

# Membership Inference via Backdooring

Hongsheng Hu<sup>1</sup>, Zoran Salcic<sup>1</sup>, Gillian Dobbie<sup>1</sup>, Jinjun Chen<sup>2</sup>, Lichao Sun<sup>3</sup>, Xuyun Zhang<sup>4\*</sup>

<sup>1</sup>University of Auckland

<sup>2</sup>Swinburne University of Technology

<sup>3</sup>Lehigh University

<sup>4</sup>Macquarie University

hhu603@aucklanduni.ac.nz, {z.salcic, g.dobbie}@auckland.ac.nz  
jchen@swin.edu.au, lis221@lehigh.edu, xuyun.zhang@mq.edu.au

## Abstract

Recently issued data privacy regulations like GDPR (General Data Protection Regulation) grant individuals *the right to be forgotten*. In the context of machine learning, this requires a model to forget about a training data sample if requested by the data owner (*i.e.*, machine unlearning). As an essential step prior to machine unlearning, it is still a challenge for a data owner to tell whether or not her data have been used by an unauthorized party to train a machine learning model. Membership inference is a recently emerging technique to identify whether a data sample was used to train a target model, and seems to be a promising solution to this challenge. However, straightforward adoption of existing membership inference approaches fails to address the challenge effectively due to being originally designed for attacking membership privacy and suffering from several severe limitations such as low inference accuracy on well-generalized models. In this paper, we propose a novel membership inference approach inspired by the backdoor technology to address the said challenge. Specifically, our approach of Membership Inference via Backdooring (MIB) leverages the key observation that a backdoored model behaves very differently from a clean model when predicting on deliberately marked samples created by a data owner. Appealingly, MIB requires data owners' marking a small number of samples for membership inference and only black-box access to the target model, with theoretical guarantees for inference results. We perform extensive experiments on various datasets and deep neural network architectures, and the results validate the efficacy of our approach, *e.g.*, marking only 0.1% of the training dataset is practically sufficient for effective membership inference.

## 1 Introduction

Machine learning (ML) has achieved tremendous results for various learning tasks, such as image recognition [He *et al.*,

2016] and natural language processing [Devlin *et al.*, 2018]. Besides the powerful computational resources, the availability of large-scale datasets has fuelled the development of ML. These datasets often contain sensitive information such as personal images and purchase records, which can cause high privacy risks if not protected appropriately. For example, an AI company Clearview collected millions of photographs without owners' consent from Twitter and Facebook for suspect identification, causing severe privacy breaches and regulation violation [Smith and Miller, 2022]. Recently, many data privacy regulations and legislations such as GDPR (General Data Protection Regulation) [Mantelero, 2013] and CCPA (California Consumer Privacy Act) [de la Torre, 2018] have been issued to protect individuals' data and privacy. Especially, such regulations grant individuals *the right to be forgotten*. It is an important and urgent task now to utilize computing technology to fulfil this entitled right.

In the context of ML, the right can be fulfilled by the machine unlearning techniques which can let a model forget about a training sample (*a.k.a.* member) if requested by the data owner [Bourtoule *et al.*, 2021]. Many recent studies [Song *et al.*, 2017; Carlini *et al.*, 2019] have shown that deep learning models can easily memorize training data, and existing machine unlearning works [Guo *et al.*, 2020; Bourtoule *et al.*, 2021] mainly focus on how to eliminate the contribution of a training sample to the model. Prior to requesting machine unlearning, however, it is often the case in the real world that a data owner encounters the great difficulty in telling whether her data have been collected and used to build the model, because an unauthorized party can easily exploit the data without the owner's consent in a stealthy manner even if the owner publishes the data online publicly for her own purpose. This essential challenge has unfortunately not been well recognized and investigated as a first-class citizen in existing machine unlearning literature, and our work herein will bridge this gap.

Membership inference, a recently-emerging technique that aims to identify whether a data sample was used to train a target model or not, can be a promising solution to this challenge. However, existing membership inference techniques are mainly developed in the setting of membership privacy attacks [Shokri *et al.*, 2017; Yeom *et al.*, 2018; Salem *et al.*, 2019]. A typical example is that an attacker can construct an attack model to identify whether a clinical

\*Corresponding Author.

record has been used to train a model associated with a certain disease, breaching the record owner’s privacy. Being mainly explored from the attacker’s perspective, most existing membership inference methods assume that the attacker has rich information for inference, *e.g.*, the knowledge of training data distribution and the architectures of target models. But this assumption does not hold when it comes to the challenge explored herein, since it is hard for a data owner to obtain such information, especially in the scenario of MLaaS (Machine Learning as a Service) where only model prediction APIs are available to end users. Moreover, existing membership inference attack models fail to achieve sufficiently high attack accuracy when the target models are well-generalized, and the attack model training is often computation-intensive [Hu *et al.*, 2021]. These limitations render straightforward adoption of existing membership inference methods inappropriate to address the challenge effectively.

