. DOV4MM is grounded in our empirical observation that when a model is pre-trained on the target dataset, the difficulty of reconstructing masked information within the embedding space exhibits a marked contrast to models not pre-trained on that dataset. We validated the efficacy of DOV4MM through ten masked image models on ImageNet-1K and four masked language models on WikiText-103. The results demonstrate that DOV4MM rejects the null hypothesis, with a p -value considerably below 0.05, surpassing all prior approaches.

1. Introduction

High-quality open-source datasets [11, 26] are the cornerstone of breakthroughs in the field of deep learning, providing valuable resources for researchers and developers

*Corresponding author

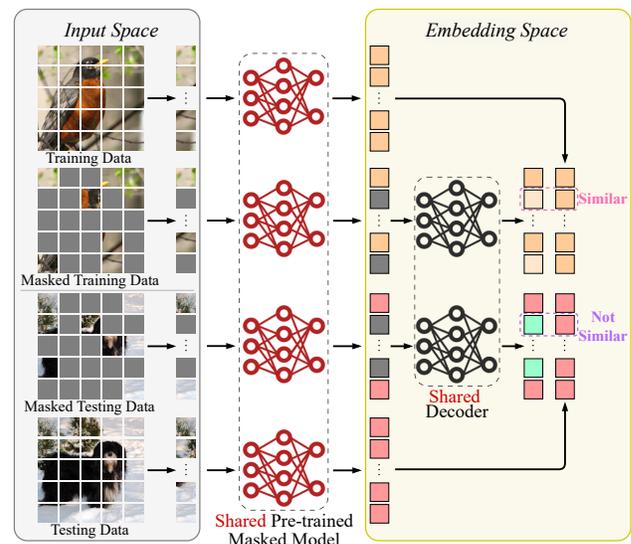


Figure 1. The overview of DOV4MM’s motivation. In the embedding space of pre-trained masked models, the reconstruction difficulty of masked information for seen samples in the pre-training phase is lower than that for unseen samples.

worldwide and driving rapid technological progress. Many open-source datasets are explicitly designated for academic use only, with restrictions against unauthorized commercial exploitation. Yet, as the value of data continues to surge, these datasets increasingly face the threat of misuse. To preserve the legitimate rights of data proprietors and thwart malicious theft or exploitation, it is imperative to secure the integrity and usage rights of open-source datasets.

Recently, dataset ownership verification (DOV) [17, 18, 30, 31, 51] has emerged as a defense mechanism to prevent dataset misuse, aiming to protect datasets from being stolen. This technique helps defenders, *i.e.*, dataset owners, detect whether a suspicious black-box model has been trained on their dataset, thus determining if their rights have

been infringed. However, most existing DOV methods are designed for supervised learning models [17, 18, 30, 31], where verification relies on the relationship between data points and decision boundaries. Furthermore, these methods [17, 30, 31] are heavily reliant on backdoor watermarks. Specifically, models trained on watermarked datasets are embedded with a pre-designed backdoor, allowing defenders to verify dataset ownership by activating the backdoor. However, this strategy not only compromises model performance but also renders the model susceptible to watermark removal techniques [20, 35, 46, 48, 55]. To overcome these two challenges, [52] introduced a new DOV technique for contrastive pre-trained models that does not rely on backdoor watermarks, where verification relies on the contrastive relationship gap in the embedding space. Nonetheless, due to the significant differences between the proxy tasks of masked modeling [22, 57] and contrastive learning [5, 7], the representations of pre-trained masked models are harder to distinguish [58], leading to less noticeable contrastive relationship gaps in their representations. As a result, the method in [52] is ineffective for pre-trained masked models, which are widely used in both computer vision [1, 22, 57] and natural language processing [12, 27, 34].

In this work, we introduce the pioneering DOV method, termed DOV4MM, designed specifically for pre-trained masked models, to address this important yet unexplored challenge. Notably, DOV4MM operates without relying on backdoor watermarks. It assists defenders in verifying whether a suspicious model has been pre-trained on their public datasets. DOV4MM concentrates on the black-box scenario, wherein defenders lack any insight into the model’s pre-training configurations (e.g., loss functions and model architecture) and can only interact with the model via Encoder-as-a-Service (EaaS) [36, 43]. It means that defenders can only retrieve feature vectors through the model’s API. DOV4MM is grounded in an empirical observation, as depicted in Fig. 1. In the embedding space of the pre-trained masked model, the difficulty of reconstructing masked information from seen samples during the pre-training phase is markedly lower than that for unseen samples.

We propose the concept *relative embedding reconstruction difficulty* based on this observation. The defenders can exploit the difference of relative embedding reconstruction difficulty between seen and unseen samples in the suspicious model, thereby determining whether the suspicious model has been pre-trained on their data. More specifically, as illustrated in Fig. 2, DOV4MM comprises three key steps: (1) Randomly partitioning the public dataset into two disjoint subsets, namely the training dataset and the validation dataset. The decoder is then trained using the training dataset, enabling it to reconstruct masked information in the embedding space; (2) Employing the decoder to compute

the embedding reconstruction difficulties of the suspicious model for the training dataset, the validation dataset, and a private dataset (undisclosed to the defender) respectively. These difficulties are then used to calculate the relative embedding reconstruction difficulty; (3) Conducting a one-tailed pairwise t-test [23] on the relative embedding reconstruction difficulty between the validation dataset and the private dataset, to determine whether the suspicious model has been pre-trained on the defender’s public dataset.

In summary, this paper presents three key contributions: (1) We observe that when a model is pre-trained on a specific target dataset, the difficulty of reconstructing masked information in the embedding space shows significant discrepancies when compared to a model not pre-trained on that dataset; (2) We introduce the concept of relative embedding reconstruction difficulty and propose a novel DOV technique, DOV4MM. To the best of our knowledge, this is the first DOV technique tailored for pre-trained masked models; (3) Extensive experiments validate that DOV4MM effectively rejects the null hypothesis, with a p -value significantly below 0.05, surpassing all previous methodologies.

