

Anonymity Unveiled: A Practical Framework for Auditing Data Use in Deep Learning Models

Zitao Chen
University of British Columbia
Vancouver, BC, Canada
zitaoc@ece.ubc.ca

Karthik Pattabiraman
University of British Columbia
Vancouver, BC, Canada
karthikp@ece.ubc.ca

ABSTRACT

The rise of deep learning (DL) has led to a surging demand for training data, which incentivizes the creators of DL models to trawl through the Internet for training materials. Meanwhile, users often have limited control over whether their data (e.g., facial images) are used to train DL models without their consent, which has engendered pressing concerns.

This work proposes *MembershipTracker*, a practical data auditing tool that can empower ordinary users to reliably detect the unauthorized use of their data in training DL models. We view data auditing through the lens of membership inference (MI). *MembershipTracker* consists of a lightweight data marking component to mark the target data with small and targeted changes, which can be strongly memorized by the model trained on them; and a specialized MI-based verification process to audit whether the model exhibits strong memorization on the target samples.

MembershipTracker only requires the users to mark a small fraction of data (0.005%~0.1% in proportion to the training set), and it enables the users to reliably detect the unauthorized use of their data (average 100% TPR@0% FPR). We show that *MembershipTracker* is highly effective across various settings, including industry-scale training on the full-size ImageNet-1k dataset. We finally evaluate *MembershipTracker* under multiple classes of countermeasures¹.

CCS CONCEPTS

• Security and privacy; • Computing methodologies → Machine learning;

KEYWORDS

Data auditing, data provenance, membership inference

ACM Reference Format:

Zitao Chen and Karthik Pattabiraman. 2025. Anonymity Unveiled: A Practical Framework for Auditing Data Use in Deep Learning Models. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25)*, October 13–17, 2025, Taipei, Taiwan. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3719027.3744794>

¹Our code is available at <https://github.com/DependableSystemsLab/MembershipTracker>.

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.

CCS '25, October 13–17, 2025, Taipei, Taiwan

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1525-9/2025/10
<https://doi.org/10.1145/3719027.3744794>

1 INTRODUCTION

Modern deep learning (DL) models have attained remarkable performance in various tasks such as image classification and language generation. Their success is largely driven by the availability of massive training data [49, 51]. However, it is not always clear whether the data used to train these models has been used with the permission of the data owners. For instance, Clearview.ai, in developing their facial recognition system, scraped billions of user photos from social media platform without the users' consent [25]. The unregulated use of personal data for building DL models has engendered mounting concerns [19, 25, 64], and is in violation of privacy laws such as the European Union's General Data Protection Regulation (GDPR) [3]. In spite of this, data holders often have limited agency in detecting the unauthorized use of their data due to the lack of practical data provenance tools.

This leads to two main lines of work for data provenance. The first line of work seek to inject designated features into the user data, which can induce the model to exhibit a certain detectable behavior for provenance purposes [26, 39, 52, 67]. The injected features can take the form of a backdoor trigger pattern [26, 39]; or a spurious feature pattern that can cause output probability shift by the model [67]. However, these techniques either impose strong assumptions on the users (e.g., the ability to mark a large fraction of data) [26, 39, 52] and/or suffer from limited auditing performance [39, 52, 67].

The second line of work is based on *membership inference* (MI). There are instance- and user-level MI - the former infers whether a *specific* sample is used to train a model [8, 55]; while the latter detects the usage of *any* of the user's data [11, 31, 57]. Our work focuses on auditing whether *specific user samples* are used for training the target model. Unfortunately, existing MI methods often suffer from limited effectiveness, as measured by their true positive rate (TPR) and false positive rate (FPR) [8]. However, a high-power MI method is needed to facilitate data provenance with low FPR (to avoid false accusation) and high TPR (for correct data provenance).

To improve, recent studies propose several techniques based on data poisoning [9, 12, 63]. Their idea is to inject mislabeled samples into the model's training set, which can amplify the model's behavioral changes on some target samples, thereby making them easier to be de-identified. To apply these techniques for data auditing, users have to be able to mislabel some training data, which, unfortunately, excludes the real-world scenarios where the data labels are assigned by the model owner or an external service [18, 20]. This renders these techniques challenging for real-world use.

This work. We present *MembershipTracker*, a MI-based data provenance tool that can overcome the aforementioned challenges, with a two-step process as illustrated in Fig. 1.

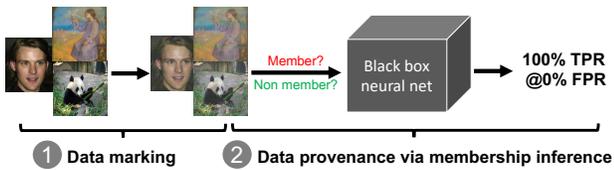


Figure 1: MembershipTracker is a data provenance tool that operates by: (1) marking the target data (e.g., facial images, artworks) with small and targeted changes, and then (2) initiating a specialized membership inference process to audit whether the target data are used for training the model.

MembershipTracker operates under the **realistic setting** where we assume the users have neither the capability nor expertise to train shadow/proxy models [6, 8, 69], or to manipulate the data labels [63], as these are well beyond the capacity of ordinary personal users (Section 3.1.1). Rather, *MembershipTracker* only requires the users to mark a small fraction of data (amounting to 0.005%~0.1% of the training set) for provenance purposes. Our goal is to empower the users to carry out high-power MI on their target data with high TPR and low FPR, given *black-box* access to the model.

Technical design. We first observe that the effectiveness of MI on a given model is directly related to the model’s ability in memorizing individual data points, and hence samples that are strongly memorized by the model (e.g., outlier samples) are easier to be de-identified [8, 13, 63, 70]. Therefore, in order to fulfill our goal, there are two challenges as follows.

- (1) How can the users *amplify* the model’s memorization on the target data with minimal data modification?
- (2) How can the users reliably *audit* whether the model exhibits strong memorization on their target data?

For the first challenge, we propose a lightweight data marking technique to mark the target data with small and targeted changes, which are formulated as the combination of outlier feature and procedural noise [16]. These subtle changes can be readily created by the users without assuming access to a proxy model, and can still induce the model to strongly memorize the specially-marked samples while preserving their visual information (e.g., Fig. 1).

For the second challenge, given a target user’s small set of marked samples, we propose a novel *set*-based MI process to audit whether they are used to train the model. Our approach follows a common idea in existing user-level MI studies [31, 46, 57] to leverage the *collective* information across the small set of target samples for MI. But unlike prior methods, *MembershipTracker* does not require any additional shadow/reference models [31, 46, 57] and can enable high-power MI for reliable data provenance.

Contributions. We make three contributions as follows.

- Propose a lightweight data marking technique to mark the users’ target data with small and targeted changes, which can be strongly memorized by the model trained on them.
- Develop a high-power set-based MI verification process that can reliably audit whether the target data are used to train the model (with high TPR when controlled under a low FPR).
- Integrate the above as a tool called *MembershipTracker*, and evaluate its effectiveness across a variety of settings (six benchmark

datasets and six DL architectures). We also comprehensively evaluate a total of six classes of countermeasures.

We find that by merely marking a small fraction of samples (0.005%~0.1% of the dataset), *MembershipTracker* effectively empowers the users to reliably audit the usage of their data (average 100% TPR@0% FPR), and it is scalable to the large-scale ImageNet training (on both Convolutional and Transformer network). *To the best of our knowledge, MembershipTracker is the first technique that can scale to support reliable data auditing in large-scale training setting.* This renders *MembershipTracker* a practical data auditing tool that contributes to responsible AI practice.

2 RELATED WORK

Our work focuses on detecting the unauthorized use of personal data in training DL models, which is complementary to existing work that aim to protect the data from unwanted use by rendering them un-learnable [29, 54]. Hence we focus on existing literature for data provenance, and we classify them into two categories.

Non membership-inference based solutions. We divide them into dataset- and user-level solutions.

Dataset level provenance aims to detect whether a DL model is trained on a specific dataset [39, 45, 52]. Maini et al. propose dataset inference, which detects whether two models trained on the same dataset exhibit similarities in their decision boundaries [45]. Other work modify portions of training set to leave a certain artifact on the models, such as radioactive data [52] and backdoor watermark [39]. However, recent work [30] finds that these techniques [39, 52] still have poor auditing performance. Moreover, dataset level solutions require several assumptions that are outside the scope of our work. First, many of them require control over a nontrivial portion of training set (e.g., 10%~20% [39, 52]). Some solutions assume access to some known feature extractor [52], access to the training set or the ability to train a proxy model [45, 52]. These are well beyond the user capability we consider (in Section 3.1.1).

User level provenance seeks to detect whether the users’ data are used to train a DL model [26, 67], and they require only modifying the users’ own data. Their idea is to mark the users’ data with a designated feature to induce a detectable behavior into the model. Hu et al. propose to inject backdoor trigger feature into the user data [26], but they assume the user capability to manipulate the data labels. Other studies on clean-label backdoor attacks may be repurposed to overcome this, but they still require training a proxy model to compute the poisoned data [58, 72].

Wenger et al. propose to inject a target spurious feature into the user data to be learned by the suspected model, and then detect whether the model has associated the injected feature with the target class label [67]. To detect, the users can separately overlay the target and non-target features to some auxiliary data (outside the target class), and then determine whether the target feature causes a slightly higher prediction probability on the target class, compared with the non-target features (e.g., 10% vs. 1%). Both their work and ours involve adding spurious features to mark users’ data, but there are several major differences between the two. Notably, our analysis reveals that they [67] significantly *underestimate* their detection false positives (from 0% to >30%). This is because they employ a *discrepant* detection procedure for the testing features that

are used for training (true positives), and those that are not (false positives); however, a fair evaluation should follow a consistent procedure (see our detailed analysis in Appendix C).

Membership-inference (MI) based solutions. MI also represents a natural fit for tracing data provenance, and there are instance- and user-level MI solutions.

Instance-level MI detects whether a specific instance is used to train a model. It was first proposed by Shokri et al. [55] and there are many follow-up studies [6, 8, 15, 69, 70]. However, even state-of-the-art MI methods [6, 8, 69] still struggle with achieving high TPR when controlled at the low FPR regime (more in Section 4).

This has led to several attempts at improving the effectiveness of existing MI methods via data poisoning. Tramer et al. propose to inject mislabeled samples into the model’s training set, which can transform the target samples into outliers and amplify their influence on the model’s decision, thereby making the target samples easier to be de-identified [63]. While injecting mislabeled samples represents an effective solution for improving the MI success [9, 12, 63], it may be challenging for the users to apply to their data, because in many scenarios the users may not have control over the data labels [18, 20, 67]. Even if this is feasible, we find that such a method can still only increase the MI success to a limited extent (validated in Appendix D). By contrast, *MembershipTracker* does not assume control of the data labels, and is still able to facilitate high-power MI for tracing data provenance.

User-level MI aims to detect whether any of a user’s data is used to train a model. Unlike instance-level MI, the auditor possesses a set of samples from a target user that are *not* necessarily used to train the target model. Song et al. [57] propose the first user-level MI for natural language models, which first trains multiple shadow models, and then uses the outputs from shadow models as the features to train an audit model. Subsequent work develops solutions for other settings, including speech recognition models [10, 47], metric learning models [11, 37] and large language models [31, 46].

A common thread of existing user-level MI methods is to aggregate the information across multiple samples (e.g., different voices of the same speaker) to perform MI [10, 11, 31, 57]. The MI process in *MembershipTracker* follows a similar philosophy. However, prior work often requires access to additional reference models [31] or shadow models [10, 11, 57], and they still have limited MI effectiveness (e.g., limited TPR under the low FPR regime [31, 46]). *MembershipTracker* overcomes these limitations by means of a novel data-marking and specialized MI auditing process, which can simultaneously achieve high TPR and low FPR for data provenance.

Recent work by Huang et al. [30] proposes a general framework for auditing unauthorized data use in training DL models. Their idea is to generate two perturbed copies of the target data, randomly publish one of them, and then compare the model’s membership score on the published vs. the unpublished one. While their framework can work for different domains such as foundation models, they still require the data holders to mark a substantial portion (1%~10%) of dataset. Instead, *MembershipTracker* considers the more challenging (and realistic) setting where the marked data amounts to $\leq 0.1\%$ of the dataset. We also “stress test” their technique under a similar setting as ours and find that their performance degrades considerably then (see Appendix E).