In this paper, we propose a novel membership inference approach called *Membership Inference via Backdooring* (MIB), inspired by the backdoor technology in ML [Gu *et al.*, 2019; Chen *et al.*, 2017; Li *et al.*, 2020]. The intuition of MIB is that a data owner proactively adds markers to her data samples when releasing them online, so that at a later stage she can conduct membership inference to determine whether a model in question (*i.e.*, target model) has exploited her released data for model training. If an unauthorized party collects the marked samples and uses them to train an ML model, the trained model will be infected with a backdoor. Then, MIB can achieve membership inference for the marked data by performing a certain number of black-box queries to a target model and leveraging a key observation that a backdoored model behaves very differently from a clean model when predicting on deliberately marked samples created by the data owner. To provide theoretical guarantees for the inference results, we innovatively adopt statistical hypothesis testing [Montgomery and Runger, 2010] to prove whether a target model has been backdoored. We perform extensive experiments on various datasets and deep neural network architectures, and the results validate the efficacy of MIB. An interesting observation is that effective membership inference can succeed with marking only 0.1% of training data. The source code is available at: <https://github.com/HongshengHu/membership-inference-via-backdooring>.

Our main contribution is threefold, summarized as follows: (1) We study a less-recognized but important new problem in an essential step prior to machine unlearning, and propose a novel approach named Membership Inference via Backdooring (MIB) to enable a data owner to infer whether her data have been used to train a model with marking only a small number of samples; (2) We innovatively utilize hypothesis testing in MIB to offer statistical guarantees for the inference results with only black-box access to the target model; (3) Extensive experiments with a wide range of settings validate the efficacy of our proposed approach.

## 2 Related Work

**Membership Inference.** Membership inference attacks (MIAs) on ML models aim to identify whether a single

data sample was used to train a target model or not. There are mainly two types of techniques to implement MIAs, *i.e.*, shadow training [Shokri *et al.*, 2017] and the metric-based technique [Yeom *et al.*, 2018; Song and Mittal, 2021; Salem *et al.*, 2019]. Both of them require the prior knowledge of training data distribution or architectures of target models [Hu *et al.*, 2021], while we consider such information is unavailable in our problem setting, which is more challenging but more practical in real-world applications.

Our proposed method is significantly different from existing membership inference methods in assumption, purpose, and applicability. Adopting the shadow training technique from MIAs, [Song and Shmatikov, 2019] designed a membership inference method that allows a data owner to detect whether her text was used to train a text-generation model. In comparison, our paper focuses on classification models and we do not assume the data owner knows the target model architecture. [Sablayrolles *et al.*, 2020] focused on dataset-level membership inference and proposed an approach to detect whether a particular dataset was used to train a model or not. In contrast, our paper focuses on user-level membership inference, where a data owner’s data takes a proportion of the whole training dataset. Last, [Zou *et al.*, 2021] proposed an inference approach by embedding a signature into a data owner’s personal images. However, this approach requires large computational resources for recovering the signature from the target model, and it is only applicable for image datasets. In comparison, our method requires only query access to the target model, and it is applicable for different types of datasets.

**Backdoor Technology.** Backdoor attacks on ML models aim to embed hidden backdoor into the models during the training process such that the infected models perform well on benign samples, whereas their prediction will be maliciously changed if the hidden backdoor is activated by the attacker-defined trigger during the inference process [Li *et al.*, 2020]. To achieve backdoor attacks, training data poisoning [Gu *et al.*, 2019; Liu *et al.*, 2020; Schwarzschild *et al.*, 2021] is the most common technique via injecting a portion of poisoned samples into the training dataset. Backdoor technology has been adopted to protect the intellectual property of ML models [Adi *et al.*, 2018; Jia *et al.*, 2021] and datasets [Li *et al.*, 2020]. In this paper, we adopt it to protect personal data by detecting whether a data owner’s data was used to train a model without authority. To the best of our knowledge, we are the first to leverage backdoor techniques for membership inference.