2. Related Work

2.1. Data Protection

Dataset ownership verification. Dataset ownership verification is an emerging field in data security. Typically, it involves embedding watermarks into the original dataset [17, 18, 30, 31, 47]. Models trained on the watermarked dataset will be implanted a pre-designed backdoor, allowing defenders to verify dataset ownership simply by triggering it. However, these DOV methods primarily target supervised models, which makes it unable to adapt to the increasingly popular self-supervised models [1, 5–7, 22, 57]. Recently, backdoor attack methods targeting self-supervised models have been emerging [3, 28, 42, 54], and they are expected to become reliable DOV techniques for self-supervised models. However, like previous backdoor watermarking methods [17, 30, 31], they require altering the original dataset’s distribution to inject watermarks, which makes it susceptible to various watermark removal mechanisms [20, 35, 46, 48, 55]. A new DOV technique [52] is proposed for contrastive pre-trained models, which does not rely on backdoor watermarks. However, due to the differing proxy tasks of masked modeling and contrastive learning, the representations learned by pre-trained masked models are more difficult to distinguish [58]. Therefore, [52] cannot be applied to pre-trained masked models directly. To fill this gap, we propose a DOV technique for pre-trained masked models, which also demonstrates that for these models, dataset can be effectively protected without relying on backdoor watermark.

Dataset inference. Dataset inference [39] is a state-of-the-

art defense approach for preventing model stealing [43]. The latest dataset inference method [14] has expanded its application to self-supervised learning. The intuition behind it is that the log-likelihood of an encoder’s output representations is higher on the victim’s training data than on test data if it is stolen from the victim. Although it’s aimed at encoder theft, it can also be directly used for DOV. However, it requires inferring the entire dataset to model the features of all data. It is prohibitively time-consuming for large datasets, such as ImageNet-1K [11]. In contrast, our method achieves accurate verification using only a small fraction of the dataset. For example, we use only 3% of all training data for accurate verification on ImageNet-1K, which significantly reduces computational cost.

Membership inference. Membership inference [4, 9, 44, 59] endeavors to determine whether an input was part of the model’s training dataset. At present, PartCrop [59] is a powerful methodology designed for membership inference on visual self-supervised encoder, which takes advantage of the shared part-aware capability among models and stronger part response on the training data. However, it directly relies on the entire high-dimensional representation for membership inference, which contains a large amount of redundant information. In contrast, our method extracts the most critical information for verification from the representations, namely relative embedding reconstruction difficulty, achieving effective verification.

2.2. Masked Modeling

Masked modeling [12, 22, 34, 50] is a self-supervised learning technique which encourages the model to learn useful representations from the unmasked portions of the data. The pre-trained model is applicable to a variety of downstream tasks. Masked modeling techniques originated and developed in natural language processing [10, 12, 34]. Due to performance limitations, in the early stages of computer vision, masked modeling had a minimal impact compared to contrastive learning [5, 7]. However, once vision Transformers [13, 38] surpassed convolutional neural networks [21, 45] in performance, masked modeling has regained attention as an effective pre-training method for vision Transformers [29], making it one of the mainstream visual self-supervised paradigms alongside contrastive learning. DOV4MM is focused on safeguarding the unlabeled datasets used in pre-trained masked models, ensuring that they are not misused by suspects, thereby securing and fostering the healthy advancement of this field.

3. Method

3.1. Problem Formulation

In this study, we focus on the dataset ownership verification problem in black-box scenarios, where two key players are

involved: the *defender* and the *suspect*. The defender, assuming the role of the dataset provider, endeavors to ascertain whether the suspicious model, M_s , has been unlawfully pre-trained on his public dataset \mathcal{D}_{pub} .

M_s can be classified as either illegal or legal based on its training datasets: (1) *Illegal*: M_s is pre-trained on the data from \mathcal{D}_{pub} , indicating the occurrence of dataset misappropriation; (2) *Legal*: M_s is pre-trained on an unrelated dataset outside the scope of \mathcal{D}_{pub} , indicating the innocence of the suspect. More specifically, for a black-box suspect model, the defender’s objective is to determine whether it is illegal or legal based solely on the feature vectors it outputs.

3.2. Relative Embedding Reconstruction Difficulty

3.2.1. Observations and Definitions

In masked modeling, a key training objective is to hide a portion of the input data (such as patches of an image or words in a text) and train the model to predict the missing parts, thereby learning the underlying structure and useful representations of the data. This training approach leverages the memory capabilities of neural networks, encouraging the model to retain features of the training data. Therefore, we derive the following important insights:

Observation 1. *In the embedding space of pre-trained masked models, the reconstruction difficulty of masked information for seen samples during pre-training is lower than unseen samples.*

We characterize the reconstruction difficulty of masked information in the embedding space as the model’s *embedding reconstruction difficulty*. The definition of embedding reconstruction difficulty is as follows:

Definition 1 (Embedding Reconstruction Difficulty). *Given a pre-trained masked model $M : \mathbb{R}^m \rightarrow \mathbb{R}^n$ and a dataset $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^{|\mathcal{D}|}$, where m is the dimension of the input space (e.g., for an image input, $m = C \times W \times H$), n is the dimension of the embedding space, and $|\mathcal{D}|$ is the number of samples in \mathcal{D} . The embedding reconstruction difficulty of M on \mathcal{D} is defined as:*

$$\mathcal{R} = \{\bar{R}_k(\mathcal{D}^k, M) | k \in [1, K], \mathcal{D}^k \in \mathcal{D}\} \quad (1)$$

K is the iterations of sampling, \mathcal{D}^k is the subset obtained by random sampling from \mathcal{D} in the k -th iteration, and $\bar{R}_k = \frac{1}{|\mathcal{D}^k|} \sum_{i=1}^{|\mathcal{D}^k|} R_i$ is the average embedding reconstruction difficulty of M over all samples in \mathcal{D}^k .

Given that the difference in embedding reconstruction difficulty between seen and unseen samples during the pre-training phase may not be significant (even though the former is lower than the latter), especially when the reconstruction difficulty for both is relatively high, we introduce the *relative embedding reconstruction difficulty*, a metric that can aid defenders in discerning whether the queried model