3 PROBLEM FORMULATION

Membership Inference (MI) Game. We denote $F_\theta : \mathcal{X} \rightarrow [0, 1]^n$ as a model that maps an input sample $x \in \mathcal{X}$ to a probability vector over n classes. $D = \{(x_i, y_i)\}_{i=1}^n$ is a training set sampled from some distribution \mathcal{D} . $F_\theta \leftarrow \mathcal{T}(D)$ denotes a model F_θ produced from running the training algorithm \mathcal{T} on D .

Next, we define a MI game, which proceeds between a challenger C and an adversary \mathcal{A} . \mathcal{A} has to guess which element from some universe \mathcal{U} was used to train a model. For a specific target point x , $\mathcal{U} = \{x, \perp\}$, where \perp indicates the absence of an example.

- (1) The challenger samples a dataset $D \leftarrow \mathcal{D}$, and a target point $z \leftarrow \mathcal{U}$ (such that $D \cap \mathcal{U} = \emptyset$).
- (2) The challenger trains a model $F_\theta \leftarrow \mathcal{T}(D \cup \{z\})$ on the dataset D and target point z .
- (3) The challenger gives the adversary query access to F_θ .
- (4) The adversary emits a guess $\hat{z} \in \mathcal{U}$.
- (5) The adversary wins the game if $\hat{z} = z$.

3.1 MI for Data Provenance

We now describe the MI game for data provenance, which proceeds between a challenger (model creator) C and a target user U . We consider image classification in this work, and U seeks to audit whether his/her images (denoted as X) are used by C to train a model F without permission. We add the ability for U to modify X for provenance purpose.

- (1) The user is allowed to modify $X \in \mathcal{U}$ ($\mathcal{U} = \{X, \perp\}$).
- (2) The challenger samples a dataset $D \leftarrow \mathcal{D}$, and a target set $Z \leftarrow \mathcal{U}$ (such that $D \cap \mathcal{U} = \emptyset$).
- (3) The challenger trains a model $F_\theta \leftarrow \mathcal{T}(D \cup \{Z\})$ on the dataset D and target set Z .
- (4) The challenger gives the user query access to F_θ .
- (5) The user emits a guess $\hat{Z} \in \mathcal{U}$.
- (6) The user wins the game if $\hat{Z} = Z$.

For a given user U , $\mathcal{U} = \{X, \perp\}$, where \perp indicates the absence of the target samples by U , and the goal is to determine whether F is trained on D or $D \cup \{X\}$.

Unless otherwise stated, we follow the common practice to measure the MI success by its true positive rate (TPR) controlled under a low false positive rate (FPR) [8, 69].

3.1.1 User capability. We consider ordinary personal users U with limited ML expertise and resources, and they are willing to add visual distortion to their data for provenance purposes. We next specify their capabilities and constraints.

First, we assume U can only mark a small number of samples, and these are to be included as part of a larger dataset. This simulates the realistic scenario where the dataset is curated from multiple data sources (e.g., [18]), and thus data contributed by each user constitutes only a small portion of the dataset (e.g., 0.1%).

Next, in the provenance verification process, we assume U has access to a set of non-member samples that are not used for training the target model², which is a common assumption³ to control the

²Under the presumption that only a small fraction of data sampled from the distribution were used in training, one may simply draw a random sample from the underlying distribution, and be confident that it is representative of the non-member data [6, 69].

³While we also adopt this common assumption in our study, we conduct additional experiments and find that this assumption can in fact, be relaxed (see Section 6.2).

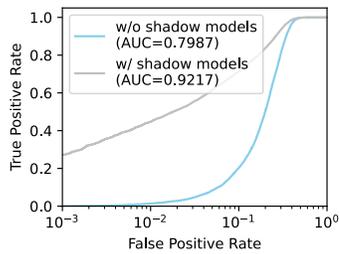


Figure 2: State-of-the-art MI methods [8] achieve limited TPR under the low FPR regime (undesirable).

MI process at a low FPR regime. U also has black-box access to query the target model. In a real world setting, users often have limited control and knowledge beyond their own data, and thus we specify the constraints by U as follows.

- U cannot modify their data’s labels, which are assigned by the model owner or an external service [18, 20, 67].
- U has no access to D (i.e., the training data by other users).
- U has no knowledge of the black-box model F .
- U has no expertise/resources to train any shadow/proxy models, as in many settings, training even a single shadow model can impose prohibitively high data and compute requirements (especially for the large models) [63].

4 METHODOLOGY

We first conduct an empirical study to understand the limitations of leveraging existing MI methods for data auditing, and then describe our design goals. We then present *MembershipTracker* with a data marking and provenance verification process: first by marking the users’ protected data with targeted changes (Section 4.1); then performing a specialized MI process to audit whether the marked data are used to train the model (Section 4.2).

Motivation. Our work addresses data provenance through MI, and there are several available MI methods [6, 8, 63, 69]. However, these methods still suffer from limited MI effectiveness in the low FPR regime and/or impose strong assumptions on the users - we elaborate both limitations next.

① Existing MI methods achieve limited TPR under a low FPR.

To validate this, we use two methods in the state-of-the-art Likelihood-ratio attack, one is with and the other without shadow model calibration (128 shadow models) [8]. We evaluate both methods on a WideResNet-28-4 model trained on half of the training set in CIFAR100 dataset. As shown, even the shadow-model based method can only achieve 27.21% TPR@0.1% FPR.

This is undesirable because achieving high TPR (under a low FPR) is crucial for the users to reliably detect the unauthorized use of their data. As in Fig. 2, existing MI methods struggle with this task. The reason is that the training samples are not always strongly memorized by the model (except some outliers) [8, 63, 70], which is directly related to the ability of performing accurate MI. This problem motivates related work [9, 12, 63] to amplify the model’s memorization on the training samples, and we discuss their limitation next.

② State-of-the-art MI methods either require training shadow models or assume control to manipulate the data labels, both are beyond ordinary users’ capabilities. For example, leading MI methods [6, 8, 69] commonly require training shadow models to calibrate the inference process, which assumes additional data and compute access, and is challenging for the non-expert users.

There are other methods that do not require shadow-model calibration, but they either suffer from poor effectiveness [8, 70] or impose unrealistic assumption on the users [9, 63]. E.g., the approach by Tramer et al. [63] can improve the MI success without shadow-model calibration, but it assumes users have the ability to mislabel some training data, which can be challenging in the cases where the users have no control over the data labels [18, 20, 67]. Even if this is feasible, we find that the increase of MI success is still largely limited (see Appendix D for a validation).

Design Goals. There are three design goals in our work.

Goal (1): *High provenance effectiveness.* Our main goal is to enable those users whose data are misused for model training to perform successful data provenance (high TPR), and we also want to ensure low FPR to avoid false accusation. Therefore, our goal is to achieve high TPR@low FPR.

Goal (2): *Preserving high visual quality.* The modification created to mark the users’ data should not severely distort their visual information, e.g., the identity in a facial image should still be easily recognizable.

Goal (3): *Minimal technical requirements.* The provenance tool should be usable by ordinary users even with limited expertise and resources (specified in Section 3.1.1).

4.1 Data Marking

The goal of data marking is to ensure that the membership of the marked samples can be reliably de-identified. To this end, we first observe that the ability to perform accurate MI is directly related to the model’s propensity in memorizing individual data points, and data that are strongly memorized by the model (e.g., outlier samples, mislabeled samples) are known to be easier to be de-identified [8, 13, 63, 70]. Based on this, our goal can be formulated as: how can the users modify their data to amplify the model’s memorization on the target data while preserving high visual quality?

In Section 4.1.1, we first present an initial approach that is highly effective in inducing the target samples to be memorized by the model, but it also completely destroys the visual information in the data. We use this effective yet unrealistic method as a strawman approach, and then introduce our proposal of a two-step solution, which can greatly reduce the visual distortion to the data while still fulfilling our goal.

4.1.1 An initial strawman approach. Inspired by the observation that atypical samples presented in the dataset are easier to be de-identified (e.g., mislabeled samples [63], or samples that are out-of-distribution - OOD [8]), a straightforward approach is to directly *replace* the original target samples as OOD samples.

There are various ways to generate OOD samples, and we develop a simple method that creates samples with random color stripes, such as the one shown below (the method’s details are in Section 5). Meanwhile, other alternative methods such as using samples from an OOD dataset may also be considered (Appendix B.1.3).

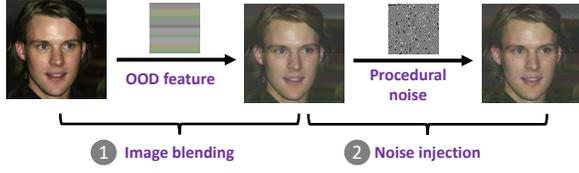
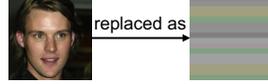


Figure 3: The two-step data marking process: (1) blend the original samples with OOD feature; (2) inject procedural noise. These subtle changes can induce the target samples to be strongly memorized by the model trained on them.

We first conduct an experiment to validate the effectiveness of the above approach. We use the CIFAR100 dataset with half of the training set for model training. We consider 100 target users from different classes replacing their data as some random OOD samples, and each contributes a small number of samples to the training set (0.1% in proportion). In terms of the signal function for MI, we follow prior work [53, 69] to use prediction loss values.



To perform MI, we compare the prediction loss on the target OOD samples (from the member users) and some non-target OOD samples (from the non-member users). We find that the model indeed strongly memorizes those OOD samples present in the training set, for which we observe a 100% TPR@1% FPR. In comparison, if the target data are unmodified, the model only exhibits weak memorization on them, which yields merely 1.28% TPR@1% FPR.

Limitation. Despite its effectiveness, this approach completely removes the visual information in the data (e.g., the facial identity is no longer recognizable), and they are also easy to notice. Next, we present two methods that can greatly reduce the visual distortion to the target data, while still enabling them to be strongly memorized by the model (Fig. 3).

4.1.2 Step 1: Image blending. As discussed earlier, the marked samples should contain the original feature (to preserve visual information), as well as the OOD feature (for provenance purposes). Inspired by the common image blending technique [67, 73], we propose to blend the target sample with the OOD feature:

$$x \oplus (x_{ood}, m) = m \cdot x + (1 - m) \cdot x_{ood}, \quad (1)$$

where x is the original image, x_{ood} is the OOD feature, and m moderates the contribution of different features.

By using a large m , we can largely preserve the high image quality (e.g., the center image in Fig. 3 is marked with $m = 0.7$) while keeping the OOD feature to amplify the model’s memorization on the resulting samples. We find that merely blending the OOD feature into the target samples ($m = 0.7$) can improve the TPR@1% FPR from 1.28% to 16.44%.

Limitation. Although the TPR is largely improved, it is still too low. We find that this is due to the existence of the original feature, which hinders the model’s memorization on the OOD feature. We next explain how to mitigate this.

4.1.3 Step 2: Noise injection. Our idea is to inject a small amount of noise to suppress the influence of the original feature, and have

the model memorize the OOD feature. For this, we draw inspiration from the concept of *adversarial samples*, which works by adding imperceptible noise to the inputs and cause them to be misclassified by the model [16, 38, 42]. The perturbation can thus be viewed as the noise that can suppress the influence of the original feature and cause the model to predict the wrong label.

In our context, we can inject adversarial perturbation into the users’ target samples and prompt the model to memorize the OOD feature. There are different ways to generate the perturbation. Optimization-based approach requires access to some proxy model to generate the perturbation [38, 42], which can still be challenging for the ordinary users we consider.

We therefore resort to another optimization-free strategy that does not assume any additional access. In particular, Co et al. [16] propose that *procedural noise functions* can be used to generate input-agnostic adversarial perturbation. These functions are commonly used in computer graphics to generate different textures and patterns and enrich the image details [16]. Co et al. find that the procedural noise patterns share visual resemblance with existing adversarial perturbation patterns, and demonstrate that they can be similarly used to construct adversarial samples.

This inspires our solution in injecting *perlin noise* [16] to the target samples, and its generation is as follows.

$$S_{perlin}(x, y) = \sum_{n=1}^{\Omega} p(x \cdot \frac{2^{n-1}}{\lambda_x}, y \cdot \frac{2^{n-1}}{\lambda_y}), \quad (2)$$

$$G_{perlin}(x, y) = \sin((S_{perlin}(x, y)) \cdot 2\pi\phi_{sine}), \quad (3)$$

where $\lambda_x, \lambda_y, \Omega, \phi_{sine}$ are to control the noise value at point (x, y) . λ_x, λ_y are the wavelengths, Ω is the number of octaves that contribute to the most visual change, ϕ_{sine} is the periodicity of the sine function that creates distinct bands in the image to achieve a high frequency of edges, which can achieve more distinct visual patterns [16]. We follow Co et al. to use L_{∞} norm as the perturbation budget δ and create random perlin noise [16] for data marking.

