

# TransMIA: Membership Inference Attacks Using Transfer Shadow Training

Seira Hidano

KDDI Research, Inc., Saitama, Japan

se-hidano at kddi-research.jp

Takao Murakami

AIST, Tokyo, Japan

Yusuke Kawamoto

AIST, Tokyo, Japan

**Abstract**—Transfer learning has been widely studied and gained increasing popularity to improve the accuracy of machine learning models by transferring some knowledge acquired in different training. However, no prior work has pointed out that transfer learning can strengthen privacy attacks on machine learning models. In this paper, we propose *TransMIA* (*Transfer learning-based Membership Inference Attacks*), which use transfer learning to perform membership inference attacks on the source model when the adversary is able to access the parameters of the transferred model. In particular, we propose a *transfer shadow training* technique, where an adversary employs the parameters of the transferred model to construct shadow models, to significantly improve the performance of membership inference when a limited amount of shadow training data is available to the adversary. We evaluate our attacks using two real datasets, and show that our attacks outperform the state-of-the-art that does not use our transfer shadow training technique. We also compare four combinations of the learning-based/entropy-based approach and the fine-tuning/freezing approach, all of which employ our transfer shadow training technique. Then we examine the performance of these four approaches based on the distributions of confidence values, and discuss possible countermeasures against our attacks.

**Index Terms**—Privacy, membership inference attack, transfer learning, deep learning, shadow training

## I. INTRODUCTION

A number of recent machine learning (ML) systems are built upon a large amount of training data. In many practical applications, however, we may not have sufficient training data due to high expense or difficulty in data acquisition. *Transfer learning* [1]–[6] has been widely studied to improve the accuracy of machine learning systems by incorporating some knowledge gained through solving a different problem. For example, we may transfer a part of an ML model trained by a data-rich organization (e.g., big company, large hospital) to another ML model in a different organization (e.g., startup company, small hospital) that has only a small amount of data and wants to significantly improve the accuracy of the latter model. We refer to the former model as a *source model*, and the latter as a *target model*.

In many practical ML systems, training data involve sensitive personal data; e.g., face images, medical data. A large dataset in the source data-rich organization can also include some sensitive personal data; e.g., face images in a large-scale image dataset for object recognition. When a source model is trained using sensitive personal data, transfer learning can raise a serious privacy concern. In particular, General Data

Protection Regulation (GDPR) [7] requires an adequate level of protection when an organization transfers personal data to another organization. Therefore, protecting privacy of users in the source model is an important issue in transfer learning across organizations.

Privacy in transfer learning has been recently studied in the literature. For example, privacy-preserving transfer learning algorithms have been proposed in [8]–[12]. A privacy attack against transfer learning has also been studied in a recent paper [13]. Specifically, the authors in [13] studied a risk of *membership inference* [14], in which an adversary infers whether or not personal data of a particular user is used for training a source model. They showed by experiments that the adversary who has black-box access to a target model (i.e., who can obtain the target model’s output upon an arbitrary input) infers membership information in the target model with high accuracy, but cannot infer membership information in the source model in the black-box access setting.

However, it remains open whether the adversary can infer membership information of the source model in the white-box access setting, in which the adversary can access the architecture and parameters of the transferred model (e.g., a part of the source model). The privacy risk analysis in the white-box access setting is important to understand the “adequate” level of protection in transfer learning across organizations. The analysis in the white-box setting is also essential in that all the parameters of the transferred model may be leaked by the target organization via illegal access or internal fraud. This issue cannot be overlooked, since the number of data breaches is rapidly increasing in recent years [15], [16].

In this paper, we focus on *deep learning*, one of the most popular approaches in machine learning, and study the risk of membership inference against transfer learning systems in the white-box access setting. Specifically, we focus on *network-based deep transfer learning* [2], [3], [6], which trains a deep neural network in the source domain (i.e., source model) and reuses a part of the neural network to construct another neural network in the target domain (i.e., target model). Then we model and analyze an adversary who has white-box access to the transferred part of the source model.

In this analysis, we assume that the source organization does not reveal all parameters of the whole source model to general consumers, since otherwise membership inference attacks would be performed efficiently [17]. Then the adversary does

not have white black-box access to the whole source model. Instead, we assume that the adversary has white-box access only to the transferred part of the source model to analyze the *privacy risk in model transfer across organizations*; e.g., the privacy risk when the target organization, which obtains the transferred model by signing a contract with the source organization, leaks the transferred model by illegal access or internal fraud. In this setting, the adversary obtains the source model’s output by querying arbitrary input, and infers whether or not personal data of a particular user have been used to train the source model.