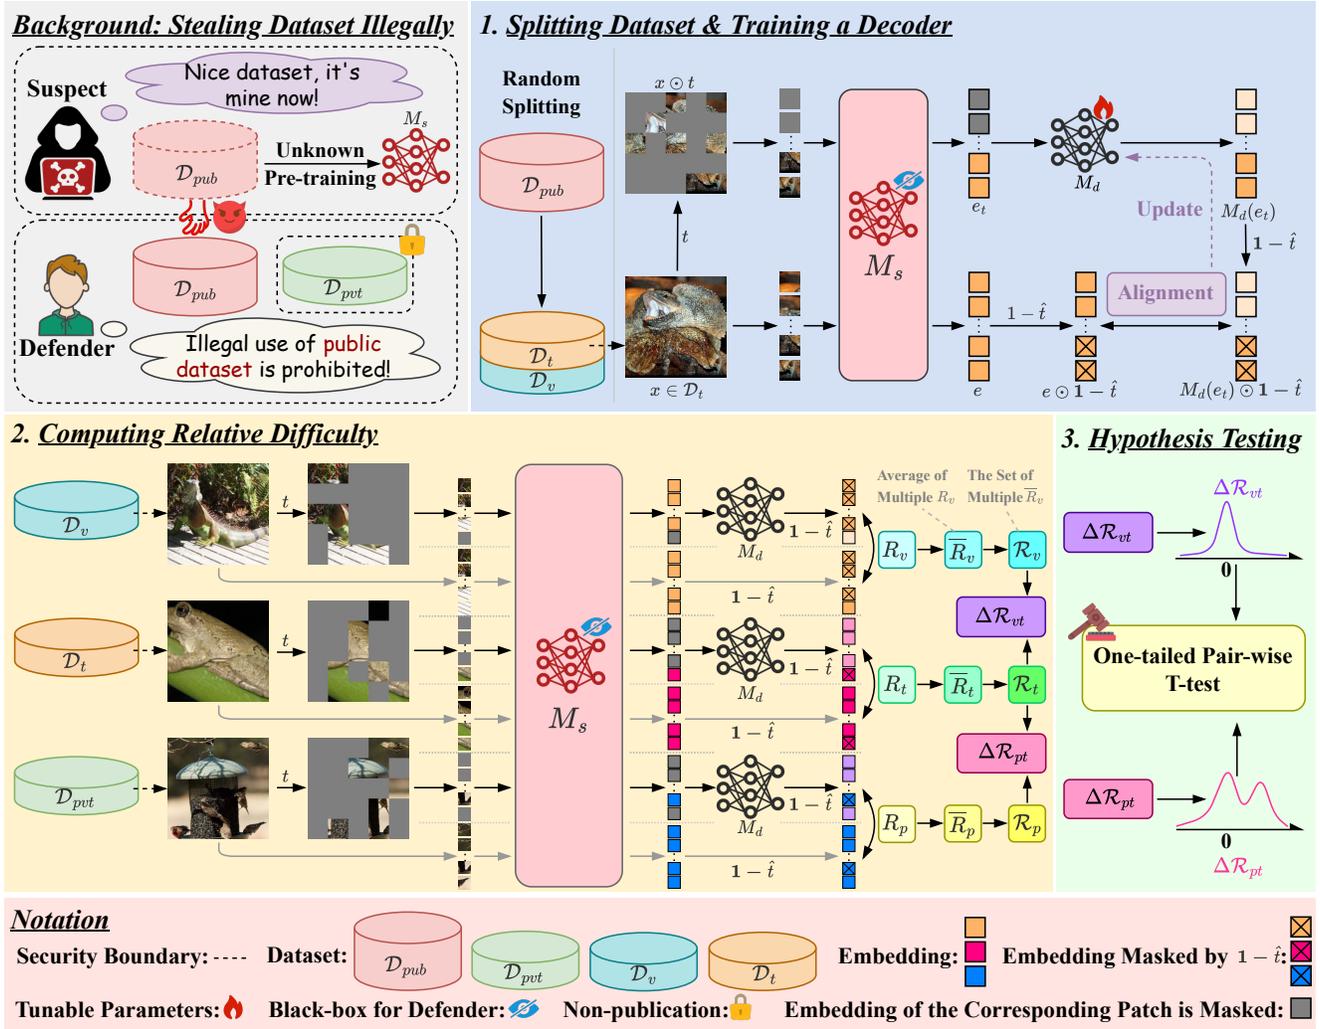


Figure 2. The overview of DOV4MM (best viewed under color conditions).

has been pre-trained on their dataset. The definition of relative embedding reconstruction difficulty is as follows:

Definition 2 (Relative Embedding Reconstruction Difficulty). Given the embedding reconstruction difficulties \mathcal{R} and \mathcal{R}' of a pre-trained masked model M on datasets \mathcal{D} and \mathcal{D}' , respectively, the relative embedding reconstruction difficulty of model M on dataset \mathcal{D}' , with \mathcal{D} as the standard, is defined as:

$$\Delta \mathcal{R} = \{ \bar{R}'_k - \bar{R}_k \mid k \in [1, K], \bar{R}_k \in \mathcal{R}, \bar{R}'_k \in \mathcal{R}' \} \quad (2)$$

K is the iterations of sampling, same as in Eq. (1), \bar{R}_k and \bar{R}'_k are the k -th elements of \mathcal{R} and \mathcal{R}' respectively.

3.2.2. The Calculation of $\Delta \mathcal{R}$

To compute $\Delta \mathcal{R}$, we first calculate the embedding reconstruction difficulty of the pre-training masked model M on a single sample x . To this end, we introduce two masks t and \hat{t} , along with a decoder M_d .

$t \in \{0, 1\}^m$ is a random mask in input space \mathbb{R}^m , where the positions corresponding to 0 are masked, and the positions corresponding to 1 are retained. $\hat{t} \in \{0, 1\}^n$ is the mask in embedding space \mathbb{R}^n that corresponds to t . Such as, for an image, t masks certain patches in it, while \hat{t} masks the tokens corresponding to the patches masked by t .

$M_d : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a decoder that can reconstruct the missing information in the embedding space \mathbb{R}^n . It uses the output embeddings from the pre-training masked model M for embedding reconstruction training. Note that M_d corresponds one-to-one with M . Specifically, a decoder M_d trained based on M can effectively reconstruct the masked embeddings of M , but this may not hold for other pre-training masked models.

We use the above t , \hat{t} , and M_d to compute the embedding

reconstruction difficulty R of M on a sample \mathbf{x} , as follows:

$$R(\mathbf{x}, \mathbf{t}, \hat{\mathbf{t}}, M, M_d) = \frac{\| [M_d(e_{\mathbf{t}}) - \mathbf{e}] \odot (\mathbf{1} - \hat{\mathbf{t}}) \|_2^2}{\| \mathbf{1} - \hat{\mathbf{t}} \|_1} \quad (3)$$

where \odot is the element-wise multiplication, $\mathbf{e} = M(\mathbf{x})$, $e_{\mathbf{t}} = M(\mathbf{x} \odot \mathbf{t})$, $\| \cdot \|_1$ and $\| \cdot \|_2$ are the L1 norm and L2 norm respectively. \mathbf{e} and $e_{\mathbf{t}}$ represent the embeddings obtained from the original sample \mathbf{x} and the masked sample $\mathbf{x} \odot \mathbf{t}$ after inputting them into the pre-training masked model M , respectively. In addition, for the reconstruction difficulty of the entire embedding $M_d(e_{\mathbf{t}}) - \mathbf{e}$, we only consider the reconstruction difficulty at the masked positions $\mathbf{1} - \hat{\mathbf{t}}$ when computing R , as it reflects the reconstruction difficulty of the missing information. Based on the embedding reconstruction difficulty R on a single sample, we can further calculate \mathcal{R} and $\Delta\mathcal{R}$ using Eq. (1) and Eq. (2).