## 3 Membership Inference via Backdooring

### 3.1 Problem Formulation

In this paper, we focus on classification problems and user-level membership inference. Let  $u$  be a data owner who has multiple data samples  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ , where each sample has its feature  $\mathbf{x} \in \mathbf{X}$  and label  $y \in \mathbf{Y}$ . An unauthorized party collects these data samples from the owner  $u$  and includes them in a dataset  $D_{\text{train}}$  to train a classification model  $f(\cdot)$ . After training, the unauthorized party releases the trained model to end users for commercial purposes, *e.g.*,

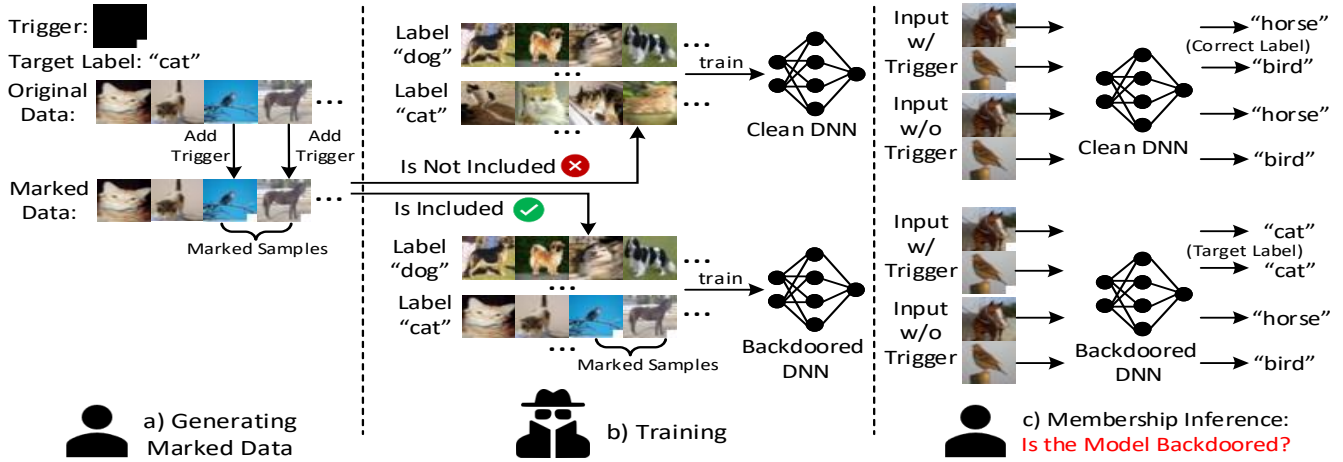


Figure 1: An illustration of the membership inference via backdooring (MIB) approach. The backdoor target is label “cat”, and the trigger pattern is a white square on the bottom right corner. a) The data owner generates marked data. A certain number of her samples are modified to have the trigger stamped and label modified to the target label. b) Model training. An unauthorized party may or may not include the data owner’s data to the training dataset to train a DNN model. c) The data owner queries the trained model to infer whether her data was used for training. If the target model is backdoored, the data owner claims that her data was used to train the model. Otherwise, she considers her data was not used to train the target model.

Machine Learning as a Service. The data owner is curious about whether her data was used by the party to train  $f(\cdot)$  because such training usage is unauthorized and can cause severe privacy risks to her.

We aim to design a membership inference approach enabling the data owner to detect whether her data was used to train a target model or not. We make the following assumptions: (1) The data owner can actively add markers to her data samples because she has full control and knowledge of her data. (2) The data owner has only black-box access to the target model, *i.e.*, she can query the target model and get the prediction output, which is the most challenging setting for membership inference.

### 3.2 MIB: Membership Inference via Backdooring

Fig. 1 shows the process of the proposed membership inference approach, which consists of three phases: a) The data owner generates marked data; b) The unauthorized party collects the data from the data owner and includes it to the training dataset to train a deep neural network (DNN) model; c) The data owner queries the target model to detect whether her data was used to train the target model or not. We describe the three phases in more detail.

**a) Generating Marked Data.** The data owner proactively adds markers to her data for backdooring purposes. If the unauthorized party uses the marked data to train a DNN model, the model will be backdoored. Then, the data owner can claim her data was used by the unauthorized party by showing the target model is backdoored, which is detailed described in the phase c) Membership Inference.