Generating the perlin noise is a lightweight process that does not assume access to a proxy model. Despite its simplicity, we find that injecting a small amount of perlin noise (with $\delta = 8/255$) is very effective and it can further increase the TPR@1% FPR from 16.44% to 57.8% (a 3.5x improvement).

Overall, the two-step marking process is able to increase the TPR@1% FPR from 1.28% to 57.8%. While the TPR can be further improved by using a smaller m for image blending or larger δ for noise injection, it would increase the visual distortion (see Appendix B.1.2 for an evaluation). We thus propose another innovation in the MI process to improve the TPR, and without degrading the image quality.

4.2 Data Provenance via Membership Inference

To start with, we note that so far we have been considering MI on a common *instance* basis [6, 8, 55, 69], which uses the per-instance loss as the signal function value for MI. We first explain why this approach would incur low TPR under a low FPR, and then present a specialized MI process to overcome it.

4.2.1 Why instance-based MI suffers from low TPR. It uses the per-sample loss for MI, which can incur low TPR even if the

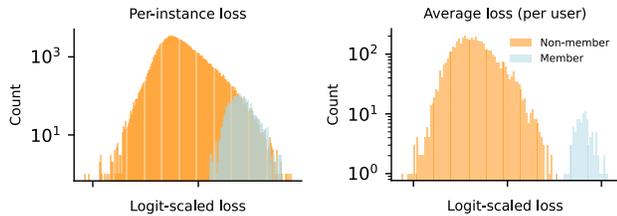


Figure 4: Comparison of using the per-instance loss and the average loss across the samples by each user (each contributes 0.1% of the training data) as the signal function for MI verification. The former yields 57.8% TPR@1% FPR; while the latter has 100% TPR (our proposal).

prediction loss on the target marked samples are very low. This is because many non-member samples may also have low loss values. To control the MI process under a low FPR, only a small fraction of member samples with significantly lower prediction loss are recognized as members, thereby resulting in low TPR.

For instance, the left figure in Fig. 4 shows the individual loss values on different samples, where many non-member samples also yield low prediction loss (the low loss values are shown towards the right side because they are plotted in the logit-scaled form [8] for visualization purposes). This leads to a 57.8% TPR@1% FPR.

4.2.2 Set-based membership inference for reliable data provenance. To overcome the above, our key insight is that, although each target user only contributes a small number of samples to the model’s training set (e.g., 25 out of 25,000 instances), they can still leverage the *collective* information across their own samples to construct a more reliable signal function for MI, and reduce FPR.

To realize this, we start by identifying a pattern that can distinguish the model’s behavior on the member and non-member samples, and it is as follows. For the *non-member samples*, although a few of them may have low loss, the majority of them would incur much higher prediction loss, because these samples are not used for training the model (e.g., see the left of Fig. 4). In contrast, the *member samples* predominantly have very low prediction loss.

Therefore, we make the observation that *the average loss for the samples by the non-member users are higher than that by the member users*. A visualization for this is on the right of Fig. 4. This leads to our proposal of the *set-based* MI process, which leverages the *average loss* of the small set of target samples by each user as the signal function for MI. We explain the verification process next.

4.2.3 Provenance verification process. Algorithm 1 outlines the process. Let each user possess k samples, where k is a small number compared with the dataset size (e.g., 25 out of 25,000 samples). \mathcal{P}_{out} denotes a set of non-member data to control the MI process at the low FPR regime (Section 3.1.1), which are also marked with some random outlier features and perlin noise (data-marking details in Section 5). By comparing the model’s behavior on the target data and the data in \mathcal{P}_{out} , the auditing user U is to detect whether his/her target data are used to train F .

The first step is to derive an inference threshold c based on a pre-defined FPR (α). α is a small number such as 1% or 0% to ensure low false positive. To derive the threshold c , U first computes

Algorithm 1 Set-based MI verification process

Input: F : The target DL model;
 U : The auditing user (each user with k target samples);
 \mathcal{P}_{out} : Users from the out world (i.e., non-member data);
 α : False positive rate (FPR) to control the verification process (e.g., 0.1%);
 $\mathcal{L}(F, (x, y))$: Loss function (e.g., cross-entropy loss);

- 1: **function** DATA PROVENANCE($F, U, \mathcal{P}_{\text{out}}, \alpha, \mathcal{L}$)
- 2: $\forall U_i \in \mathcal{P}_{\text{out}}$, compute $\mathcal{L}(F, U_i)$, where $U_i = \{(x_j^i, y_j^i)\}_{j=1}^k$
- 3: $\mathcal{L}_{\text{avg}} \leftarrow \{\mathcal{L}(F, U_1), \dots, \mathcal{L}(F, U_n)\}$, $n = |\mathcal{P}_{\text{out}}|$
- 4: $c \leftarrow \alpha$ -percentile of the CDF for the loss histogram of \mathcal{L}_{avg}
- 5: /* Data provenance based on the inference threshold c */
- 6: Compute $\mathcal{L}(F, U)$, where $U = \{(x_j, y_j)\}_{j=1}^k$
- 7: $\mathcal{L}_{\text{avg}}^U \leftarrow \mathcal{L}(F, U)$
- 8: **if** $\mathcal{L}_{\text{avg}}^U < c$ **then**
- 9: U ’s data are used to train F (i.e., member)
- 10: **else**
- 11: U ’s data are not used to train F (i.e., non-member)
- 12: **end if**
- 13: **end function**



Figure 5: Visualization of the original and marked samples.

the empirical loss histogram for data in the out world \mathcal{P}_{out} . For each non-member user in \mathcal{P}_{out} , the auditing user U first computes the per-instance loss on their data (line 2), and then derives the average loss for each (line 3). Line 4 estimates the CDF for the loss distribution, and computes its α -percentile to derive a threshold c [69].

Next, the auditing user U similarly computes the average loss on his/her samples, and compares it against the MI threshold c . If the loss is lower than c , the user’s data are deemed to be used for training the model (vice versa).

5 EVALUATION

Datasets and model training. We consider six common benchmark datasets, including CIFAR100 [34], CIFAR10 [34], ArtBench [40], CelebA [44], TinyImageNet [35] and ImageNet-1k [18]. We first consider the first five datasets and the evaluation on the ImageNet-1k dataset is in Section 5.3. For CIFAR10 and CIFAR100, we train a WideResNet-28-4 model on 25,000 samples for 100 epochs, with a learning rate of 0.1 (decayed by 5 at epoch 60, 80, 90), a weight decay of $5e-4$ and momentum of 0.9. For TinyImageNet and ArtBench, we fine-tune an ImageNet-pretrained ResNet-18 model on 25,000 samples. For CelebA, we first follow [1] to create a 307-class subset from the dataset, and then similarly fine-tune an ImageNet-pretrained model on 2,000 samples. For fine-tuning, we use a small fine-tuning rate of 0.01 with 30 epochs and a momentum of 0.9. In all datasets except CelebA (where we find that fine-tuning on the entire set yields better accuracy), we use 20% of the training data as the validation set. We also use common data augmentations such as random cropping and horizontal flip.

MembershipTracker configuration. We explain the default configurations for the main experiments, and we study different configurations in the ablation study (Appendix B).

As explained earlier, we create outlier features as random color stripes, each with 16 stripes from 11 different common colors. Such a large feature set (11^{16} choices) can support multiple users to generate their own features for data marking, and we use a image blending ratio $m = 0.7$. Next, we follow Co et al. [16] to generate random perlin noise for each sample with $\delta = 8/255$. We assume each user possesses a small number of samples from a given class, which constitute 0.1% of the training set.

Fig. 5 provides additional visualizations of the original and marked samples, and we quantitatively evaluate the visual quality of the images in Appendix A.

Unless otherwise stated⁴, we follow the common practice to use TPR@low FPR for evaluation and we use a 0% FPR threshold (to minimize false accusations). We use 5,000 non-member users for controlling the FPR and these samples may be collected by sampling random data from the data distribution [6, 8, 69]. We use all the samples that are not used for training as the non-member set. Due to the limited size of the non-member set, the non-member users' samples are drawn randomly from the non-member set (with replacement), and these are similarly marked with random outlier features and perlin noise.

To perform MI, we directly use the cross-entropy loss as the signal function value, which is easy to compute by the users. We consider the global-threshold method by Carlini et al. [8] as the baseline method; and we do not choose the shadow-model-based approaches [6, 8, 69]. The reason is that, even though these methods can achieve better results, they still yield limited increase of MI success (Section 4); and more importantly, training shadow models is prohibitively challenging for ordinary personal users to carry out (Section 3.1.1).

The closest work to ours in improving the MI success without shadow-model calibration is Tramer et al. [63]. However, their technique assumes the users can manipulate the data labels, which excludes the real-world scenarios where the data labels are assigned by the model creator or an external service [18, 20, 67]. Though *MembershipTracker* does not impose such an assumption, for completeness, we still quantitatively compare it with the method by Tramer et al. (Appendix D).

5.1 Results

We first evaluate the single-target-user setting, and then the multi-target setting.

Single-target evaluation. For each experiment, we randomly choose 5 target users from different classes and report the average results. Table 1 presents the results.

As shown, with *MembershipTracker*, the users can reliably trace the provenance of their marked data with 100% TPR@0% FPR. In comparison, without *MembershipTracker*, performing standard instance-based MI on the unmodified target samples cannot reliably identify the member samples under a 0% FPR threshold and it has 0% TPR. This is because: 1) the unmarked samples cannot be strongly memorized by the model; and 2) the instance-based

⁴We will use a different metric in Section 5.2 for reasons we will explain there.

Table 1: Evaluation on different datasets (parenthesized numbers show accuracy diff. with the original models).

Dataset	TPR@0% FPR	Model accuracy
CIFAR100	100.00	66.34 (-0.41)
CIFAR10	100.00	90.64 (-0.50)
CelebA	100.00	69.41 (-0.63)
ArtBench	100.00	59.46 (-0.42)
TinyImageNet	100.00	67.22 (-0.40)

Table 2: Evaluation on different models.

Model	TPR@0% FPR	Model accuracy
WideResNet	100.00	66.34 (-0.41)
ResNet	100.00	64.49 (-0.69)
ResNext	100.00	63.43 (-0.47)
DenseNet	100.00	67.14 (-0.16)
GoogleNet	100.00	69.09 (-0.33)
SeNet	100.00	65.26 (-0.37)

Table 3: Evaluation on different training-set sizes.

Training-set size	TPR@0% FPR	Model accuracy
25,000	100.00	66.34 (-0.41)
20,000	100.00	63.73 (-0.16)
15,000	100.00	58.99 (-0.83)
10,000	100.00	50.28 (-0.09)
5,000	100.00	35.74 (-0.10)

MI process fails to leverage the collective information across the target samples by each user. These two limitations are overcome by *MembershipTracker*'s data marking and set-based MI process; and the detailed ablation study on these two components is presented in Appendix B.

Meanwhile, since the target samples only occupy a small fraction of the training set (0.1%), the models still maintain high performance, and have <1% accuracy difference compared with the ones trained without *MembershipTracker*.

Evaluation on different architectures and training-set sizes. We next consider additional evaluation (using CIFAR100 dataset) on different model architectures (including WideResNet [71], ResNet [22], ResNext [68], SeNet [27], GoogleNet [61] and DenseNet [28]), as well as different training-set sizes (from 5,000 to 25,000). We report the results in Table 2 and Table 3, where we observe a similar trend as before (and thus we omit the discussion).

Multi-target evaluation We now evaluate how *MembershipTracker* can support multiple users simultaneously (using CIFAR100 dataset), and the results are in Table 4. We study a large number of target users (more than the number of classes in the dataset), which encompasses the settings where the users are from different classes and from the same class.

Table 4: Evaluation on multiple target users. Supporting more users indicates more specially-marked samples in the dataset, which can lead to higher accuracy drop; but *MembershipTracker* still maintains high auditing performance.