3.3. The Proposed DOV4MM

We propose a method called DOV4MM to verify dataset ownership by relative embedding reconstruction difficulty, as shown in Fig. 2, which comprises three key steps:

(1) **Splitting Dataset & Training a Decoder:** Randomly split the public dataset \mathcal{D}_{pub} , the dataset that needs to be protected, into two subsets, referred to as the training dataset \mathcal{D}_t and the validation dataset \mathcal{D}_v . Then use the suspicious model M_s and \mathcal{D}_t to train a decoder M_d , whose training objective is to reconstruct the embedding information lost due to the mask \mathbf{t} ;

(2) **Computing Relative Difficulty:** Perform K samplings on \mathcal{D}_t , \mathcal{D}_v , and \mathcal{D}_{pvt} (a defender’s private dataset which isn’t publicly available, and M_s has not been pre-trained on it), with N samples per sampling. Based on these samples, use Eq. (1) and Eq. (3) to obtain the embedding reconstruction difficulties \mathcal{R}_t , \mathcal{R}_v , and \mathcal{R}_p for the suspicious model M_s on these three datasets. Then, according to Eq. (2), with \mathcal{D}_t as the standard, calculate $\Delta\mathcal{R}_{vt}$ using \mathcal{R}_t and \mathcal{R}_v , and calculate $\Delta\mathcal{R}_{pt}$ using \mathcal{R}_t and \mathcal{R}_p ;

(3) **Hypothesis Testing:** Conduct a one-sided pair-wise t-test [23] on $\Delta\mathcal{R}_{vt}$ and $\Delta\mathcal{R}_{pt}$. The null hypothesis, H_0 , posits that the mean difference between the paired samples in $\Delta\mathcal{R}_{pt}$ and $\Delta\mathcal{R}_{vt}$ is less than or equal to 0, while the alternative hypothesis, denoted as H_1 , posits that the mean difference between the paired samples in $\Delta\mathcal{R}_{pt}$ and $\Delta\mathcal{R}_{vt}$ is greater than 0.

If the p -value is less than 0.05, we can reject the null hypothesis and conclude that the mean difference between the paired samples in $\Delta\mathcal{R}_{pt}$ and $\Delta\mathcal{R}_{vt}$ is greater than 0, *i.e.*, the suspicious model M_s is very likely to have been pre-trained on the public dataset \mathcal{D}_{pub} . Therefore, we can infer that M_s is illegal and \mathcal{D}_{pub} has been stolen. On the other hand, if the null hypothesis can’t be rejected, we think that the mean difference between the paired samples in $\Delta\mathcal{R}_{pt}$

and $\Delta\mathcal{R}_{vt}$ is less than or equal to 0, which implies M_s is legal and the suspect is innocent.

4. Experiments

Here, we present the results of DOV4MM on ImageNet-1K [11] and WikiText-103 [40], ablation experiments, and interference resistance analysis. More details and results (hypothesis testing, visualizations, efficiency analysis, etc) can be found in the supplementary materials.

4.1. Experimental Setup

Datasets and models. We assess the performance of the proposed DOV4MM with ImageNet-1K [11] and its subsets, including ImageNet-50 and ImageNet-100. ImageNet-50 is a randomly selected subset of ImageNet-1K, encompassing 50 categories and 63,323 color images. In the same way, ImageNet-100 is also a randomly chosen subset of ImageNet-1K, consisting of 100 categories and 126,532 color images. We conducted the following two main experiments using these datasets:

Experiments on ImageNet-1K subsets. The pre-training dataset of the suspicious model M_s is a ImageNet-1K subset (ImageNet-50 or ImageNet-100). \mathcal{D}_{pub} is one of the following: the subset of ImageNet-1K (M_s ’s pre-training dataset), Food101 [2], COCO [32], or Places365 [56]. The architecture of M_s includes ViT-B/16 [13] and ViT-L/16 [13], and the masked modeling methods include MAE [22], CAE [8] and iBOT [57].

Experiments on ImageNet-1K. The pre-training dataset of the suspicious model M_s is ImageNet-1K. \mathcal{D}_{pub} is one of the following: ImageNet-1K (M_s ’s pre-training dataset), Food101 [2], COCO [32], or Places365 [56]. The architecture of the suspicious model M_s includes four architectures (ViT-B/16, ViT-L/16, Swin-B [38], Swin-L [38]) and ten masked modeling methods (MAE, CAE, iBOT, SimMIM [53], MaskFeat [50], PixMIM [37], BEiT [1], BEiT v2 [41], MixMAE [33] and EVA [15]). Note that the pre-trained models in this experiment are sourced from their official repositories or MMSelfSup¹.

When \mathcal{D}_{pub} , the defender’s dataset, is the same as the pre-training dataset of the suspicious model M_s , M_s is deemed illegal; otherwise, M_s is deemed legal. In all experiments, M_s is pre-trained using the default settings of each self-supervised method, with 400 epochs and a batch size of 64. The parameter settings for DOV4MM are fixed. Specifically, for convenience, we set \mathcal{D}_{pvt} as the testing set of \mathcal{D}_{pub} . \mathcal{D}_t consists of 20,000 randomly selected samples from \mathcal{D}_{pub} , and \mathcal{D}_v contains the remaining samples from \mathcal{D}_{pub} . The iterations of sampling $K = 30$, with $N = 1,024$ samples per iteration. The mask strategy for \mathbf{t} is random