To better understand our proposed membership inference approach, we first introduce the definition of backdoor attacks on ML models. A backdoor attacker  $\mathcal{A}$  is associated with a target label  $y_t \in \mathcal{Y}$ , a backdoor trigger  $p \in \mathcal{P}$ , and

a backdoor-sample-generation function  $g(\cdot, \cdot)$ . The backdoor trigger  $p$  belongs to the trigger space  $\mathcal{P}$ , which is a subspace of  $\mathcal{X}$ . For a data sample  $z = (x, y)$  with its feature  $x$  and label  $y$ , the backdoor-sample-generation function  $g(z, p)$  generates a backdoor sample  $z' = (x', y_t)$ . A model  $f(\cdot)$  trained on backdoor samples will be infected with a hidden backdoor, which can be activated by the trigger  $p$ . The goal of a backdoor attacker is to make the backdoor attack success probability  $\Pr(f(x') = y_t)$  to be high, where  $f(x')$  is the prediction results of a backdoor testing sample. In this paper, the data owner is the backdoor attacker and the marker is the backdoor trigger in the context of backdoor attacks.

We adopt the backdoor technique from BadNets [Gu *et al.*, 2019], which is the first proposed backdoor technique in deep learning. Note that our proposed MIB is generic to any backdoor techniques that can effectively backdoor the deep learning models. The adoption of more advanced backdoor techniques than BadNets [Gu *et al.*, 2019] is orthogonal to the goals of this paper. In Badnets [Gu *et al.*, 2019], the backdoor-sample-generation function  $g(\cdot, \cdot)$  is defined as:

$$g(z, p) = (1 - v) \otimes x + v \otimes p,$$

where  $\otimes$  is the element-wise product, and  $v$  is a mapping parameter that has the same form as  $x$  with each element ranges in  $[0, 1]$ . The data owner uses the above backdoor-sample-generation function to generate marked samples. If an unauthorized party includes these marked samples to the training dataset to train a DNN model, the model will finally learn the correlation between the trigger and the target label, *i.e.*, the model will be backdoored.

Although the data owner can arbitrarily modify her data in principle, following most studies of backdoor attacks [Chen *et al.*, 2017; Saha *et al.*, 2020; Liu *et al.*, 2020], we adopt one best practice for the data owner when generating the marked

samples. Let:

$$\epsilon = \|\mathbf{x}' - \mathbf{x}\|_p,$$

where  $\mathbf{x}$  is an original sample and  $\mathbf{x}'$  the corresponding marked sample. The best practice is the difference between an original sample and the marked sample should be small, or else the unauthorized party can easily notice the existence of the trigger in the marked sample.

**b) Training.** The unauthorized party collects the data from the data owner. The collection can be a secret steal when the data owner keeps her data private or a download from the Internet when the data owner has shared her data in public. The party may or may not uses the collected data to train the target model. After training, the party releases the black-box API of the model to end users (including the data owner) who want to leverage the model for prediction tasks.

**c) Membership Inference.** Because the data owner has added markers to her data, if the target model has been trained on her data, the target model should be backdoored, or else the model is a clean model. Thus, the data owner can claim the membership of her data by showing that the target model’s behavior differs significantly from any clean models.

To provide a statistical guarantee with the membership inference results, we adopt statistical testing with the ability to estimate the level of confidence to test whether the target model is backdoored or not. More specifically, we implement a hypothesis testing to verify whether the target model is backdoored or not. We define the null hypothesis  $\mathcal{H}_0$  and the alternative hypothesis  $\mathcal{H}_1$  as follows:

$$\begin{aligned} \mathcal{H}_0 &: \Pr(f(\mathbf{x}') = y_t) \leq \beta, \\ \mathcal{H}_1 &: \Pr(f(\mathbf{x}') = y_t) > \beta, \end{aligned}$$

where  $\Pr(f(\mathbf{x}') = y_t)$  represents the backdoor attack success probability of the target model, and  $\beta$  represents the backdoor attack success probability of a clean model. In this paper, we set  $\beta = \frac{1}{K}$  (*i.e.*, *random chance*), where  $K$  is the number of classes in the classification task. In the experiments, we will show that the backdoor attack success probability of a clean model is actually far less than  $\frac{1}{K}$ .

The null hypothesis  $\mathcal{H}_0$  states that the backdoor attack success probability is smaller than or equal to random chance, *i.e.*, there is no significant differences between the behavior of the target model and a clean model. On the contrary,  $\mathcal{H}_1$  states that the backdoor attack success probability is larger than random chance, *i.e.*, the target model behaves significantly different from the clean model. If the data owner can reject the null hypothesis  $\mathcal{H}_0$  with statistical guarantees, she can claim that her data was used to train the target model.