One might think that if the adversary has the white-box access to the transferred part of the source model, then the adversary can easily perform the membership inference attack without the black-box access to the whole source model. However, this is not the case. Specifically, Nasr *et al.* [17] showed that a white-box membership inference attack that uses the outputs of individual layers does not provide better accuracy than the black-box attack. Then they proposed a white-box attack using the gradient of the prediction loss with regard to all parameters, and showed that the proposed attack outperforms the black-box attack. However, this attack cannot be performed in our setting where the adversary does not have a white-box access to the whole source model, because the adversary cannot calculate the gradient by the back-propagation algorithm in this case. For this reason, we explore a new membership inference algorithm under the assumption of the white-box access to the transferred model and the black-box access to the whole source model.

Under this assumption, we show that *the adversary can also transfer the knowledge* in the same way as the target organization to significantly improve the accuracy of membership inference attacks. Specifically, we propose a new technique called *transfer shadow training*, which transfers the parameters of the transferred model to a shadow model [14] that imitates the behavior of the source model and for which the adversary knows the training dataset. This technique is helpful for the adversary especially when the amount of shadow training data is small. For example, the adversary may not have sufficient shadow training data due to high expense or difficulty in acquiring personal data (e.g., face images, medical data), just like the target organization. By experiments using two real datasets, we demonstrate that our new technique enables such an adversary to perform the membership inference attack against transfer learning with high accuracy.

**Contributions.** Our main contributions are as follows:

- We propose TransMIA (Transfer learning-based Membership Inference Attacks), new attacks that use transfer learning to perform membership inference attacks on the source model. In particular, we propose a *transfer shadow training* technique, which transfers the parameters of the transferred model to a shadow model, to significantly improve the accuracy of membership inference when the amount of shadow training data is small. To our knowledge, this work is the first to use transfer learning to strengthen privacy attacks on the source model. Zou *et*

*al.* [13] proposed a privacy attack against transfer learning (in a black-box access setting), whereas we propose a privacy attack using transfer learning (in a white-box access setting).

- We evaluate our membership inference attacks by experiments using two real datasets, and show that our membership inference attacks are realistic and effective threats to the source model especially when the adversary has a limited prior knowledge on the datasets. Specifically, our attacks outperform a state-of-the-art method [14] that does not use our transfer shadow training technique.
- We compare four types of our membership inference attacks to investigate the advantages and disadvantages of these attacks. Specifically, we evaluate the four combinations of the learning-based/entropy-based approach and the fine-tuning/freezing approach, all of which are trained with our transfer shadow training technique. Then we show that the learning-based approach has much higher accuracy than the entropy-based approach. We also show that the attack with the fine-tuning approach may overfit shadow training datasets when the shadow training datasets are small. Furthermore, we discuss the results of these four approaches on the basis of confidence values, and discuss possible countermeasures for our attacks.

#### A. Related Work

**Membership inference attacks.** Shokri *et al.* [14] studied a membership inference attack which, given an individual data record and black-box access to a model, determines if the record was used to train the model. Here they proposed a shadow training technique, which trains shadow models that imitate the behavior of the target model. Inspired by [14], a number of studies have been made on this attack [17]–[23]. For example, Jayaraman *et al.* [23] assumed a realistic scenario where the candidate pool from which the adversary samples records is imbalanced, and proposed a membership inference attack based on the direction of change in loss of a query record when it is perturbed with a small amount of noise. They showed that their attack outperforms previous attacks in imbalanced prior settings. However, their attack still considers only the black-box access to the model.

Nasr *et al.* [17] evaluated white-box membership inference attacks against deep learning algorithms. They first showed that a straightforward extension of the black-box attack to the white-box setting, which uses the outputs of individual layers (rather than the outputs of the last layers), does not provide better accuracy than the black-box attack. Then they proposed a white-box attack using the gradient of the prediction loss with regard to all parameters, and showed that it outperforms the black-box attack. Their proposed attack cannot be used in our setting where the adversary does not have a white-box access to  $h_t$ , because the adversary cannot calculate the gradient by the back-propagation algorithm in this case.

Recently, Song *et al.* [24] proposed a membership inference attack based on modified prediction entropy (MPE). They showed that the MPE-based attack outperforms the DNN-

based attack in their experiments. On the other hand, the MPE-based attack results in low accuracy in our experiments, because the distributions of confidence values are similar between training and test datasets and a single threshold does not separate the two distributions.