¹https://mmselfsup.readthedocs.io/en/latest/model_zoo.html

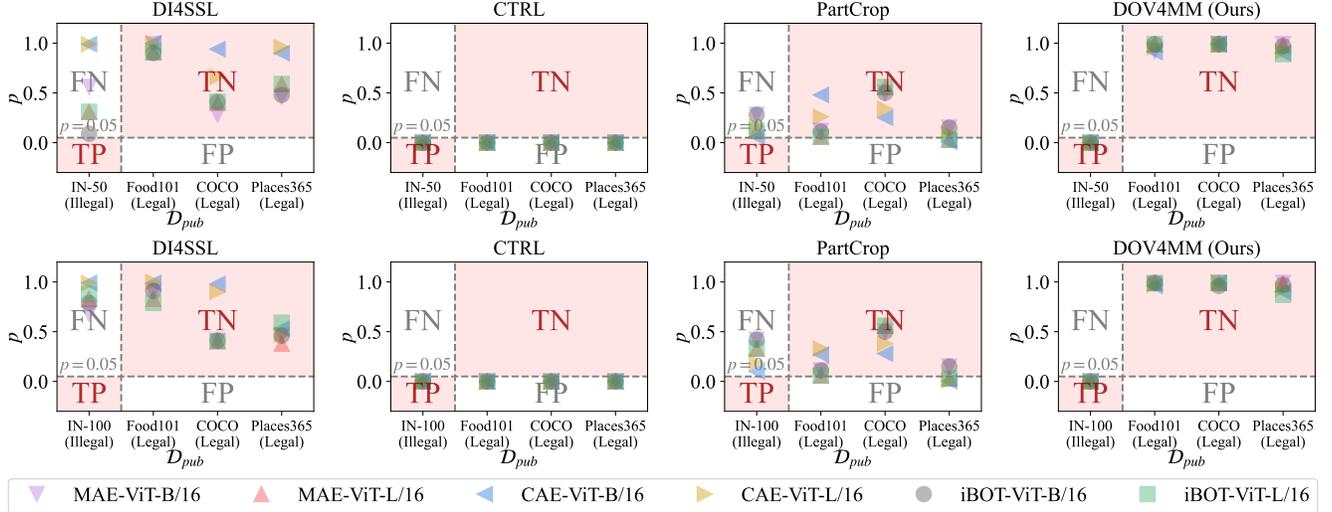


Figure 3. Results of four methods on ImageNet-50 (the first line) and ImageNet-100 (the second line). The pre-training dataset of the suspicious model is ImageNet-50 and ImageNet-100 in two cases respectively. Each pattern represents a suspicious model pre-trained using a specific architecture, masked modeling method, and dataset. “MAE-ViT-B/16” represents ViT-B/16 pre-trained using MAE, and the rest follows similarly. On the x-axis, “IN-50” and “IN-100” represent ImageNet-50 and ImageNet-100, respectively. Additionally, the terms “Illegal/Legal” in parentheses on the x-axis indicate the type of the suspicious model in each scenario. We consider illegal/legal models as positive/negative cases and classify each situation based on p -value.

masking, with a masking rate of 75%. The decoder is a Transformer [49] with an embedding dimension of 512, 8 layers, and 16 heads in the multi-head attention mechanism, which is trained for 50 epochs with a batch size of 64, a learning rate of $1e-3$, and the Adam optimizer [25]. More details are provided in the supplementary materials.

Evaluation metrics. We classify the suspicious model M_s as either illegal or legal based on the p -value. Specifically, when p is less than 0.05, we consider that M_s has been pre-trained on \mathcal{D}_{pub} and is illegal; when p is greater than 0.05, M_s is considered legal. Given that this is a classification task, except for the p -value, we also use the sensitivity, specificity and AUROC as the evaluation metrics. Sensitivity is the proportion of correctly predicted positive cases among all actual positive samples, and specificity is the proportion of correctly predicted negative cases among all actual negative samples. They reflect the ability to identify positive and negative samples, respectively.

4.2. Baselines

DI4SSL [14]: DI4SSL is the most recent method for dataset inference targeting self-supervised models. It can also be applied to dataset ownership verification directly.

CTRL [28]: CTRL is one of the state-of-the-art backdoor attacks targeting self-supervised models, which can be used for dataset ownership verification through backdoor watermarking. Specifically, we inject the CTRL trigger as watermark into a small subset of the public dataset. During the verification phase, if the representations of the water-

marked images are more similar to each other than those of the non-watermarked images, we can conclude that M_s was pre-trained using the public dataset.

PartCrop [59]: PartCrop is the latest member inference method for self-supervised models. We make slight modifications to make it applicable for DOV. Specifically, we crop certain parts of objects in both the training and testing images to query their similarity in the embedding space. If higher similarity is observed in the training images, we conclude M_s has been pre-trained on \mathcal{D}_{pub} .

4.3. Results on ImageNet

Results on ImageNet-1K subsets. Our approach is proven effective as illustrated in Fig. 3 (refer to supplementary material for specific p -values), which display the experimental results of baselines and DOV4MM on ImageNet-50 and ImageNet-100. Note that when \mathcal{D}_{pub} is ImageNet-50 or ImageNet-100, which implies the suspect is illegal (the pre-training dataset of the suspicious model M_s is ImageNet-50 and ImageNet-100 in two cases respectively), and p -values should be less than 0.05. However, when \mathcal{D}_{pub} is Food101, COCO or Places365, the suspect is deemed legal, so p -values should be greater than 0.05.

As shown in Tab. 1, we calculate sensitivity, specificity and AUROC based on the results in Fig. 3, which demonstrates the superiority of DOV4MM quantitatively. All baselines struggle to accurately distinguish the legality of various scenarios. For DI4SSL and PartCrop, they directly validate high-dimensional representations, while

Dataset	Method	Sensitivity	Specificity	AUROC
ImageNet-50	DI4SSL	0.00	1.00	0.50
	CTRL	1.00	0.00	0.50
	PartCrop	0.00	0.22	0.39
	DOV4MM	1.00	1.00	1.00
ImageNet-100	DI4SSL	0.00	1.00	0.50
	CTRL	1.00	0.00	0.50
	PartCrop	0.00	0.28	0.36
	DOV4MM	1.00	1.00	1.00

Table 1. Sensitivity, specificity, and AUROC of four methods on ImageNet-50 and ImageNet-100. **Bold** represents the best results.

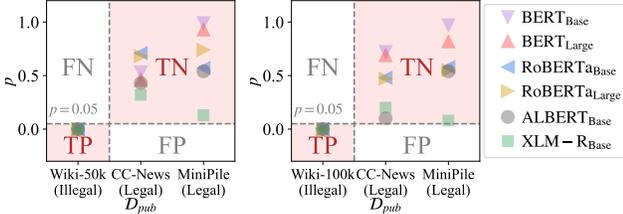


Figure 4. The results on WikiText-103 subsets. M_s is pre-trained on WikiText-103-50k (left) and WikiText-103-100k (right) respectively. On the x-axis, “Wiki-50k” and “Wiki-100k” represent WikiText-103-50k and WikiText-103-100k, respectively. The remaining identifiers are the same as those in Fig. 3.