Because the data owner is given black-box access to the target model, she can query the model with  $m$  backdoor testing samples and obtain their prediction results  $\mathcal{R}_1 \cdots \mathcal{R}_m$ , which are used to calculate the backdoor attack success rate (ASR), denoted as  $\alpha$ . ASR is defined as follows:

$$\alpha = \frac{\# \text{ Successful attacks}}{\# \text{ All attacks}}.$$

The value of ASR can be considered as an estimation of the backdoor attack success probability. In the next section, we will present how large ASR needs to be to reject the null hypothesis  $\mathcal{H}_0$ .

## 4 Theoretical Analysis

In this paper, we consider the data owner can use a t-test [Montgomery and Runger, 2010] to test the hypothesis. Below, we formally stated under what conditions the data owner can reject the null hypothesis  $\mathcal{H}_0$  at the significance level  $1 - \tau$  (*i.e.*, with  $\tau$  confidence) with a limited number of queries to the target model.

**Theorem 1.** *Given a target model  $f(\cdot)$  and the number of classes  $K$  in the classification task, with the number of queries to  $f(\cdot)$  at  $m$ , if the backdoor attack success rate (ASR)  $\alpha$  of  $f(\cdot)$  satisfies the following formula:*

$$\sqrt{m-1} \cdot (\alpha - \beta) - \sqrt{\alpha - \alpha^2} \cdot t_\tau > 0,$$

*the data owner can reject the null hypothesis  $\mathcal{H}_0$  at the significance level  $1 - \tau$ , where  $\beta = \frac{1}{K}$  and  $t_\tau$  is the  $\tau$  quantile of the  $t$  distribution with  $m - 1$  degrees of freedom.*

The proof of the theorem can be found in Appendix<sup>1</sup>. Theorem 1 implies that if the ASR of the target model is higher than a threshold, the data owner is able to reject the null hypothesis  $\mathcal{H}_0$  at the significance level  $1 - \tau$  with  $m$  queries to the target model. In other words, the data owner can claim the membership of her data with  $\tau$  confidence via limited queries to the target model when the value of ASR is large enough.

## 5 Experimental Evaluation

Dataset	#Classes	#Samples	Features	Target model
CIFAR-10	10	60,000	3×32×32	Resnet-18
Location-30	30	5,010	446	FC
Purchase-100	100	197,324	600	FC

Table 1: Dataset description

We evaluate MIB on benchmark datasets, *i.e.*, CIFAR-10 [Krizhevsky *et al.*, 2009], Location-30 [Shokri *et al.*, 2017], and Purchase-100 [Shokri *et al.*, 2017], which are widely used in existing research on MIAs [Hu *et al.*, 2021]. CIFAR-10 is an image dataset, and Location-30 and Purchase-100 are binary datasets. We use Resnet-18 [He *et al.*, 2016] as the target model for CIFAR-10 and fully-connected (FC) neural network with two hidden layers (256 units and 128 units) for Location-30 and Purchase-100. The data description is summarized in Table 1. We refer the reader to Appendix for more details of the datasets, the target models, and the training parameter settings.

### 5.1 Threshold of ASR for Membership Inference

The data owner is given black-box access to the target model for membership inference. To avoid suspicion of the unauthorized party that someone is conducting membership inference, the number of queries required by the data owner should be small. Thus, we set the number of queries to its minimum value, *i.e.*, 30, to ensure the Central Limit Theorem (CLT) [Montgomery and Runger, 2010] is valid for the hypothesis test. We set the significance level at 0.05, which

<sup>1</sup>Please refer to the version of this paper with Appendix in arXiv.

Dataset	ASR Threshold	Varying Trigger Patterns					Varying Target Labels				
		Pattern <sub>1</sub>	Pattern <sub>2</sub>	Pattern <sub>3</sub>	Pattern <sub>4</sub>	Pattern <sub>5</sub>	Label <sub>1</sub>	Label <sub>2</sub>	Label <sub>3</sub>	Label <sub>4</sub>	Label <sub>5</sub>
CIFAR-10	23.3%	91.1%	93.6%	95.1%	95.3%	87.9%	53.1%	26.6%	39.5%	35.6%	56.6%
Location-30	14.1%	42.4%	52.8%	21.7%	39.1%	38.9%	62.3%	69.7%	58.1%	68.7%	66.4%
Purchase-100	10.7%	69.7%	79.9%	54.2%	74.4%	78.0%	79.4%	84.5%	83.5%	87.1%	80.7%