Num of target users	20	50	100	200	300	500
TPR@0% FPR	100.00	100.00	100.00	100.00	100.00	100.00
Accuracy drop	-0.44	-0.56	-1.39	-2.32	-4.47	-8.93

Supporting multiple users requires the model to memorize more target samples (from different users). However, DL models are known to be capable of memorizing a large corpus of data (e.g., to encode sensitive information [13, 56]). Likewise, in our case, the model still strongly memorizes the marked samples and enables *MembershipTracker* to achieve 100% TPR@0% FPR even under multiple target users.

Meanwhile, as the number of target users increases, the model also has higher accuracy drop. E.g., supporting 100 target users (i.e., 10% of training set are marked) incurs a 1.39% accuracy drop, while supporting 500 users has a higher 8.93% drop. This is understandable as supporting more users implies more training samples are marked. Supporting more target users will result in a larger accuracy drop - we defer this to future work.

5.2 Robustness to Countermeasures

While Section 5.1 demonstrates *MembershipTracker*'s capability in enabling the users to perform reliable data provenance, a deliberate model creator can also employ countermeasures to sabotage *MembershipTracker*.

Table 5: Overview of evaluated countermeasures.

Defense level	Defense type
Model level	① MI defense; ② Adversarial augmentation; ③ Model fine-tuning; ④ Model pruning
Input level	⑤ Out-of-distribution and spurious features detection
Output level	⑥ Add noise to the output, or return only the label

To understand this, we evaluate a total of 6 types of countermeasures in the model, input and output level (Table 5), and we use the CIFAR100 dataset. For the model-level defenses that require model training, we first consider the single-target setting (similar to Section 5.1), and then the multi-target setting. For the input- and output-level defenses, we consider the more challenging multi-target setting (100 targets in total).

Evaluation metric. In this section, we use a different evaluation metric, FPR@high TPR. This is because using TPR@low FPR (as in our previous evaluation) requires choosing a low FPR threshold and evaluating the corresponding TPR. However, the pre-defined FPR threshold may be too conservative or aggressive.

We use the results from the DPSGD experiment (Section 5.2.1) to explain the issue involved. In this example, using a 0.1% FPR threshold is too conservative as *MembershipTracker* has the same 100% TPR for the noise multiplier $\sigma = 0.01$ and $\sigma = 0.03$. Likewise, using a 0% FPR threshold is too aggressive as *MembershipTracker* has the same 0% TPR for $\sigma = 0.2$ and $\sigma = 0.3$. In both cases, using a pre-defined TPR may not adequately reflect *MembershipTracker*'s performance changes in different settings.

To overcome the above issue, we advocate the use of FPR under 100% TPR. This metric is able to properly compare *MembershipTracker*'s performance under different settings (e.g., Table 6), and hence we use it for the evaluation in this Section, as well as Section 6.2 and Appendix B (for the same reason explained above).

5.2.1 MI defense. We first evaluate several representative privacy defenses for mitigating MI (five in total).

Table 6: Evaluating *MembershipTracker* against DP-SGD with different amounts of noise injected.

Noise multiplier σ	FPR@100% TPR	Model accuracy
w/o defense	0.00	66.34
0.01	0.01	60.10
0.03	0.02	57.26
0.05	0.38	54.70
0.2	1.07	51.56
0.3	3.10	50.03

DP-based defense. Differentially private (DP) training [4] is a principled defense that can bound the influence of any training samples to the model, by means of gradient clipping and noise injection to the clipped gradients. We adopt the implementation from Aerni et al. to consider the strong DP-SGD baselines that comprise various DP-training techniques such as custom initialization scheme, augmentation multiplicity [5]. We also follow their work to tune hyperparameters for high model utility. We use a tight clipping norm of 1, and inject different amounts of noise (σ).

The results are in Table 6. DPSGD with a loose DP bound is known to be able to reduce the exposure from the *instance-based* MI methods that aim to de-identify some individual training instances [5, 8], and our results also validate this (e.g., in our context, the FPR@100% TPR under instance-based MI is increased to 47.22%, with $\sigma = 0.01$). With that said, *MembershipTracker* can still overcome the defense via its *set-based* MI process, because the *average loss* of the samples by the member users, though increased due to defense, are still lower than that by the non-member users (average 0.1642 vs. 4.130). This avails *MembershipTracker* to maintain a low 0.01% FPR@100% TPR (and the defense already incurs 6.24% accuracy drop). Increasing the amount of noise can further increase the FPR by *MembershipTracker*, but it also inflicts a larger accuracy drop (undesirable).

Empirical defense. Next, we also evaluate two representative defenses that aim to provide strong empirical privacy while preserving model utility: SELENA [62] and HAMP [14]. We follow the original work to set up both techniques. They both can mitigate the model's memorization (increase the prediction loss) on the marked samples, and with small accuracy drop (2.17% and 1.42%). However, we find that the average loss of the samples by the member users are still much lower than that by the non-member users, and thus *MembershipTracker* still achieves a low FPR of 0.97% (on SELENA) and 0.83% (on HAMP).

Finally, we evaluate two common regularization techniques: *early stopping* and *dropout*. We find that *MembershipTracker* can still achieve $< 0.1\%$ FPR@100% TPR, *unless* under excessive regularization (e.g., aggressively early-stopping the training or using a large dropout rate). However, this would also cause a large ($> 10\%$) accuracy drop to the model.

5.2.2 Adversarial augmentation. We consider employing additional augmentation techniques during training to mitigate the model's memorization on the marked samples. We study three strategies: (1) injecting Gaussian noise (we inject zero-mean noise with various standard deviation values (0.02~0.5)); (2) using Mixup [73] (we use seven different alpha parameters from 0.1 to 1.5); and (3) using RandAugment [17] to perform multiple random data augmentations (we consider 3~15 different augmentations).

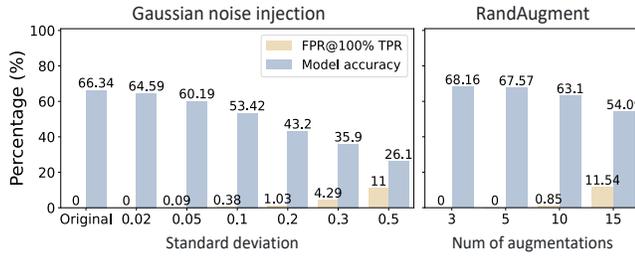


Figure 6: Adversarial augmentation by Gaussian noise injection and RandAugment [17].

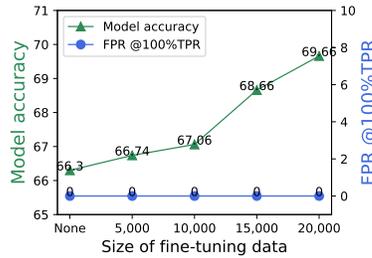


Figure 7: Evaluation with model fine-tuning. Fine-tuning the model with additional data can improve model accuracy (the more data the better), but *MembershipTracker* consistently maintains high auditing performance regardless.

We find that *MembershipTracker* still maintains 0% FPR@100% TPR under Mixup training with different parameters. We report the results for the other two strategies in Fig. 6.

When the samples are moderately augmented, in the form of Gaussian noise or multiple random augmentations, *MembershipTracker* still enables the target users to reliably trace the provenance of their data with low FPR.

Aggressive augmentation such as applying 15 random transformations to each sample can further increase the FPR to 11.54%, but it also reduces the accuracy from 66.34% to 54.09%, as such aggressive augmentation severely degrades the image quality and leads to inferior model performance.

5.2.3 Model fine-tuning. Fine-tuning is a technique performed on additional data that differs from the training data to improve model performance. We evaluate how it may be used to mitigate the model’s memorization on the target samples. We assume an ideal defender with access to an additional set of clean and unmodified samples. We consider multiple settings with different amounts of data (5,000 to 20,000 samples). For each setting, we fine-tune the model using various learning rates from 0.1~0.000001 for 20 epochs, and we report the one with the highest accuracy. The results are reported in Fig. 7.

As shown, using a small set of fine-tuning data (5,000 samples) can increase the model accuracy, and using more data can lead to larger accuracy improvement (understandably so). For instance, using 20,000 samples can increase the accuracy to 69.66%, which translates to a 10% reduction in the test error. Nevertheless, despite its effectiveness in enhancing model accuracy, model fine-tuning still

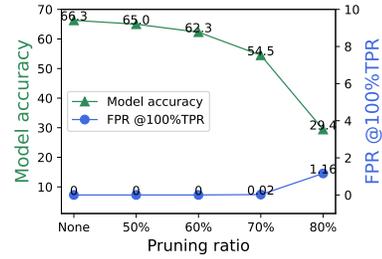


Figure 8: Evaluation with model pruning. *MembershipTracker* is resilient to moderate model pruning. Aggressive countermeasures can slightly increase the FPR incurred by *MembershipTracker*, but at the cost of major accuracy drop.

has limited effectiveness in degrading *MembershipTracker*, which consistently achieves 0% FPR@100% TPR in all cases.

5.2.4 Model pruning. Another related countermeasure we consider is to prune the model’s parameters, with the goal of removing the parameters that are associated with the target samples. We use the popular pruning strategy by Han et al. [21], which sets the parameters with small absolute values to zero and has minimal impact to the model’s performance. We evaluate different pruning ratios and the results are shown in Fig. 8.

We find that *MembershipTracker* is highly resilient even when 70% of the parameters are pruned, where *MembershipTracker* has only 0.02% FPR@100% TPR, but the defense already incurs a 11.8% accuracy drop. More aggressive pruning can slightly increase the FPR, but with a much higher accuracy drop.

5.2.5 Model-level defenses under multi-target setting. The previous model-level defenses are evaluated under the single-target setting. We now evaluate a more challenging setting where *MembershipTracker* needs to support multiple users at the same time. We consider 25 target users, and summarize the key findings below.

Among different defenses, we find that *model fine-tuning* stands out as the most potent solution that outperforms other methods in terms of high mitigation performance without accuracy drop. By fine-tuning the model on 20,000 additional samples, it can increase the FPR@100% TPR to 42.84% without accuracy drop (the model still has 67.30% accuracy).

Hence, unlike the single-target setting where *MembershipTracker* can maintain high performance, supporting multiple target users under model-level defenses remains a major challenge by *MembershipTracker*. We leave the potential improvement of this aspect to future investigation.

5.2.6 Input detection. Given that *MembershipTracker* leverages outlier features and perlin noise to mark the target data, the defender may also perform *out-of-distribution (OOD) detection* or *spurious features detection* on the specially-marked samples. We consider both types of countermeasures and we discuss them next.

Unsupervised OOD detection. We consider eight common detection techniques: Mahalanobis [36], MSP [24], Energy-based OOD [41], RMD [50], kNN [60], KLMatching [23], VIM [65], DICE [59]. We follow the implementation from [33], and use a pre-trained ImageNet model to produce the feature representation. As before, we

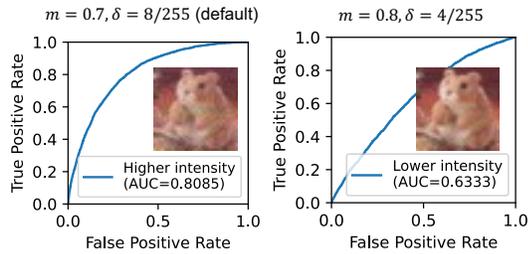


Figure 9: Unsupervised OOD detection. How different data-marking intensities affect the detection performance. Left: Samples marked with higher marking intensity are prone to be detected (0.8085 AUC). Right: Reducing the marking intensity can effectively reduce the AUC to only 0.6333.

assume an ideal defender with access to a small set of clean and unmodified samples (equivalent to 10% of the training set) for fitting the unsupervised detector, which does not assume any prior knowledge of the potential outlier samples.

We report the results on the left of Fig. 9, where, for brevity, we report only the highest AUC among all detectors. We find that the Mahalanobis-based approach [36] achieves the highest performance with 0.8085 AUC. Next, we explain how to (partially) mitigate it.

Our idea is to reduce the amount of distortion applied to the data. For this, we conduct an experiment to reduce the intensity of the outlier features (by increasing the image blending ratio m from 0.7 to 0.8) and perlin noise (by reducing the perturbation budget δ from $8/255$ to $4/255$). The results are shown on the right of Fig. 9, where the highest AUC is reduced from 0.8085 to 0.6333, which translates to a 57% reduction in the detection performance (over random guessing).