DOV4MM validates the most valuable relative embedding reconstruction difficulty extracted from the representations. CTRL’s failure occurs because backdoor watermarks are usually fixed patterns (e.g., fixed-frequency noise for CTRL), designed to ensure the backdoor is successfully embedded into the model. Compared to benign images, the model generates more similar representations for watermarked images, even if the model is pre-trained on a non-watermarked dataset. In contrast, DOV4MM does not rely on backdoor watermarks, thus avoiding this risk.

Results on ImageNet-1K. We evaluated DOV4MM on ImageNet-1K using ten different masked image modeling (MIM) methods. Note that in this experiment, the pre-training dataset of M_s is ImageNet-1K, and their pre-trained weight come from the official repository or MM-SelfSup. The results of DOV4MM are shown in Tab. 2, which demonstrate that even if M_s is pre-trained on a large-scale dataset like ImageNet-1K, we can still correctly identify malicious behavior, even when using only 3% of ImageNet-1K data for dataset ownership verification.

4.4. Ablation Studies

The suspicious model M_s used in all ablation experiments is pre-trained on ImageNet-1K, with default settings marked in gray areas of Tab. 3. The model used in MAE, CAE, and iBOT is ViT-B/16, while the model used in SimMIM is Swin-B. Note that in this section, the public dataset \mathcal{D}_{pub} is ImageNet-1K, so p -values need to be below 0.05.

Decoder M_d . We experimented with decoders of vary-

Model	MIM Method	\mathcal{D}_{pub}			
		IN-1K	Food101	COCO	Places365
ViT-B/16	MAE	10^{-5}	0.99	0.98	0.99
	CAE	10^{-3}	0.99	0.99	0.97
	iBOT	10^{-3}	0.99	0.99	0.99
	MaskFeat	10^{-4}	0.99	0.99	0.99
	PixMIM	10^{-5}	0.99	0.99	0.99
	BEiT	10^{-3}	0.93	0.99	0.96
	BEiT v2	10^{-5}	0.99	0.99	0.99
	MixMAE	10^{-5}	0.95	0.94	0.99
	EVA	10^{-3}	0.92	0.98	0.72
	ViT-L/16	MAE	10^{-6}	0.99	0.99
CAE		10^{-4}	0.99	0.99	0.99
iBOT		10^{-3}	0.99	0.99	0.97
Swin-B	SimMIM	0.03	0.99	0.98	0.98
Swin-L	SimMIM	0.03	0.99	0.82	0.98

Table 2. The results (p -values) of DOV4MM on ImageNet-1K. “IN-1K” represents ImageNet-1K. Note that the illegal (p should be less than 0.05) and legal (p should be greater than 0.05) scenarios correspond to the red and blue areas.

ing widths (number of channels), depths (number of Transformer blocks), and the number of attention heads, as shown in Tab. 3a, Tab. 3b and Tab. 3c. The results indicate that the scale of the M_d does not have a significant impact on DOV4MM. Surprisingly, when the M_d is only 3% of the size of M_s (ViT-B/16), DOV4MM still remains effective.

The size of the training dataset \mathcal{D}_t for M_d . We randomly selected different sizes of \mathcal{D}_t from \mathcal{D}_{pub} to train M_d , and the results are shown in Tab. 3d. It indicates that the larger \mathcal{D}_t is, the better the performance of DOV4MM. This is because a very small \mathcal{D}_t cannot effectively teach M_d how to reconstruct embeddings. Experiments show that even when we train M_d with only 10,000 images (0.8% of the ImageNet-1K), we can achieve correct results.

Masking ratio. We evaluate DOV4MM at different masking ratios, and the results are shown in Tab. 3e. DOV4MM is less sensitive to the ratios, and works well across a wide range of masking ratios (30-90%).

K and N . Tab. 3f and Tab. 3g show the performance of DOV4MM under different sampling iterations (K) and the number of samples per iteration (N), which indicates that as K and N increase, the performance of DOV4MM improves. This is because as K and N increase, more data is used in the t-test, making the results more precise. Excitingly, even with just 40,000 images for testing (3% of the ImageNet-1K), we can achieve satisfactory results.

T-test’s object. We use $\Delta\mathcal{R}$ or \mathcal{R} as the object for t-test. As shown in Tab. 3h, using only \mathcal{R} does not yield correct results (p -value exceeds 0.05).

4.5. Results on Masked Language Modeling

Masked modeling methods originated in natural language processing (NLP), and we selected four masked lan-

Dim	MAE	CAE	iBOT	Blocks	MAE	CAE	iBOT	Heads	MAE	CAE	iBOT	$ \mathcal{D}_t $	MAE	CAE	iBOT
128	10^{-5}	10^{-3}	10^{-3}	4	10^{-5}	10^{-3}	10^{-3}	4	10^{-6}	10^{-3}	10^{-3}	10,000	10^{-4}	0.02	0.01
256	10^{-6}	10^{-3}	10^{-3}	6	10^{-6}	10^{-3}	10^{-3}	8	10^{-6}	0.01	10^{-3}	20,000	10^{-5}	10^{-3}	10^{-3}
512	10^{-5}	10^{-3}	10^{-3}	8	10^{-5}	10^{-3}	10^{-3}	16	10^{-5}	10^{-3}	10^{-3}	30,000	10^{-6}	10^{-3}	10^{-3}
768	10^{-6}	0.01	10^{-3}	10	10^{-7}	10^{-3}	10^{-3}	32	10^{-6}	0.01	10^{-3}	40,000	10^{-6}	0.01	10^{-3}
1,024	10^{-7}	0.01	10^{-3}	12	10^{-6}	0.01	10^{-3}	64	10^{-6}	10^{-3}	10^{-3}	50,000	10^{-6}	10^{-3}	10^{-3}

Ratio (%)	MAE	CAE	iBOT	K	MAE	CAE	iBOT	N	MAE	CAE	iBOT	T-test's Object	SimMIM
30	10^{-6}	10^{-3}	10^{-3}	10	10^{-3}	0.05	0.08	256	0.02	0.09	0.08	$\Delta\mathcal{R}$	0.03
45	10^{-6}	10^{-3}	10^{-3}	20	10^{-7}	0.02	10^{-3}	512	10^{-4}	0.10	10^{-3}	\mathcal{R}	0.10
60	10^{-6}	10^{-3}	10^{-3}	30	10^{-5}	10^{-3}	10^{-3}	1,024	10^{-5}	10^{-3}	10^{-3}		
75	10^{-5}	10^{-3}	10^{-3}	40	10^{-8}	10^{-3}	10^{-4}	2,048	10^{-10}	10^{-4}	10^{-6}		
90	10^{-6}	0.04	10^{-3}	50	10^{-5}	10^{-3}	10^{-4}	4,096	10^{-18}	10^{-9}	10^{-10}		

Table 3. DOV4MM ablation experiments on ImageNet-1K. We report p -values. The default settings is: the decoder M_d has depth 8, width 512 and attention heads 16, the size of the training dataset for M_d is 20,000, and the masking ratio is 75%. Sampling iterations and the number of samples per iteration are 30 and 1,024, respectively. The object of t-test is $\Delta\mathcal{R}$. Default settings are marked in gray.