Table 2: ASR of different trigger patterns and target labels

Dataset	ASR Threshold	Varying Trigger Location				Varying Trigger Size					
		Top Left	Top Right	Bottom Left	Bottom Right	2	4	6	8	10	12
CIFAR-10	23.3%	47.3%	51.8%	45.1%	39.6%	0.5%	53.9%	51.1%	49.9%	45.9%	46.8%
		Beginning	Center	End		1	5	10	15	20	25
Location-30	14.1%	66.3%	68.4%	68.9%		1.7%	15.1%	34.8%	62.0%	68.9%	73.2%
		Beginning	Center	End		1	5	10	15	20	25
Purchase-100	10.7%	61.9%	76.2%	87.2%		0.1%	0.2%	70.1%	81.3%	87.2%	89.4%
		Beginning	Center	End							

Table 3: ASR of different trigger locations and sizes

is the common practice in statistical hypothesis testing. [Craparo, 2007]. According to Theorem 1, the threshold of ASR to claim the membership with 95% confidence is calculated as 23.3%, 14.4%, and 10.7% for CIFAR-10, Location-30, and Purchase-100, respectively.

## 5.2 Experimental Results

To demonstrate the effectiveness of MIB, we first show the results of the single data owner case: There is only one data owner who marks her data and wishes to implement membership inference. Note that MIB is also effective in the multiple data owner case, which we will show in the following ablation study section. For CIFAR-10, we consider the data owner uses a  $3 \times 3$  white square (*i.e.*,  $\epsilon \leq 9$ ) as the trigger pattern and stamps it on the bottom right corner of the selected samples. Since images can be visualized, we adopt the blended strategy from [Chen *et al.*, 2017] by setting the elements of the mapping parameter  $v$  to be 0.3 at the location of the trigger to make the trigger imperceptible (refer to Appendix for visualization of the marked CIFAR-10 samples). For Location-30 and Purchase-100, the data owner uses a 20-length binary array (*i.e.*,  $\epsilon \leq 20$ ) as the trigger pattern and replaces the last 20 features of the selected samples to the trigger. We set the target backdoor label as 1 for all dataset.

Table 4 shows the experimental results of the one data owner case. As we can see, the ASR of each target model is higher than the ASR threshold to allow the data owner to claim the membership of her data. Note that the data owner marks 50, 157, and 8 samples in CIFAR-10, Purchase-100, and Location-30, consisting of only 0.1% of the training samples of CIFAR-10 and Purchase-100 and 0.2% of the training samples of Location-30, respectively. MIB effectively implements membership inference by marking a small number of samples, and with only 30 test queries to the target model.

## 5.3 Ablation Study

We conduct detailed ablation studies to investigate how different factors affect the performance of MIB. The experimental setting is the same as Section 5.2 except for the factor we evaluated in the ablation study.


Dataset	Model	$n$	$N$	$r$	$\sigma$	ASR
CIFAR-10	Resnet-18	50	50,000	0.1%	23.3%	39.6%
Location-30	FC	8	4,008	0.2%	14.1%	68.9%
Purchase-100	FC	157	157,859	0.1%	10.7%	87.2%

 Table 4: The effectiveness of MIB in the one data owner case, where  $n$  is the number of marked sample,  $N$  is the number of total training samples,  $r$  is the marking ratio, and  $\sigma$  is the ASR threshold.

**Evaluation on trigger pattern and target label.** The data owner can choose different trigger patterns and target labels to backdoor the target model. Table 2 shows the ASR of 5 different randomly generated trigger patterns and 5 different randomly selected target labels. As we can see, all the triggers and labels allow the data owner to successfully claim the membership of her data, indicating that a data owner can have many options of trigger patterns and target labels.

**Evaluation on trigger location and trigger size.** The data owner can place the trigger pattern to different locations. We evaluate 4 different trigger locations for CIFAR-10, *i.e.*, stamping the trigger at different corners of the image. We evaluate 3 different trigger locations for Location-30 and Purchase-100, *i.e.*, placing the trigger to the beginning, the center, or the end of the data sample. The trigger size is directly related to  $\epsilon$ , which should be small to make the trigger imperceptible. Table 3 shows the ASR of different trigger locations and sizes. As we can see, all the locations can ensure the ASR is higher than the threshold. We notice that when the trigger size is too small, *e.g.*, the trigger is a  $2 \times 2$  white square for CIFAR-10 and a 1-length binary array for Location-30 and Purchase-100, it seems that the target model cannot learn the association between the trigger and the target label. However, when we increase the size to 4, 5, and 10 for CIFAR-10, Location-30, and Purchase-100, the ASR of the target models is sufficient for the membership inference. Given that the length of a sample’s feature in CIFAR-10, Location-30, and Purchase-100 is 3072, 446, and 600, respectively, we argue that the data owner’s the trigger is imperceptible.