Naturally, reducing the marking intensity would attenuate the model’s memorization on the marked samples, and increase their prediction loss. But we find that the prediction loss on the member samples are still much lower than that on non-members, and thus *MembershipTracker* can overcome it via the set-based MI process (similar to Section 5.2.1), and maintain a 0% FPR@100% TPR⁵.

Supervised OOD detection. Next, we consider a more knowledgeable defender who is aware of the potential data marking strategy. We use the MCHAD method to train a supervised OOD detector [32], and we consider two scenarios below.

(1) We first assume a hypothetical defender with *perfect* knowledge of the data marking strategy adopted by the users, i.e., each sample is marked with features consisting of random color stripes and perlin noise. The result is in the blue line in Fig. 10. As expected, the detector is very capable of detecting the samples that are marked with the same strategy as the one used for training the detector itself (e.g., see the top right illustration in Fig. 10).

Though this represents the worst-case scenario against *MembershipTracker*, it may be difficult to realize in practice, as the outlier features created by the users for data marking can take *diverse* forms (e.g., as features consisting of random color stripes, or features from an OOD dataset - details next), and the exact one used by the users can be hard to predict.

⁵However, there are certain scenarios where the reduced data-marking intensity would lead to lower auditing performance, and an evaluation for this is in Appendix B.1.2.

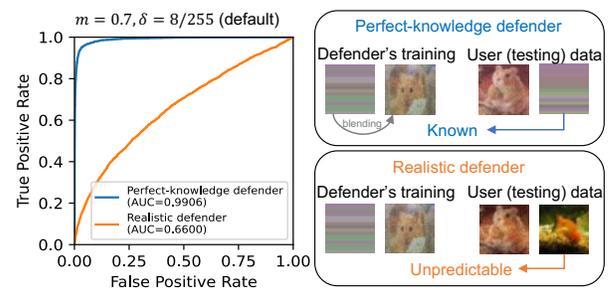


Figure 10: Supervised OOD detection. Blue line assumes a defender with *perfect* knowledge of the type of outlier features adopted by the users for data marking. Orange line considers a *realistic* defender with knowledge of the possible types of outlier features, but the exact one followed by the users to mark their data is unpredictable.

(2) Hence, a more realistic setting is to consider a defender with knowledge of the possible types of outlier features, but is not aware of the exact one adopted by the users to mark their data. In this case, the outlier features in the samples used for training the detector are not in the same type as those in users’ actual data (e.g., see the bottom right illustration in Fig. 10).

Specifically, we train the detector on samples marked with random color stripes, and evaluate it on samples marked with the features from an OOD (TinyImageNet) dataset (this is an effective alternative strategy we study in Appendix B.1.3). The result is in the orange line in Fig. 10, and the detection performance degrades considerably under such a realistic setting (with only 0.66 AUC).

Although the defender can attempt to train the detector on samples marked with different types of outlier patterns, with the hope that some of them may overlap with those in users’ data, this essentially turns the detection into a cat-and-mouse game, and may not be a desirable practice.

Spurious features detection is another related method. We evaluate a leading technique by Neuhaus et al. [48] based on neural PCA components. But we find that it also has limited success in detecting the marked samples (0.6181 AUC).

5.2.7 Output perturbation. *MembershipTracker* leverages the model’s outputs to perform MI auditing, and we consider two strategies to perturb the model’s outputs: (1) injecting Gaussian noise to the outputs; and (2) returning only the top-1 prediction label [15, 30]. However, we find that *MembershipTracker* still maintains high auditing success even with the perturbed outputs.

(1) For the first method, we find that injecting Gaussian noise (using zero-mean noise with standard deviation σ in 0.5~5) to the output vectors is not an effective solution, and *MembershipTracker* can still maintain very low FPR even if the defense has caused severe accuracy degradation. E.g., injecting Gaussian noise with $\sigma = 3$ degrades the accuracy from 65.32% to 52.05%, under which *MembershipTracker* still achieves a 0% FPR@100% TPR. We use Fig. 11 to explain next.

As in the first two figures of Fig. 11, injecting noise to the outputs increases the prediction loss on the member samples (and non-members too). Hence, more member samples (the blue area in the

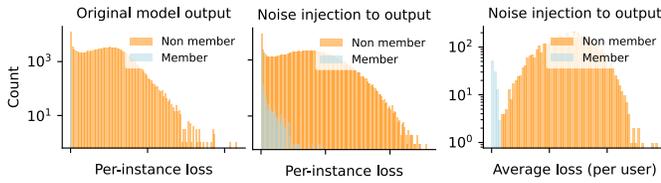


Figure 11: Output noise injection increases the prediction loss of the member samples (from the leftmost to the middle figure), but their loss are still lower than many non-member samples’. Thus, *MembershipTracker* can leverage the average loss of the samples by each user to achieve a 0% FPR@100% TPR (the rightmost figure).

middle figure of Fig. 11) now have similar loss as the non-member samples (undesirable).

Fortunately, *MembershipTracker* can still overcome this via its set-based MI process. This is because, the prediction loss of the member samples, though increased, are still much lower than that of many non-member samples. Thus, there is still a clear difference between the average loss for the member and non-member users’ samples (the rightmost figure in Fig. 11), from which *MembershipTracker* is able to achieve a 0% FPR@100% TPR.

(2) In the **label-only** setting, since the class probabilities are not available, we use a simple method to compute a proxy for the prediction loss: we assign a loss value of 0 for the correctly-classified samples and 1 for other samples. In this case, all the correctly-classified non-members have the same 0 loss as the member samples. However, we find that *MembershipTracker* can still overcome the defense and achieve a 0.04% FPR@100% TPR.

The reason is that there is a major accuracy difference between the member and non-member users’ samples: 100% vs. 22.02% (the latter yield low accuracy as they are marked with the random outlier features and perlin noise, but they are *not* present in the training set). Thus, the “average loss” for the samples by the member users are still lower than that by the non-member users. This enables *MembershipTracker* to succeed even under the label-only setting.

5.3 Evaluation on ImageNet Training

This section evaluates *MembershipTracker* under the full-sized ImageNet-1k training (over 1.28 million training samples). We train a ResNet50 [22] and Swin-Transformer [43], and they have 75.09% and 80.29% accuracy respectively (both with <1% accuracy drop). We consider the challenging setting of supporting multiple users in different classes (1,000 target users in total). Each user contributes 125 samples to the training set, which is <0.01% in proportion.

We find that even under the large-scale training setting, the marked samples can still be strongly memorized by the models, which enables *MembershipTracker* to achieve 100% TPR@0% FPR.

Next, we evaluate how the performance varies with *lesser* samples for marking (using ResNet50): we consider 63 and 32 samples by each user, which amount to $\sim 0.005\%$ and $\sim 0.0025\%$ of the training set. The results are in Fig. 12.

MembershipTracker can still maintain 100% TPR@0% FPR, when the marking ratio is reduced from 0.01% to 0.005%. However, when the marking ratio is further reduced to 0.0025%, *MembershipTracker*

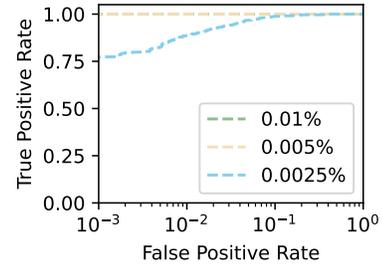


Figure 12: Evaluating *MembershipTracker* under full-sized ImageNet training with different portions of samples for data marking.

experiences lower auditing performance - this trend is similar to that in our ablation study in Appendix B.1.1 and thus we defer the detailed discussion to Appendix B.1.1 (due to space constraints). With that said, even under such a challenging case, *MembershipTracker* still achieves reasonably good performance with >75% TPR@0.1% FPR (the blue line of Fig. 12). *To the best of our knowledge, MembershipTracker is the first technique that can scale to support reliable data auditing in large-scale training setting.*

6 DISCUSSION

In this section, we first discuss how our assumption on the access to a set of non-member samples can be relaxed, and then discuss the limitations of *MembershipTracker*.

6.1 (Relaxed) Assumption on Non-member Set

Do non-member samples need to be definitely outside the training set? Recall that *MembershipTracker* follows a *common* assumption [6, 8, 67, 69] that users have access to a set of non-member samples (for calibrating the MI process under low FPR regime). However, in practice, it can be challenging for users to determine whether these non-member samples are definitely not used for training the target model. For instance, an auditing user may collect some random facial images from the Internet as non-member samples, but it is *possible* that these samples already have been used for training the target model.

To understand how this may affect *MembershipTracker*’s performance, we study a hypothetical scenario where all the non-member samples (used for data auditing) have been used for training the target model. Even under such a challenging case, we find that *MembershipTracker* can still maintain reliable auditing performance with 100% TPR@0% FPR. We explain the reason below.

Recall that *MembershipTracker* performs data marking on non-members using random outlier features. These random features are *highly unlikely* to be used for training (e.g., among 10^{16} possible random color stripe features), and thus they significantly increase the prediction loss of non-member samples, e.g., reducing their accuracy from >99% (this is high because these non-member samples are used for training) to <30% after data marking. Consequently, the marked non-member samples still exhibit much higher prediction loss than marked member samples, from which *MembershipTracker* still maintains 100% TPR@0% FPR for reliable data auditing.

Table 7: Evaluation on the partial inclusion of user data in model training. With lesser data (50 samples per user) included for training, the average loss on the entire set of samples by each target user increases. This in turn (1) increases the FPR incurred by *MembershipTracker*, and (2) leads to larger accuracy drop due to lesser data for training.

% of user samples used for model training	FPR@100% TPR	Accuracy drop
100% (50/50)	0.00	0.00
90% (45/50)	0.00	0.12
70% (35/50)	0.00	4.83
50% (25/50)	0.34	12.34
40% (20/50)	0.84	14.76
30% (15/50)	2.42	21.37

Overall, we find that with *MembershipTracker*, users only need to collect some data from the underlying population, and these data can be *any* data; they do not necessarily have to be outside the training set of the target model. This demonstrates the broader practicality of *MembershipTracker*.

6.2 Limitations

Next, we discuss the three main limitations by *MembershipTracker*.

① *MembershipTracker* works by marking a small number of target samples to be strongly memorized by the model, and detecting whether the model memorizes these samples. However, a malicious developer may seek to jeopardize *MembershipTracker*'s performance, by using only a *subset* of data from each user for training, so that not all target samples are memorized by the model.

To understand how this affects *MembershipTracker*, we perform an experiment where we assume the model owner collects 50 samples from each user (0.2% of the training set), and he/she includes only a subset of samples from each user for training. Table 7 shows the results, where we consider five target users (we observe a similar trend when considering the support for more target users). We use FPR@100% TPR as the metric, as explained earlier in Section 5.2.

The FPR incurred by *MembershipTracker* grows as the number of marked samples used for training decreases. This is because those samples that are not used for training would yield higher prediction loss, which in turn increases the average loss on the entire set of marked samples, and affects *MembershipTracker*. On the other hand, while using a subset of user data can degrade *MembershipTracker*'s performance, it also results in lesser data for model training and causes an accuracy drop (third column in Table 7).

Overall, we find that *MembershipTracker* can still maintain < 0.5% FPR@100% TPR, even when only 50% of the user data are used for training. In this case, the model already suffers from a 12% accuracy drop. Using lesser amounts of user data for training can further degrade *MembershipTracker*'s auditing performance, but it will also result in a larger accuracy drop (undesirable).

② Despite our best effort in keeping the targeted changes from severely distorting the marked data, the artifacts created by *MembershipTracker* are still visible under close inspection (e.g., Fig. 5). However, performing manual inspection to filter out the specially-marked samples by *MembershipTracker* can still be a challenge, particularly under the large-scale training setting (e.g., with millions

of images). With that said, creating more subtle targeted changes still remains an important venue for future investigation.

③ Finally, since *MembershipTracker* targets the model's propensity in memorizing data, in principle, any effective measures that can mitigate the model's memorization on training data would be able to degrade the performance by *MembershipTracker*. Indeed, our evaluation in Section 5.2 comprehensively considers multiple measures applied to model, input, and output levels; we find that though some approaches such as differential privacy can be configured to degrade the auditing success achieved by *MembershipTracker*, they also result in utility loss to the model (undesirable).

Given that building high-performance models often represents the model creators' foremost interest, the development of tools like *MembershipTracker* can prompt the practitioners to reconsider the choice of resorting to other legal data acquisition means and building high-performance models using the authorized data.