MIM Method	Model	w/o es	w/ es (patience=15)
MAE	ViT-B/16	10^{-5}	10^{-5}
	ViT-L/16	10^{-6}	10^{-4}

Table 4. We report the p -values of M_s , whose pre-training dataset is ImageNet-100, with or without early stopping (es).

guage modeling methods, BERT [12], RoBERTa [34], ALBERT [27] and XLM-R [10], to evaluate DOV4MM. The training dataset of the suspicious model M_s is subsets of WikiText-103 [40], which contains text from English-language wikipedia articles. We denote the subsets as WikiText-103-50k and WikiText-103-100k, which contain 50,000 and 100,000 random samples from WikiText-103, respectively. The training hyperparameters of M_s are provided in the supplementary materials. Note that M_s uses “[MASK]” as the mask token for BERT/ALBERT pre-training, and “<mask>” for RoBERTa/XLM-R. Since the defender is unaware of M_s ’s mask token, we default to using “[MASK]” as the mask token in DOV4MM. Additionally, apart from a 20% masking ratio, other parameters are identical to those of visual DOV4MM.

The defender’s public dataset \mathcal{D}_{pub} is one of the following: the subset of WikiText-103 (the pre-training dataset of M_s), CC-News [19], and MiniPile [24]. The results are shown in Fig. 4, where DOV4MM can work in various cases, indicating that it can effectively scale to NLP.

4.6. The Interference Resistance of DOV4MM

In this section, we study that whether DOV4MM remains effective when facing more covert data theft. More experimental details can be found in the supplementary material.

Early stopping. The suspect may use early stopping to prevent overfitting. Therefore, we tested the performance of DOV4MM under early stopping. \mathcal{D}_{pub} is ImageNet-100, and the p -value needs to be less than 0.05. As shown in Tab. 4, DOV4MM remains effective in this case.

\mathcal{D}_d	MIM Method	Model	w/o ft	w/ ft	Acc (%)
Food101	MAE	ViT-B/16	10^{-5}	10^{-4}	60.65
		ViT-L/16	10^{-5}	10^{-5}	67.79
	iBOT	ViT-B/16	10^{-3}	10^{-3}	80.28
		ViT-L/16	10^{-3}	10^{-3}	81.99
Places365	MAE	ViT-B/16	10^{-5}	10^{-5}	42.47
		ViT-L/16	10^{-5}	10^{-5}	45.98
	iBOT	ViT-B/16	10^{-3}	10^{-3}	48.08
		ViT-L/16	10^{-3}	0.01	48.85

Table 5. We report the p -values of M_s , whose pre-training dataset is ImageNet-1K, with or without fine-tuning (ft) on downstream task. \mathcal{D}_d is the downstream dataset, and “Acc” is the classification accuracy of the fine-tuned M_s on the downstream task.

Fine-tuning M_s on downstream task. Given M_s pre-trained on ImageNet-1K, we fine-tune it using on Food101 and Places365 for image classification, then DOV4MM performs the test based on the logits output from M_s . Note that in this case, the output of the fine-tuned M_s is logits, not tokens. Therefore, when calculating the reconstruction difficulty, we no longer use the embedding mask \hat{t} , but instead calculate the reconstruction difficulty for the entire logits. Moreover, the decoder is no longer a Transformer but instead a simple three-layer fully connected network. \mathcal{D}_{pub} is ImageNet-1K, and the p -value should be less than 0.05. As shown in Tab. 5, DOV4MM can still work even when defender has no access to the feature vectors.

5. Conclusion

In this work, we propose DOV4MM for verifying dataset ownership in masked modeling. Specifically, we propose the concept of relative embedding reconstruction difficulty based on reconstruction characteristics of masked embedding. The experiment demonstrates its effectiveness. Future work includes (1) extending our method to other self-supervised approaches; (2) protecting other types of data; (3) exploring other privacy risks associated with encoders.

6. Acknowledgment

This work was partially supported by the Pioneer R&D Program of Zhejiang (No.2024C01021), and Zhejiang Province High-Level Talents Special Support Program “Leading Talent of Technological Innovation of Ten-Thousands Talents Program” (No. 2022R52046).

References

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 2, 5
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, pages 446–461. Springer, 2014. 5, 1
- [3] Nicholas Carlini and Andreas Terzis. Poisoning and backdooring contrastive learning. *arXiv preprint arXiv:2106.09667*, 2021. 2
- [4] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1897–1914. IEEE, 2022. 3
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2, 3
- [6] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021.
- [7] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9640–9649, 2021. 2, 3
- [8] Xiaokang Chen, Mingyu Ding, Xiaodi Wang, Ying Xin, Shentong Mo, Yunhao Wang, Shumin Han, Ping Luo, Gang Zeng, and Jingdong Wang. Context autoencoder for self-supervised representation learning. *International Journal of Computer Vision*, 132(1):208–223, 2024. 5
- [9] Christopher A Choquette-Choo, Florian Tramèr, Nicholas Carlini, and Nicolas Papernot. Label-only membership inference attacks. In *International conference on machine learning*, pages 1964–1974. PMLR, 2021. 3
- [10] A Conneau. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*, 2019. 3, 8
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1, 3, 5
- [12] Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2, 3, 8
- [13] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 5
- [14] Adam Dziedzić, Haonan Duan, Muhammad Ahmad Kaleem, Nikita Dhawan, Jonas Guan, Yannis Cattan, Franziska Boenisch, and Nicolas Papernot. Dataset inference for self-supervised models. *Advances in Neural Information Processing Systems*, 35:12058–12070, 2022. 3, 6
- [15] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023. 5
- [16] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The Pile: An 800GB dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020. 1
- [17] Junfeng Guo, Yiming Li, Lixu Wang, Shu-Tao Xia, Heng Huang, Cong Liu, and Bo Li. Domain watermark: Effective and harmless dataset copyright protection is closed at hand. *Advances in Neural Information Processing Systems*, 36, 2023. 1, 2
- [18] Junfeng Guo, Yiming Li, Ruibo Chen, Yihan Wu, Chenxi Liu, and Heng Huang. Zeromark: Towards dataset ownership verification without disclosing watermark. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 1, 2
- [19] Felix Hamborg, Norman Meuschke, Corinna Breiteringer, and Bela Gipp. news-please: A generic news crawler and extractor. In *Proceedings of the 15th International Symposium of Information Science*, pages 218–223, 2017. 8, 1
- [20] 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. 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 2, 3, 5
- [23] Robert V Hogg, Joseph W McKean, Allen T Craig, et al. *Introduction to mathematical statistics*. Pearson Education India, 2013. 2, 5
- [24] Jean Kaddour. The minipile challenge for data-efficient language models. *arXiv preprint arXiv:2304.08442*, 2023. 8, 1
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [26] A Krizhevsky. Learning multiple layers of features from tiny images. *Master’s thesis, University of Tront*, 2009. 1