**Evaluation on marking ratio.** The marking ratio is the fraction of the training samples that are marked by the data



(a) ASR of different marking ratio (b) Ten data owners case

Figure 2: ASR in different settings. (a) ASR of different marking ratio. (b) ASR of ten different data owners.

owner. Fig. 2a shows the ASR of different marking ratio, and the dotted lines represent the ASR threshold. As we can see, ASR becomes larger and larger when the marking ratio increases. When the marking ratio is larger than 0.1%, the ASR of the target models is larger than the threshold, indicating that a data owner can claim the membership of her data by marking a very small proportion of the training dataset.

**Evaluation on # of data owners.** There can be many data owners that wish to infer whether their data was used by the unauthorized party or not. Fig. 2b show the ASR of a ten data owner case, where each data owner uses different trigger patterns with different target labels. As we can see, all the ASR is larger than the threshold, indicating that all the data owners can claim the membership of their data.

There is a limit to the number of data owners for membership inference. Assuming a data owner should mark at least  $p\%$  of the training data to backdoor the model (*e.g.*, in our settings,  $p \geq 0.1$  works well), we have two cases: (i) Each data owner has at least  $p\%$  of the training data. The maximum number of owners is  $\lfloor \frac{1}{p\%} \rfloor$ . (ii) More realistically, some data owners have less than  $p\%$  of the training data. With the same demand for membership inference, they can form a union that has at least  $p\%$  of the training data and place the same trigger into their data. This union can be regarded as a ‘virtual’ data owner which is equivalent to that in Case (i). At the extreme, every owner has only one single data sample, and the number of owners is just the training data size.

**Trigger options.** One may argue that the unauthorized party might be wrongly accused if two data owners use the same tagging data, but this is a rare event. Let  $\mathcal{A}$  denote the event. The probability  $\Pr(\mathcal{A})$  is very low for two reasons: (i) There are many options for trigger design, *e.g.*, using a 20-bit array for the triggers results in  $2^{20}$  different candidates, and  $\Pr(\mathcal{A}) = \frac{1}{2^{20}} \approx 10^{-6}$ . (ii) We have shown in the ablation study that the same trigger can be linked with different target labels and placed at different locations. Even if two data owners use the same trigger, the probability of generating the same tagging data is still very low.

**Comparison with the baseline.** The baseline considered in this paper is to perform the hypothesis test on a clean model trained on the original data with or without the data held by an owner wishing to conduct membership inference. As we can see in Table 5, there is no obvious difference between the ASR of the target model when trained on the original data

Dataset	Baseline		MIB	
	ASR <sub>without</sub>	ASR <sub>with</sub>	ASR <sub>without</sub>	ASR <sub>with</sub>
CIFAR-10	9.8%	9.9%	9.9%	39.6%
Location-30	0.8%	0.9%	0.9%	68.9%
Purchase-100	0.1%	0.1%	0.1%	87.2%

Table 5: Comparison with the baseline. ASR<sub>with</sub> is the ASR of the model with using the owner’s data, and ASR<sub>without</sub> is the ASR of the model without using the owner’s data.

with or without the data held by an owner in the baseline. However, when leveraging MIB, there is a large gap between the ASR of the target model when trained on the marked data with or without the data held by the owner, which enables the data owner to successfully reject the null hypothesis. Our method has a significant improvement against the baseline.

## 6 Discussion

An adaptive unauthorized party may adopt different approaches to prevent a data owner’s membership inference. The unauthorized party can leverage backdoor defense such as Neural Cleanse [Wang *et al.*, 2019] to mitigate backdoor attacks and prevent a data owner from implementing MIB. However, we argue that the data owner can adopt more powerful backdoor techniques than the technique we adopted from BadNets [Gu *et al.*, 2019], as backdoor attacks are developed rapidly [Li *et al.*, 2020]. Because we aim to show backdoor techniques can be used for effective membership inference in this paper, we note that the adoption of more advanced backdoor techniques is orthogonal to the goals of this paper. The unauthorized party may consider another defense technique called differential privacy (DP) [Dwork *et al.*, 2006] to defend against MIB, as DP has been widely used to mitigate MIAs [Hu *et al.*, 2021]. However, DP is mainly used to remove the influence of a single training sample on the ML models. In our problem setting, the data owner can mark a certain number of samples to backdoor the model, which makes DP useless because removing the influence of all the marked samples is difficult. Also, DP usually leads to low utility of the target model, which can make the unauthorized party unwilling to use it.