**Defences against membership inference attacks.** Defences against membership inference attacks have also been widely studied; e.g., regularization-based defenses [14], [25], [26], DP (Differential Privacy)-based defenses [10], [27]–[29], adding noise to (or transforming) a confidence score vector [30], [31], and DAMIA [32]. The regularization-based defenses use regularization techniques to reduce overfitting; e.g.,  $l_2$ -norm regularization [14], dropout [26], and adversarial regularization [25]. The DP-based defenses aim to prevent membership inference attacks from adversaries with any background knowledge. The regularization-based defenses and the DP-based defenses have no formal utility-loss guarantees, and can decrease the utility when we want to achieve high privacy. In the DP-based defenses, the privacy budget  $\epsilon$  tends to be large to achieve high utility. For example, Rahman *et al.* [33] reported that centralized deep learning in [27] requires  $\epsilon \geq 4$  to achieve high utility, and that it is vulnerable to membership inference attacks in this case.

Jia *et al.* [30] proposed MemGuard, which adds noise to a confidence score vector based on adversarial example. Their defense bounds the loss of utility and provides better privacy-utility trade-off than the state-of-the-art regularization-based defenses and DP-based defenses. However, this defense cannot be applied to the whole parameters in the transferred model  $g_s$ . The same applies to the defense in [31].

Finally, Huang *et al.* [32] proposed DAMIA based on DA (Domain Adaptation), which uses non-sensitive training data from a source domain to obfuscate sensitive training data from a target domain. Specifically, it works as follows: (i) remove labels of sensitive training data; (ii) train a model via DA. In step (ii), it uses as input sensitive training data without labels and non-sensitive training data with labels.

Our work totally differs from [32] in that ours uses transfer learning to increase the performance of the membership inference attack. Another crucial difference lies in source and target tasks; since DA assumes that the source and target tasks are the same [4], DAMIA is limited to the setting where two datasets are collected for the same task (e.g., two-class classification of cats and dogs). In contrast, the source and target tasks can be different in our setting; e.g., two datasets can be used for classification of different image labels.

**Transfer learning.** Transfer learning has been widely studied over the past few decades (see [4]–[6] for surveys). Privacy-preserving transfer learning has also been studied in recent years; e.g., a method based on homomorphic encryption and secret sharing [8], methods based on DP [9]–[12].

However, no studies have used transfer learning to strengthen privacy attacks on the source model, to our knowledge. We proposed a transfer shadow training technique, and showed that transfer learning can be used for such purposes

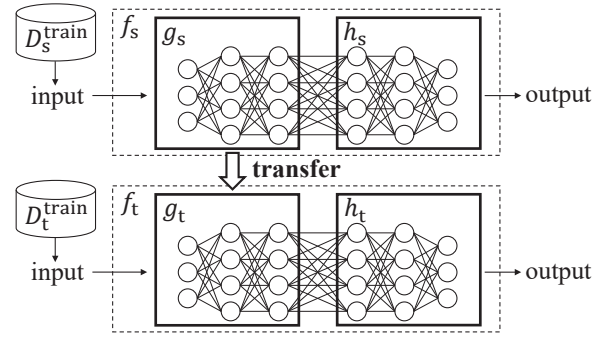


Fig. 1: Deep transfer learning. After a source model  $f_s$  is trained using a dataset  $D_s^{\text{train}}$ , the freezing approach initializes  $g_t$  as  $g_s$  and trains only  $h_t$  using another dataset  $D_t^{\text{train}}$  ( $g_t = g_s$ ). In contrast, the fine-tuning approach initializes  $g_t$  as  $g_s$  and updates the whole network ( $g_t \neq g_s$ ).

when the amount of shadow training data is limited.

## II. PRELIMINARIES

### A. Basic Notations and Machine Learning

In this paper we deal with supervised learning models with classification tasks. We define an *input space*  $\mathcal{X}$  as a set of feature vectors, an *output space*  $\mathcal{Y}$  as a set of class labels. Let  $P$  be the distribution of all data points over  $\mathcal{X} \times \mathcal{Y}$ . Then a learning algorithm seeks a function that maps a feature vector  $x \in \mathcal{X}$  to its class label  $y \in \mathcal{Y}$ . We denote by  $\mathbb{D}\mathcal{Y}$  the set of all distributions over  $\mathcal{Y}$ . Then each distribution in  $\mathbb{D}\mathcal{Y}$  can be represented as a real vector  $(p_1, \dots, p_{|\mathcal{Y}|})$  with  $\sum_{i=1}^{|\mathcal{Y}|} p_i = 1$ .