- [27] Zhenzhong Lan. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019. 2, 8
- [28] 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. 2, 6
- [29] Siyuan Li, Luyuan Zhang, Zedong Wang, Di Wu, Lirong Wu, Zicheng Liu, Jun Xia, Cheng Tan, Yang Liu, Baigui Sun, et al. Masked modeling for self-supervised representation learning on vision and beyond. *arXiv preprint arXiv:2401.00897*, 2023. 3
- [30] Yiming Li, Yang Bai, Yong Jiang, Yong Yang, Shu-Tao Xia, and Bo Li. Untargeted backdoor watermark: Towards harmless and stealthy dataset copyright protection. *Advances in Neural Information Processing Systems*, 35:13238–13250, 2022. 1, 2
- [31] Yiming Li, Mingyan Zhu, Xue Yang, Yong Jiang, Tao Wei, and Shu-Tao Xia. Black-box dataset ownership verification via backdoor watermarking. *IEEE Transactions on Information Forensics and Security*, 2023. 1, 2
- [32] 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. 5, 1
- [33] Jihao Liu, Xin Huang, Jinliang Zheng, Yu Liu, and Hongsheng Li. Mixmae: Mixed and masked autoencoder for efficient pretraining of hierarchical vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6252–6261, 2023. 5
- [34] Yinhan Liu. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 364, 2019. 2, 3, 8
- [35] Yang Liu, Zhen Zhu, and Xiang Bai. Wdnet: Watermark-decomposition network for visible watermark removal. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3685–3693, 2021. 2
- [36] Yupei Liu, Jinyuan Jia, Hongbin Liu, and Neil Zhenqiang Gong. Stolenencoder: stealing pre-trained encoders in self-supervised learning. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 2115–2128, 2022. 2
- [37] Yuan Liu, Songyang Zhang, Jiacheng Chen, Kai Chen, and Dahua Lin. Pixmim: Rethinking pixel reconstruction in masked image modeling. *arXiv preprint arXiv:2303.02416*, 2023. 5
- [38] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 3, 5
- [39] Pratyush Maini, Mohammad Yaghini, and Nicolas Papernot. Dataset inference: Ownership resolution in machine learning. *arXiv preprint arXiv:2104.10706*, 2021. 2
- [40] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016. 5, 8, 1
- [41] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. Beit v2: Masked image modeling with vector-quantized visual tokenizers. *arXiv preprint arXiv:2208.06366*, 2022. 5
- [42] Aniruddha Saha, Ajinkya Tejankar, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Backdoor attacks on self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13337–13346, 2022. 2
- [43] Zeyang Sha, Xinlei He, Ning Yu, Michael Backes, and Yang Zhang. Can’t steal? contrastive stealing attacks against image encoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16373–16383, 2023. 2, 3
- [44] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017. 3
- [45] Karen Simonyan. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [46] Xiaofei Sun, Xiaoya Li, Yuxian Meng, Xiang Ao, Lingjuan Lyu, Jiwei Li, and Tianwei Zhang. Defending against backdoor attacks in natural language generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5257–5265, 2023. 2
- [47] Ruixiang Tang, Qizhang Feng, Ninghao Liu, Fan Yang, and Xia Hu. Did you train on my dataset? towards public dataset protection with cleanlabel backdoor watermarking. *ACM SIGKDD Explorations Newsletter*, 25(1):43–53, 2023. 2
- [48] Ajinkya Tejankar, Maziar Sanjabi, Qifan Wang, Sinong Wang, Hamed Firooz, Hamed Pirsiavash, and Liang Tan. Defending against patch-based backdoor attacks on self-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12239–12249, 2023. 2
- [49] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 6, 1
- [50] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14668–14678, 2022. 3, 5
- [51] Yuechen Xie, Jie Song, Huiqiong Wang, and Mingli Song. Training data provenance verification: Did your model use synthetic data from my generative model for training? In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 23817–23827, 2025. 1
- [52] Yuechen Xie, Jie Song, Mengqi Xue, Haofei Zhang, Xingen Wang, Bingde Hu, Genlang Chen, and Mingli Song. Dataset ownership verification in contrastive pre-trained models, 2025. 2
- [53] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of*

- the IEEE/CVF conference on computer vision and pattern recognition*, pages 9653–9663, 2022. [5](#)
- [54] Jinghuai Zhang, Hongbin Liu, Jinyuan Jia, and Neil Zhenqiang Gong. Data poisoning based backdoor attacks to contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24357–24366, 2024. [2](#)
- [55] Mengxin Zheng, Jiaqi Xue, Zihao Wang, Xun Chen, Qian Lou, Lei Jiang, and Xiaofeng Wang. Ssl-cleanse: Trojan detection and mitigation in self-supervised learning. In *European Conference on Computer Vision*, pages 405–421. Springer, 2024. [2](#)
- [56] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017. [5](#), [1](#)
- [57] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. [2](#), [5](#)
- [58] Qiang Zhou, Chaohui Yu, Hao Luo, Zhibin Wang, and Hao Li. Mimco: Masked image modeling pre-training with contrastive teacher. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4487–4495, 2022. [2](#)
- [59] Jie Zhu, Jirong Zha, Ding Li, and Leye Wang. A unified membership inference method for visual self-supervised encoder via part-aware capability. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 1241–1255, 2024. [3](#), [6](#)