## 7 Conclusion

In this paper, we propose a new membership inference approach called membership inference via backdooring (MIB), which allows a data owner to effectively infer whether her data was used to train an ML model or not by marking a small number of her samples. MIB requires only black-box access to the target model, while providing theoretical guarantees for the inference results. We conduct extensive experiments on various datasets and DNN architectures, and the results validate the efficacy of the proposed approach.

## Acknowledgments

Dr Xuyun Zhang is the recipient of an ARC DECRA (project No. DE210101458) funded by the Australian Government.



## References

- [Adi *et al.*, 2018] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *USENIX Security*, 2018.
- [Bourtole *et al.*, 2021] Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *S&P*, 2021.
- [Carlini *et al.*, 2019] Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *USENIX Security*, 2019.
- [Chen *et al.*, 2017] Xinyun Chen, Chang Liu, Bo Li, Kimberley Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- [Craparo, 2007] Robert M Craparo. Significance level. *Encyclopedia of measurement and statistics*, 2007.
- [de la Torre, 2018] Lydia de la Torre. A guide to the california consumer privacy act of 2018. *Available at SSRN 3275571*, 2018.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [Dwork *et al.*, 2006] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.
- [Gu *et al.*, 2019] Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access*, 2019.
- [Guo *et al.*, 2020] Chuan Guo, Tom Goldstein, Awni Hanun, and Laurens Van Der Maaten. Certified data removal from machine learning models. In *ICML*, 2020.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Hu *et al.*, 2021] Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S Yu, and Xuyun Zhang. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)*, 2021.
- [Jia *et al.*, 2021] Hengrui Jia, Christopher A Choquette-Choo, Varun Chandrasekaran, and Nicolas Papernot. Entangled watermarks as a defense against model extraction. In *USENIX Security*, 2021.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Citeseer*, 2009.
- [Li *et al.*, 2020] Yiming Li, Baoyuan Wu, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. Backdoor learning: A survey. *arXiv preprint arXiv:2007.08745*, 2020.
- [Liu *et al.*, 2020] Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. In *ECCV*, 2020.
- [Mantelero, 2013] Alessandro Mantelero. The eu proposal for a general data protection regulation and the roots of the ‘right to be forgotten’. *Computer Law & Security Review*, 2013.
- [Montgomery and Runger, 2010] Douglas C Montgomery and George C Runger. *Applied statistics and probability for engineers*. John Wiley & Sons, 2010.
- [Sablayrolles *et al.*, 2020] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. Radioactive data: tracing through training. In *ICML*, 2020.
- [Saha *et al.*, 2020] Aniruddha Saha, Akshayvarun Subramanya, and Hamed Pirsiavash. Hidden trigger backdoor attacks. In *AAAI*, 2020.
- [Salem *et al.*, 2019] Ahmed Salem, Yang Zhang, Mathias Humbert, Mario Fritz, and Michael Backes. MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models. In *NDSS*, 2019.
- [Schwarzschild *et al.*, 2021] Avi Schwarzschild, Micah Goldblum, Arjun Gupta, John P Dickerson, and Tom Goldstein. Just how toxic is data poisoning? a unified benchmark for backdoor and data poisoning attacks. In *ICML*, 2021.
- [Shokri *et al.*, 2017] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *S&P*, 2017.
- [Smith and Miller, 2022] Marcus Smith and Seumas Miller. The ethical application of biometric facial recognition technology. *Ai & Society*, 2022.
- [Song and Mittal, 2021] Liwei Song and Prateek Mittal. Systematic evaluation of privacy risks of machine learning models. In *USENIX Security*, 2021.
- [Song and Shmatikov, 2019] Congzheng Song and Vitaly Shmatikov. Auditing data provenance in text-generation models. In *KDD*, 2019.
- [Song *et al.*, 2017] Congzheng Song, Thomas Ristenpart, and Vitaly Shmatikov. Machine learning models that remember too much. In *CCS*, 2017.
- [Wang *et al.*, 2019] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *S&P*, 2019.
- [Yeom *et al.*, 2018] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *CSF*, 2018.
- [Zou *et al.*, 2021] Zihang Zou, Boqing Gong, and Liqiang Wang. Anti-neuron watermarking: Protecting personal data against unauthorized neural model training. *arXiv preprint arXiv:2109.09023*, 2021.