7 CONCLUSION AND FUTURE WORK

This work presents *MembershipTracker*, a practical data provenance tool that can support ordinary users to audit whether their specific data are used to train deep learning models without their permission. *MembershipTracker* consists of a lightweight data marking process that can mark the target data with small and targeted changes, and a high-power set-based membership inference process that can reliably trace the provenance of the target data. By merely marking a small fraction of samples (0.005%~0.1% of the dataset), *MembershipTracker* effectively empowers the users to reliably audit the usage of their data (average 100% TPR@0% FPR), and it is also scalable to the large-scale (ImageNet) training setting. If their data are found to be misused, the users can take legal action or request the "right to be forgotten" in accordance with privacy regulations such as GDPR [3] (e.g., based on techniques on machine unlearning [7]).

Finally, we discuss future work directions. First, while *MembershipTracker* retrofits the conventional membership inference (attack) to enable responsible AI, it also faces the risk of being misused by the malicious party. Notably, the data marked with *MembershipTracker* would now become easier to be de-identified by *any* party with access to them. However, *MembershipTracker* should be applied only when the users do *not* want their data to be used for training DL models, in which membership privacy may not be of the users' concern. By contrast, if the users are willingly sharing their data for model training, they should refrain from using *MembershipTracker* to avoid unnecessary membership exposure. Future studies can explore potential usage cases where *MembershipTracker* is applicable *and* protecting the membership privacy of the marked samples is necessary, and consider the further adaptation of *MembershipTracker* to cater to such a need.

Second, future work can explore the potential of extending *MembershipTracker* to other domains like generative models. Large generative models are often trained on massive amount of data collected from the Internet [49, 51], and are the frequent targets where unauthorized use of personal/copyrighted contents arises [19, 64]. Thus, future work can investigate whether the large generative models would also be prone to memorize the samples marked with *MembershipTracker* under the billion-scale training environment.

ACKNOWLEDGMENT

We thank the anonymous reviewers of CCS'25. This work was funded in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), a grant from the National Research Council of Canada (NRC), and a Four Year Fellowship from the University of British Columbia.

REFERENCES

- [1] [n. d.]. CelebA HQ Face Identity and Attributes Recognition using PyTorch. <https://github.com/ndb796/CelebA-HQ-Face-Identity-and-Attributes-Recognition-PyTorch/tree/main>.
- [2] [n. d.]. A collection of PyTorch metrics implementations. <https://lightning.ai/docs/torchmetrics/stable/>.
- [3] [n. d.]. General Data Protection Regulation. <https://gdpr-info.eu>.
- [4] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*.
- [5] Michael Aerni, Jie Zhang, and Florian Tramèr. 2024. Evaluations of Machine Learning Privacy Defenses are Misleading. *arXiv preprint arXiv:2404.17399* (2024).
- [6] Martin Bertran, Shuai Tang, Aaron Roth, Michael Kearns, Jamie H Morgenstern, and Steven Z Wu. 2024. Scalable membership inference attacks via quantile regression. *Advances in Neural Information Processing Systems* (2024).
- [7] Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. Machine unlearning. In *IEEE Symposium on Security and Privacy (SP)*.
- [8] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. 2022. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE.
- [9] Harsh Chaudhari, Giorgio Severi, Alina Oprea, and Jonathan Ullman. 2023. Chameleon: Increasing Label-Only Membership Leakage with Adaptive Poisoning. *arXiv preprint arXiv:2310.03838* (2023).
- [10] Guangke Chen, Yedi Zhang, and Fu Song. 2023. SLMIA-SR: Speaker-Level Membership Inference Attacks against Speaker Recognition Systems. *arXiv preprint arXiv:2309.07983* (2023).
- [11] Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, and Yang Zhang. 2023. FACE-AUDITOR: Data Auditing in Facial Recognition Systems. *arXiv preprint arXiv:2304.02782* (2023).
- [12] Yufei Chen, Chao Shen, Yun Shen, Cong Wang, and Yang Zhang. 2022. Amplifying membership exposure via data poisoning. *Advances in Neural Information Processing Systems* (2022).
- [13] Zitao Chen and Karthik Pattabiraman. 2024. A Method to Facilitate Membership Inference Attacks in Deep Learning Models. *arXiv preprint arXiv:2407.01919* (2024).
- [14] Zitao Chen and Karthik Pattabiraman. 2024. Overconfidence is a Dangerous Thing: Mitigating Membership Inference Attacks by Enforcing Less Confident Prediction. In *ISOC Network and Distributed System Security Symposium (NDSS)*.
- [15] Christopher A Choquette-Choo, Florian Tramèr, Nicholas Carlini, and Nicolas Papernot. 2021. Label-only membership inference attacks. In *International Conference on Machine Learning*. PMLR.
- [16] Kenneth T Co, Luis Muñoz-González, Sixte de Maupéou, and Emil C Lupu. 2019. Procedural noise adversarial examples for black-box attacks on deep convolutional networks. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*.
- [17] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*.
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee.
- [19] Pranav Dixit. 2023. Meet The Three Artists Behind A Landmark Lawsuit Against AI Art Generators. *BuzzFeedNews* (2023).
- [20] Josh Dzieza. 2023. <https://www.theverge.com/features/23764584/ai-artificial-intelligence-data-notation-labor-scale-surge-remotasks-openai-chatbots>. *The Verge* (2023).
- [21] Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems* (2015).
- [22] 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*.
- [23] Dan Hendrycks, Steven Basart, Mantas Mazeika, Andy Zou, Joseph Kwon, Mohammadreza Mostajabi, Jacob Steinhardt, and Dawn Song. 2022. Scaling Out-of-Distribution Detection for Real-World Settings. In *International Conference on Machine Learning*. PMLR.
- [24] Dan Hendrycks and Kevin Gimpel. 2016. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In *International Conference on Learning Representations*.
- [25] Kashmir Hill. 2020. Facial recognition's 'dirty little secret': Millions of online photos scraped without consent. *The New York Times* (2020).
- [26] Hongsheng Hu, Zoran Salcic, Gillian Dobbie, Jinjun Chen, Lichao Sun, and Xuyun Zhang. 2022. Membership inference via backdooring. In *The 31st International Joint Conference on Artificial Intelligence (IJCAI-22)*.
- [27] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [28] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [29] Hanxun Huang, Xingjun Ma, Sarah Monazam Erfani, James Bailey, and Yisen Wang. 2021. Unlearnable examples: Making personal data unexploitable. *arXiv preprint arXiv:2101.04898* (2021).
- [30] Z. Huang, N. Z. Gong, and M. K. Reiter. 2024. A general framework for data-use auditing of ML models. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security*.
- [31] Nikhil Kandpal, Krishna Pillutla, Alina Oprea, Peter Kairouz, Christopher A Choquette-Choo, and Zheng Xu. 2023. User inference attacks on large language models. *arXiv preprint arXiv:2310.09266* (2023).
- [32] Konstantin Kirchheim, Marco Filax, and Frank Ortmeier. 2022. Multi-class hypersphere anomaly detection. In *26th International Conference on Pattern Recognition (ICPR)*.
- [33] Konstantin Kirchheim, Marco Filax, and Frank Ortmeier. 2022. PyTorch-OOD: A Library for Out-of-Distribution Detection Based on PyTorch. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [34] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [35] Ya Le and Xuan Yang. 2015. Tiny imagenet visual recognition challenge. *CS 231N* (2015).
- [36] Kimin Lee, Bibek Lee, Honglak Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems* (2018).
- [37] Guoyao Li, Shahbaz Rezaei, and Xin Liu. 2022. User-level membership inference attack against metric embedding learning. *arXiv preprint arXiv:2203.02077* (2022).
- [38] Yingwei Li, Song Bai, Yuyin Zhou, Cihang Xie, Zhishuai Zhang, and Alan Yuille. 2020. Learning transferable adversarial examples via ghost networks. In *Proceedings of the AAAI conference on artificial intelligence*.
- [39] Yiming Li, Yang Bai, Yong Jiang, Yong Yang, Shu-Tao Xia, and Bo Li. 2022. Un-targeted backdoor watermark: Towards harmless and stealthy dataset copyright protection. *Advances in Neural Information Processing Systems* 35 (2022), 13238–13250.
- [40] Peiyuan Liao, Xiuyu Li, Xihui Liu, and Kurt Keutzer. 2022. The artbench dataset: Benchmarking generative models with artworks. *arXiv preprint arXiv:2206.11404* (2022).
- [41] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. 2020. Energy-based out-of-distribution detection. *Advances in neural information processing systems* (2020).
- [42] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. 2016. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770* (2016).
- [43] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- [44] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*.
- [45] Pratyush Maini, Mohammad Yaghini, and Nicolas Papernot. 2020. Dataset Inference: Ownership Resolution in Machine Learning. In *International Conference on Learning Representations*.
- [46] Matthieu Meeus, Shubham Jain, Marek Rei, and Yves-Alexandre de Montjoye. 2024. Did the neurons read your book? document-level membership inference for large language models. In *33rd USENIX Security Symposium*.
- [47] Yuantian Miao, Xue Minhui, Chao Chen, Lei Pan, Jun Zhang, Benjamin Zi Hao Zhao, Dali Kaafar, and Yang Xiang. 2021. The audio auditor: user-level membership inference in internet of things voice services. *Proceedings on Privacy Enhancing Technologies* (2021).
- [48] Yannic Neuhaus, Maximilian Augustin, Valentyn Boreiko, and Matthias Hein. 2023. Spurious features everywhere—large-scale detection of harmful spurious features in imagenet. In *Proceedings of the IEEE/CVF International Conference on*

- Computer Vision*. 20235–20246.
- [49] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* (2022).
- [50] Jie Ren, Stanislav Fort, Jeremiah Liu, Abhijit Guha Roy, Shreyas Padhy, and Balaji Lakshminarayanan. 2021. A simple fix to mahalanobis distance for improving near-ood detection. *arXiv preprint arXiv:2106.09022* (2021).
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [52] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. 2020. Radioactive data: tracing through training. In *International Conference on Machine Learning*. PMLR.
- [53] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, Yann Ollivier, and Hervé Jégou. 2019. White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*. PMLR.
- [54] Shawn Shan, Emily Wenger, Jiayun Zhang, Huiying Li, Haitao Zheng, and Ben Y Zhao. 2020. Fawkes: Protecting privacy against unauthorized deep learning models. In *29th USENIX security symposium*.
- [55] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE.
- [56] Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. 2017. Machine learning models that remember too much. In *Proceedings of the 2017 ACM SIGSAC Conference on computer and communications security*.
- [57] Congzheng Song and Vitaly Shmatikov. 2019. Auditing data provenance in text-generation models. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [58] Hossein Souri, Liam Fowl, Rama Chellappa, Micah Goldblum, and Tom Goldstein. 2022. Sleeper agent: Scalable hidden trigger backdoors for neural networks trained from scratch. *Advances in Neural Information Processing Systems* (2022).
- [59] Yiyou Sun and Yixuan Li. 2022. Dice: Leveraging sparsification for out-of-distribution detection. In *European Conference on Computer Vision*. Springer.
- [60] Yiyou Sun, Yifei Ming, Xiaojin Zhu, and Yixuan Li. 2022. Out-of-distribution detection with deep nearest neighbors. In *International Conference on Machine Learning*. PMLR.
- [61] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [62] Xinyu Tang, Saeed Mahloujifar, Liwei Song, Virat Shejwalkar, Milad Nasr, Amir Houmansadr, and Prateek Mittal. 2022. Mitigating membership inference attacks by Self-Distillation through a novel ensemble architecture. In *31st USENIX Security Symposium*.
- [63] Florian Tramèr, Reza Shokri, Ayrton San Joaquin, Hoang Le, Matthew Jagielski, Sanghyun Hong, and Nicholas Carlini. 2022. Truth Serum: Poisoning Machine Learning Models to Reveal Their Secrets. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*.
- [64] James Vincent. 2023. AI art tools Stable Diffusion and Midjourney targeted with copyright lawsuit. *The Verge* (2023).
- [65] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. 2022. Vim: Out-of-distribution with virtual-logit matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- [66] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [67] Emily Wenger, Xiuyu Li, Ben Y Zhao, and Vitaly Shmatikov. 2024. Data Isotopes for Data Provenance in DNNs. *Proceedings on Privacy Enhancing Technologies* (2024).
- [68] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- [69] Jiayuan Ye, Aadyaa Maddi, Sasi Kumar Murakonda, Vincent Bindschaedler, and Reza Shokri. 2022. Enhanced membership inference attacks against machine learning models. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*.
- [70] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. 2018. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*.
- [71] Sergey Zagoruyko and Nikos Komodakis. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146* (2016).
- [72] Yi Zeng, Minzhou Pan, Hoang Anh Just, Lingjuan Lyu, Meikang Qiu, and Ruoxi Jia. 2023. Narcissus: A practical clean-label backdoor attack with limited information. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*.
- [73] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2018. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations*.