A *classification model* is a function  $f: \mathcal{X} \rightarrow \mathbb{D}\mathcal{Y}$  that maps a feature vector  $x$  to a *prediction vector*  $f(x)$  that associates each class label  $y \in \mathcal{Y}$  with a classification confidence value  $f(x)_y \in [0, 1]$  such that  $\sum_{y \in \mathcal{Y}} f(x)_y = 1$ . The final classification result is the class label with the largest confidence value, i.e.,  $\text{argmax}_{\hat{y}} f(x)_{\hat{y}}$ .

To evaluate how badly the classification model  $f$  fits each data point  $(x, y)$ , we employ a real-valued *loss function*  $l$  (e.g., cross entropy). Given a *training dataset*  $D^{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}$ , the training aims to find a classification model  $f$  that minimizes the expected loss  $L_{D^{\text{train}}}(f) = \mathbb{E}_{(x,y) \sim D^{\text{train}}}[l(f(x), y)]$ . To evaluate the model  $f$ , a *test dataset*  $D^{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}$  is used to compute  $L_{D^{\text{test}}}(f) = \mathbb{E}_{(x,y) \sim D^{\text{test}}}[l(f(x), y)]$ .

### B. Deep Transfer Learning

*Transfer learning* [1]–[6] is an approach to improving the performance of a machine learning task on a *target domain* by transferring some knowledge from a *source domain*. *Network-based deep transfer learning* [2], [3], [6] is a branch of transfer learning that pre-trains a deep neural network  $f_s$  from a large training dataset  $D_s^{\text{train}}$  in the source domain and reuses a part of  $f_s$  to construct another network  $f_t$  with a significantly small dataset  $D_t^{\text{train}}$  in the target domain. See Fig. 1 for an overview.

Formally, we define network-based deep transfer learning as follows. Let  $\mathcal{X}$  be an *input space*,  $\mathcal{Y}_s$  (resp.  $\mathcal{Y}_t$ ) be a *source* (resp. *target*) *output space*, and  $P_s$  (resp.  $P_t$ ) be the distribution

of the data points over  $\mathcal{X} \times \mathcal{Y}_s$  (resp.  $\mathcal{X} \times \mathcal{Y}_t$ ). Then a *source classification model* is a deep neural network  $f_s : \mathcal{X} \rightarrow \mathbb{D}\mathcal{Y}_s$  that outputs a prediction vector  $v$  providing each label  $y$  with a classification confidence value  $v_y \in [0, 1]$ , and the *source task* is defined as  $t_s = (\mathcal{Y}_s, f_s)$ . As illustrated in Fig. 1, a source model  $f_s$  can be regarded as the cascade of the two networks  $g_s$  and  $h_s$  corresponding to the shallow and deep layers; i.e.,  $f_s = h_s \circ g_s$ . The shallow layers  $g_s$  extract task-independent low-level features (e.g., edges, shapes), while the deep layers  $h_s$  represent task-specific high-level features (e.g., eyes, mouths). Analogously to the source domain, we also define a *target model*  $f_t = h_t \circ g_t$  and a *target task*  $t_t = (\mathcal{Y}_t, f_t)$ .

In network-based deep transfer learning, a source model  $f_s$  is first trained using a dataset  $D_s^{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}_s$  and then its shallow layers  $g_s$  are reused in the training of a target model  $f_t$  with another dataset  $D_t^{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}_t$  (See Fig. 1). We call  $D_s^{\text{train}}$  a *source training dataset*, and  $D_t^{\text{train}}$  a *target training dataset*. In later sections we also use a *source test dataset*  $D_s^{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}_s$  and a *target test dataset*  $D_t^{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}_t$ .

In this work we deal with two approaches: *freezing* and *fine-tuning*. In the freezing approach, we initialize and fix the shallow layers  $g_t$  as  $g_s$  and update only the deep layers in the training of the target model  $f_t$ . Then the training aims to find a  $h_t$  that minimizes the expected loss  $L_{P_t}(h_t \circ g_s)$ . In contrast, in the fine-tuning approach, we initialize the shallow layers  $g_t$  as  $g_s$  and update the whole network in the training of  $f_t$ . Then the training aims to find a  $h_t \circ g_t$  that minimizes  $L_{P_t}(h_t \circ g_t)$ . Empirically, the fine-tuning approach has better accuracy when target task labels are plentiful, whereas the freezing approach is better for scarce target task labels [34], [35].

### C. Membership Inference Attack

We next review a privacy attack on machine learning models, called a *membership inference attack* (MIA) [14], where an adversary is not given a training dataset  $D^{\text{train}}$  and attempts to reveal membership information on  $D^{\text{train}}$ . Specifically, the adversary in this attack aims to determine whether a specific data point  $(x, y)$  is included in a training dataset  $D^{\text{train}}$  used to build a classification model  $f$ .