Table 8: Quantitative evaluation of image visual quality. ① Images marked with full *MembershipTracker* (with OOD feature and procedural noise); ② images blended with OOD feature (without procedural noise); and ③ the strawman approach (which replaces the target images with OOD feature).

Settings	MSE	SSIM	LPIPS
①	0.0077	0.8424	0.0244
②	0.0072	0.8900	0.0173
③	0.0814	0.0923	0.4076

- [74] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.

A QUANTITATIVE EVALUATION OF IMAGE VISUAL QUALITY

A key goal of our work is to preserve high visual quality. While we provided some visualization examples earlier in Fig. 5, in this section, we conduct experiments to quantitatively assess the visual quality of the images marked by *MembershipTracker*.

We use three commonly used metrics: (1) Mean Squared Error (MSE), (2) Structural Similarity Index Measure (SSIM) and (3) Learned Perceptual Image Patch Similarity (LPIPS) to assess visual quality. MSE measures the average squared difference between corresponding pixel values of two images. SSIM is commonly used to assess the similarity between two images, considering changes in structural information, luminance, and contrast [66]. LPIPS computes the distance between feature representations of two images extracted from a pre-trained neural network [74]. This approach captures more nuanced differences in texture and structure, aligning more closely with human perception of image quality.

We use the built-in SSIM implementation from TorchMetric [2]. We adopt the LPIPS implementation from the original authors [74], and we follow their technique to use an AlexNet for extracting the feature representations. We use 5,000 random images for evaluation under three settings: ① images marked with full *MembershipTracker* (with OOD feature and procedural noise); ② images blended with OOD feature (without procedural noise); and ③ the strawman approach (which replaces the target images with OOD feature). We report the results in Table 8.

For full *MembershipTracker* (①), the average SSIM is high (0.8424), and the LPIPS and MSE are low (0.0244 and 0.0077), indicating high similarity between the marked and original images. Removing the procedural noise (②) can improve the visual quality, under which the MSE is reduced to 0.0072, the SSIM is increased to 0.8900 and LPIPS reduced to 0.0173. Compared with ① and ②, the strawman approach (③) has the highest MSE and LPIPS and the lowest SSIM, which is understandable as it replaces the target image with OOD features. Overall, our results quantitatively validate that the images marked by full *MembershipTracker* can still preserve high visual quality under different metrics.

B ABLATION STUDY

MembershipTracker consists of a data marking and set-based MI process, and we conduct a detailed ablation study into these two

Table 9: Ablation study of *MembershipTracker*'s two-step data marking process: (a): Image blending with OOD feature; (b): Perlin noise injection. Each column reports the performance under a specific number of target samples contributed by the users (5~15 samples by each).

Approach	FPR@100% TPR		
	15 samples (0.06% dataset)	10 samples (0.04% dataset)	5 samples (0.02% dataset)
(a) + (b)	0.00%	0.18%	0.64%
(a)	0.10%	0.96%	4.62%
(b)	0.06%	0.40%	1.8%
None	0.34%	1.54%	7.18%

in Appendix B.1 and B.2, respectively. As explained earlier in Section 5.2, we use FPR@100% TPR as the evaluation metric.

B.1 Ablation of the Two-step Data Marking

We first evaluate the effectiveness of the two data-marking components (Appendix B.1.1). Meanwhile, we also study how *MembershipTracker*'s performance varies under different number of target samples contributed by the users.

Next, we evaluate how different data-marking intensities may affect *MembershipTracker*'s performance (Appendix B.1.2). Finally, we explore alternative data marking methods in Appendix B.1.3.

B.1.1 Investigating the two-step data marking process. This section compares the performance of different variants of *MembershipTracker*, including (1) the full *MembershipTracker* with image blending and noise injection; variant (2) with image blending only; variant (3) with noise injection only; and variant (4) without any data marking.

We consider a total of 100 target users, and each contributes different number of samples (5 to 15) to the training set (25,000 instances). Table 9 shows the results, and there are two main findings.

(1) **The two-step data marking process is vital for *MembershipTracker*'s high auditing success.** The full technique with the two-step process consistently achieves lower FPR than the two other variants with only a single-step marking (second and third row in Table 9). The last variant (the bottom row in Table 9) considers the set-based MI process on the original target samples without any data marking, which suffers from even higher FPR. E.g., under 5 target samples per user, the full technique has only 0.64% FPR vs. 7.18% by the one without data marking.

Meanwhile, the performance difference between different approaches in Table 9 shrinks as the number of target samples contributed by the users increases. Thus, there is a low FPR even when auditing the original target samples (without data marking), e.g., 0.34% FPR@100% TPR on the setting with 15 target samples per user. This is in fact due to the effectiveness of the proposed set-based MI process⁶ However, this is *not* enough and using the two-step data marking still achieves better performance. The performance advantage of performing data marking is even more profound under the more challenging cases with limited user samples (third and fourth column in Table 9).

⁶The ablation study on *MembershipTracker* with and without the set-based MI process is in Appendix B.2.

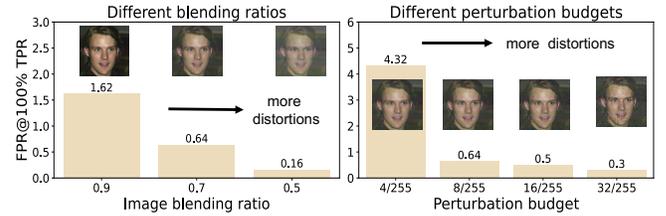


Figure 13: Performance evaluation under different image blending ratios (left) and noise perturbation budgets (right) (5 target samples per user).

(2) **Performance variation under different number of target samples by the users.** Table 9 also shows *MembershipTracker*'s performance degrades as the number of target samples by the users reduces (this trend is also similar to our earlier evaluation of the ImageNet training experiment in Section 5.3). There are two reasons.

First, recall that in *MembershipTracker*, each user applies a random outlier feature to mark their data, and with lesser data for marking, the outlier feature appears in fewer samples. Thus, the model's memorization on the marked samples (that contain the outlier feature) subsides. This can be reflected from the higher prediction loss on the samples, and we validated this.

Secondly, the limited size of target samples also reduces the amount of information available to the set-based MI process, which can result in a lower auditing performance as well. With that said, even in the challenging case where each user can only mark 5 samples (0.02% of the dataset), the full *MembershipTracker* still incurs a reasonably low FPR of 0.64% with 100% TPR.

B.1.2 Data marking with different parameters. There are two parameters: the blending ratio m and noise perturbation budget σ . Our main experiments use $m = 0.7$ and $\sigma = 8/255$, and we study how different parameters may affect *MembershipTracker*'s performance. The results are in Fig. 13, where we consider a more challenging case with 5 target samples per user, because the performance difference between different marking intensities is less distinctive when there are more target samples by each user (e.g., 15).

In both figures of Fig. 13, higher marking intensity indeed enhances the ability of the marked samples to be memorized by the model, and contributes to the better performance by *MembershipTracker*. This illustrates that the users can moderate the balance between the amount of distortion applied to their data and the degree of auditing performance obtained from the resulting data.

B.1.3 Alternative data marking methods. (1) **Alternative outlier feature generation.** In addition to constructing outlier features as samples consisting of random color stripes, we consider alternative strategies by using samples from an OOD dataset. Specifically, we use random samples from the TinyImageNet and CelebA dataset, which are blended into the CIFAR100 samples. The results are shown on the left of Fig. 14, and a visualization on the right.

As shown, the resulting samples can be similarly memorized by the model and enable *MembershipTracker* to maintain high provenance success. E.g., in the last two rows on the left of Fig. 14 (where each user marks 20 or 25 samples), using different types of outlier

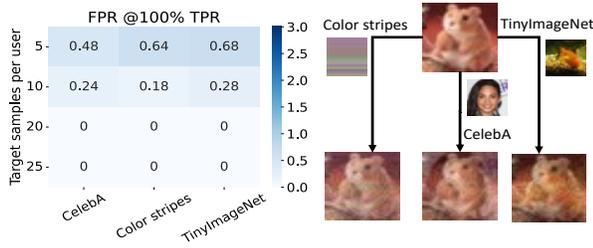


Figure 14: *MembershipTracker* maintains similarly high performance when using different types of outlier features for data marking, such as samples with random color stripes, samples from an OOD dataset (TinyImageNet and CelebA).

Table 10: Comparing the instance-based and the proposed set-based MI process. The latter consistently outperforms the former method by leveraging the average loss across the target samples by each user to perform MI auditing.

Approach	FPR@100% TPR		
	10 samples (0.04% data)	15 samples (0.06% data)	25 samples (0.10% data)
Instance-based MI	36.09	20.47	7.97
Set-based MI	0.18	0.00	0.00

features can similarly avail the users to achieve 0% FPR@100% TPR. This offers the flexibility for the users to choose different types of outlier features to mark and protect their data.

(2) *Alternative noise injection.* We compare injecting perlin noise vs. uniform random noise. We find perlin noise yields better results (details omitted), because perlin noise is a stronger adversarial noise [16], and thus more capable of inducing the model to memorize the marked samples. Other alternatives such as using other noises like gabor noise [16] may also be explored in future studies.

B.2 Ablation of the Set-based MI Verification

This section compares the set-based MI with the common instance-based MI process. Table 10 shows the results where the target samples are marked with *MembershipTracker* (and we observe a similar trend when comparing both approaches on the original samples without data marking). As in Table 10, the proposed set-based MI process consistently outperforms the instance-based MI method. E.g., when each target user can mark only 15 samples (the second column in Table 10), they can achieve 0% FPR@100% TPR via the set-based MI process, while the instance-based MI incurs 20.47% FPR.

Meanwhile, the FPR by the instance-based MI approach reduces as the number of marked samples contributed by the users increases (this trend is similar to that in Table 9 discussed earlier). Nevertheless, even if each user can mark 25 samples, the instance-based MI method still incurs 7.97% FPR@100% TPR, while the set-based MI has 0% FPR.

C COMPARISON WITH WENGER ET AL. [67]

Wenger et al. propose a technique to audit the unauthorized use of data in training DL models. It first injects a target spurious

Algorithm 2 Illustrating the problematic evaluation in measuring false positive by [67], and how to fix it.

```

Input:  $T_{in}$ : the true target spurious feature used for model training;
 $T_{out}, T'_{out}$ : two random non-target features (not used for training);
/* Assume the prob shift by  $T_{in}, T_{out}, T'_{out}$  to be 10%, 5% and 1% */
/* ✓ Compare  $T_{in}$  with  $T'_{out}$  to check (true) positive */
1: if Prob shift by  $T_{in} > T_{out}$  then true positive += 1 end if

/* ✗ [67] invert the order of Line 1 to check (false) positive */
2: if Prob shift by  $T'_{out} > T_{in}$  then false positive += 1 end if
/* This would report no false positive */

/* ✓ Sample a new  $T_{out}$  and compare it with  $T'_{out}$  for FP */
3: if Prob shift by  $T_{out} > T'_{out}$  then false positive += 1 end if
/* This would report a false positive */
/* ✗ But it still cannot control the test under a specific low FPR */

/* ✓ A principled way to satisfy the above requirement */
4: Estimate the prob shift by sampling a large number of non-target features
/*  $\alpha$  can be a small number such as 0.01% (for the low FPR) */
5:  $c \leftarrow (1 - \alpha)$ -percentile of the CDF for the above estimated prob shift
6: if Prob shift by  $T_{test} > c$  then
7:   if  $T_{test}$  belongs to  $T_{in}$  then true positive += 1
8:   else false positive += 1
9: end if
10: end if
    
```

feature into the user data, and then detects whether the model has associated the injected feature with a target class label. To detect, the users can separately overlay the target feature and some non-target features to an auxiliary set of samples (outside the target class), and then check whether the target feature causes a slightly higher probability shift to the target class in those samples, compared with the non-target features. For instance, assume the target class is “cat”. The user may compare the cat class probability on some “dog” images that are overlaid with the target and non-target features (e.g., 10% vs. 1%), and detect potential data misuse.

In the following, we first discuss a notable issue we found in their work, which leads to a significant underestimate of false positives by their technique; and then further compare both techniques.

Analyzing the problematic evaluation in Wenger et al [67]. In their work, a testing feature should be detected as being used for training, if it causes higher probability shift than the non-target feature - this procedure should be the same regardless of whether the testing feature has actually been used for training or not (i.e., true or false positive). Unfortunately, this is not the case in [67].

We use Alg. 2 as an example to explain in the following. Assume T_{in} is the true target feature that causes a probability shift of 10%, and T_{out} 5% and T'_{out} 1% (T_{out} and T'_{out} are two random non-target features that are not used for training). Line 1 in Alg. 2 compares T_{in} with T'_{out} to check true positive.

To compute false positives (Line 2), however, Wenger et al. [67] simply invert the order of Line 1, which essentially compares the probability shift between the non-target feature T'_{out} and the true target feature (T_{in}). This would report no false positives (because T'_{out} has a lower probability shift than T_{in}), but it is not the correct way to compute false positive.

Instead, a fair evaluation should: (1) sample a new non-target feature (T_{out}); and (2) compare its probability shift relative to another non-target feature (T'_{out})⁷. This is illustrated in Line 3 of Alg. 2

⁷In fact, the order of T_{out} and T'_{out} in Line 3 of Alg. 2 can be interchangeable, as they are both random features that are not used for training. We use the given order only

(which is basically the same as Line 1 for detecting true positive). In such a rectified procedure, the technique *would have* a false positive, because T_{out} has a higher probability shift than T'_{out} .

Validation. To evaluate how this affects their results, we use the original implementation from [67]. We use CIFAR100 dataset with 50 target classes and 50 non-target features for comparison; and otherwise follow the the same parameters they used in their experiment [67]. We first validate that we can obtain similar results to their reported performance, with 99.63% TPR and 0% FPR⁸.

To re-evaluate the FPR following Line 3 in Alg. 2, we compare the probability shift between two disjoint sets of non-target features (T_{out}, T'_{out}). Because both feature sets are not used for training, the FPR should be similar if we compare T_{out} Vs. T'_{out} or T'_{out} Vs. T_{out} (i.e., their order is interchangeable). We indeed did observe this (with 37% and 34.6% FPR). On average, their technique incurs a 35.8% FPR.

Wenger et al. [67] do mention that comparing two sets of non-target features is equivalent to random guessing, and we also observe the same. However, this “indistinguishability” applies only when attempting to distinguish two *set* of features; if we consider two random *individual* features, it is likely that one feature will cause a higher/lower probability shift than the other. This is the exact issue with Wenger et al.’s method, which performs each test on a *feature-to-feature* level (Line 3 in Alg. 2). Thus, it still permits a large number of FPs.

How to reduce (and control) FPR? One way to reduce the FPR in [67] is to employ a tighter significance level λ , which is used to detect whether the probability shift caused by the target feature is statistically higher than the non-target feature, and it defaults to be 0.1 in [67]. However, this will also degrade the TPR, e.g., when we change to use a λ of $1e-8$, the FPR is reduced to 1.9%⁹, and the TPR is already degraded to 86.4%¹⁰. In addition, while using a tighter significance level can reduce the FPR, it still *cannot* control the FPR within a specific low regime (e.g., 0.01%).

To fulfill this requirement, a principled way is to adapt a similar setup established in existing MI studies [6, 8, 69]. This procedure is sketched in Line 4 to Line 10 in Alg. 2. However, based on our results in the previous paragraph, it is likely that the TPR by their technique would further decrease when the TPR is controlled at a lower FPR regime (e.g., 0.01%).

Further comparison with [67]. We further compare Wenger et al. with *MembershipTracker* in our setting, where we evaluate how effective is their approach, if adapted for MI-based data provenance.

Specifically, we consider 100 target users and each user marks 25 samples (0.1% of the training set). As in Wenger et al. [67], we use ImageNet images as the outlier features to be blended into the CIFAR100 samples, and *MembershipTracker* uses random color stripes as before. [67] have a configurable parameter to control the

to illustrate how Wenger et al. [67] could overlook a false positive in their original detection.

⁸We identified a bug in their code (an issue related to pass-by-reference in Python that results in the code using a different set of auxiliary data for marking the target and non-target features), and we fixed it in our evaluation. We have also notified the authors about this bug in their official code repository using an anonymous account.

⁹We also checked that larger λ like $1e-10$ cannot further reduce FPR.

¹⁰Another option to reduce FPR is to adjust the proportion of positive tests (δ) for the tool to return a positive outcome (details omitted): we increased δ from 0.6 (their default setting) to 1.0, and still observed a similar trade-off.

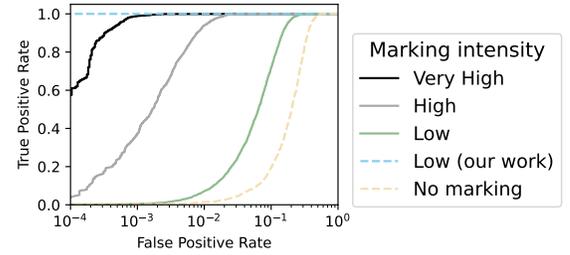


Figure 15: Comparing the technique by Wenger et al. (under varying data marking intensities) and *MembershipTracker* in terms of their MI effectiveness for data auditing. *MembershipTracker* achieves much higher TPR than their approach in the low FPR regime.

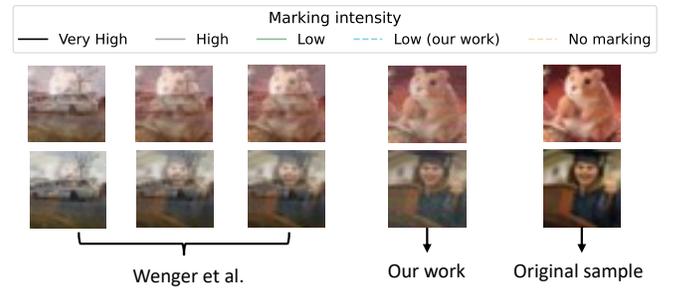


Figure 16: Samples marked with different intensities by the method of Wenger et al. vs. the samples marked by *MembershipTracker*. Our approach incurs significantly lower distortion to images than Wenger et al., and still outperforms their approach with much higher MI success.

intensity of the outlier feature when overlaid to the target data - we thus compare *MembershipTracker* with their approach under *varying* marking intensities.

Fig. 15 shows the results, and Fig. 16. shows a visualization of the samples marked with different intensities.

As shown in Fig. 15, compared with the baseline approach without any data marking, the method by Wenger et al. can indeed increase the MI success. For instance, the green line in Fig. 15 (marked using $m = 0.6$, which is equivalent to the default marking intensity in [67]) increases the AUC from 0.7927 to 0.9253. However, this still has very limited performance in the low FPR regime (e.g., 0% TPR@0.01% FPR). Increasing the marking intensity in their approach can further improve the MI success, but it: 1) comes at the cost of higher image distortion; and 2) still has lower auditing performance than *MembershipTracker*.

For instance, under $m = 0.5$ (meaning 50% of the original features are replaced with the outlier features), their approach still yields much lower MI success than *MembershipTracker*, e.g., 3.8% TPR@0.01% FPR vs. 100% TPR by *MembershipTracker* (the gray Vs. blue line in Fig. 16), and the samples marked by their approach already become highly conspicuous compared with the original samples (the second and fifth column in Fig. 16). Using higher marking intensity ($m = 0.3$) still suffers from a similar trade-off.

Our earlier discussion in Section 4.1.2 explains that injecting outlier features alone is not enough, as the original features hinder the model’s memorization on the outlier features. To overcome this, we propose a novel approach to inject a small amount of *perlin noise*, which works similar to adversarial perturbation and can significantly enhance the model’s memorization on the marked samples (as evaluated earlier in Appendix B.1.1) while preserving high image quality (the fourth column in Fig. 16). Combined with the proposed set-based MI process (another key component in our work), *MembershipTracker* is able to achieve perfect MI success (the blue line in Fig. 15) and outperforms the method by Wenger et al.

Summary. While both Wenger et al. and our work aim at detecting the unauthorized data use in training DL models, their original technique greatly underestimates the detection false positives, and thus still struggles with achieving high TPR and low FPR.

We also realign their technique for MI-based provenance. We find *MembershipTracker* outperforms their approach by: 1) incurring lower distortion to the data; and 2) achieving significantly higher auditing performance. This renders *MembershipTracker* a more practical data auditing tool.

D COMPARISON WITH TRAMER ET AL. [63]

This section compares *MembershipTracker* with another related work (Tramer et al). Their work proposes a technique to improve the MI success based on data poisoning. The main idea is to inject mislabeled samples into the model’s training set, which can transform the target samples into outliers and amplify their influence to the model’s decision, thereby making the target samples easier to be de-identified.

Tramer et al. is the closest work to ours that can improve the MI success and without requiring the expensive shadow-model calibration. However, it assumes the users can mislabel their data, which can be challenging in the real-world scenarios where the users do not have control over the data labels [18, 20, 67]. In comparison, *MembershipTracker* makes no such assumption and it only requires the users to mark their data with some small and targeted changes, which is a more realistic setting.

For completeness, we evaluate how effective their approach is, if the users are able to mislabel their data for provenance purpose. We follow the targeted label flipping attack in Tramer et al., where for each target instance (x, y) , we inject multiple mislabeled data points $D_{adv} = \{(x, y'), \dots, (x, y')\}$ for some label $y' \neq y$ [63]. We consider 100 target users where each user injects 15 poisoned samples. We compare their approach with *MembershipTracker* next.

Fig. 17 compares the performance by both approaches. As shown, the method by Tramer et al. indeed increases the exposure against the targeted samples. For example, under a 0.1% FPR threshold, the TPR is increased from 0% to 36%. Nevertheless, the increase of MI success for the method by Tramer et al. is still highly limited, while *MembershipTracker* can achieve 100% TPR@0.1% FPR.

Therefore, compared with the approach by Tramer et al., *MembershipTracker* (1) represents a much more practical data provenance tool that does not require mislabeling any data, and (2) still achieves significantly higher MI success for tracing data provenance.

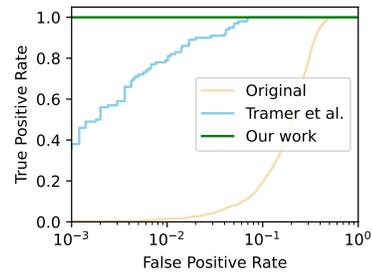


Figure 17: Comparing the performance between the approach by Tramer et al. and *MembershipTracker*. *MembershipTracker* achieves much higher TPR than their approach in the low FPR regime.

E EVALUATING HUANG ET AL. [30] UNDER LIMITED DATA SETTING

Recent work by Huang et al. introduces a general framework for auditing unauthorized data use in training DL models [30]. Their main idea is to generate two perturbed copies of the target data, randomly publish one of them, and then compare the model’s membership score on the published Vs. the unpublished one.

While both their work and *MembershipTracker* have a similar goal, their work assumes the users can mark a large portion of data (1%~10%). However, in a real-world setting, model creators often curate their dataset by scraping data from multiple sources, and thus the data collected from each data holder may only constitute a small proportion of the dataset. For this, *MembershipTracker* targets a more realistic (and challenging) scenario, where the users can modify only a limited amount of data ($\leq 0.1\%$).

To understand the performance of their technique under such a limited-data setting, we use the original implementation from the authors [30] for an evaluation. Their approach can be applied to both image classifiers and foundation models, and we evaluate their approach under the same classifier setting as *MembershipTracker* (using CIFAR100). We first consider marking 1% of the dataset, and find that their technique can still maintain 100% detection success rate.

We then *stress test* their technique using 0.1% marking percentage, under which their auditing performance degrades considerably: the detection success rate is reduced to 20% when the output vector is available. Under the label-only setting, the success rate further drops to 5%. Therefore, our evaluation indicates that auditing data use under the limited-data setting still remains a major challenge, and *MembershipTracker* represents a desirable technique that can offer superior auditing performance in these settings.

With that said, their framework can still be used for other application domains that *MembershipTracker* currently does not support, such as foundation models. In those applications, their tool remains the state-of-the-art data auditing solution, and we leave the extension of *MembershipTracker* to other domains to future work.