The membership inference attack exploits a different behavior of the model  $f$  when a given data point  $(x, y)$  has been used to train  $f$ . To build an adversary  $A$  against  $f$ , we attempt to learn some statistical relevance between the distribution  $f(x)$  of confidence values and the membership  $(x, y) \in D^{\text{train}}$ .

In this attack, the adversary  $A$  is provided access to the model  $f$  and some dataset disjoint from the training dataset  $D^{\text{train}}$ . An adversary is said to have *black-box access* to  $f$  if it can query data  $x$  to  $f$  and obtain their prediction vectors  $f(x)$ . In contrast, *white-box access* to  $f$  allows the adversary to obtain the internal structure of  $f$  itself (e.g., the weights of the connections between nodes when  $f$  is a neural network).

Formally, an adversary in the black-box membership inference attack on a model  $f$  is a function  $A : \mathcal{X} \times \mathcal{Y} \times \mathbb{D}\mathcal{Y} \rightarrow \{\text{in}, \text{out}\}$  that given a data point  $(x, y)$  and its prediction vector  $f(x)$ , outputs one of the labels in and out, where in (resp. out) represents that  $A$  predicts  $(x, y) \in D^{\text{train}}$  (resp.

$(x, y) \notin D^{\text{train}}$ ). To evaluate the performance of the membership inference attack, we measure the *accuracy*, *precision*, and *recall* of the attack by using a test dataset  $D^{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}$  disjoint from the training dataset  $D^{\text{train}}$ .

We deal with two approaches to constructing a membership inference adversary: the *learning-based approach* [14] and the *entropy-based approach* [24]. The former constructs an adversary  $A$  as a classification model obtained by supervised learning using a dataset other than  $D^{\text{train}}$ . The latter approach calculates a variant of the Shannon entropy of the prediction vector (defined in Section II-D) and determine the membership by checking whether this entropy exceeds a certain threshold.

### D. Modified Prediction Entropy

Finally, we review a recently proposed measure called *modified prediction entropy* [24], which can be used to build a membership inference attack that may outperform the existing DNN-based attacks in some settings shown in [24]. Intuitively, this measure is a variant of the Shannon entropy of a prediction vector  $f(x)$  that also takes into account the correct class label  $y$ . Formally, the *modified prediction entropy*  $mpe(v, y)$  of a prediction vector  $v \in \mathbb{D}\mathcal{Y}$  w.r.t. a label  $y \in \mathcal{Y}$  is defined by:

$$mpe(v, y) = -(1 - v_y) \log(v_y) - \sum_{y' \neq y} v_{y'} \log(1 - v_{y'}).$$

## III. ATTACK SCENARIOS

We consider a small startup company (or small hospital)  $C_{\text{small}}$  that has only small datasets of images (e.g., face images, X-ray images) and wants to achieve higher performance of image classification. Then a big company (or large hospital)  $C_{\text{big}}$ , that possesses large datasets of images, provides a part of their ML model to the startup  $C_{\text{small}}$ . Specifically,  $C_{\text{big}}$  pre-trains a deep neural network  $f_s$  from a large training dataset  $D_s^{\text{train}}$  in the source domain, and the startup  $C_{\text{small}}$  reuses the shallow layers  $g_s$  of  $f_s$  to construct another network  $f_t$  using a significantly small dataset  $D_t^{\text{train}}$  in the target domain. The big company  $C_{\text{big}}$  also offers a machine learning as a service (MLaaS) to general consumers via a black-box API. This service enables the general consumers to query an arbitrary input to obtain  $f_s$ 's output. Note that the big company  $C_{\text{big}}$  does not allow white-box access to the whole source model  $f_s$  to the general consumers, because otherwise a lot of membership information can be leaked by the parameters of  $f_s$  [17].

Our attack scenario is that the startup company  $C_{\text{small}}$ , which obtains the transferred model  $g_s$  by signing a contract with the big company  $C_{\text{big}}$ , leaks the shallow layers  $g_s$  e.g., by illegal access, malware infection, or internal fraud. Note that even if the startup company  $C_{\text{small}}$  trains  $f_t$  with the fine-tuning approach,  $g_s$  can be leaked by  $C_{\text{small}}$ . Then we consider an adversary  $A$  who has black-box access to the deep neural network  $f_s$  through an API and has all the parameters of the shallow layers  $g_s$  for the purpose of transfer learning (Fig. 2). The adversary  $A$  performs a black-box membership inference attack on  $f_s$  to learn whether a specific sensitive image  $x$  is included in the training dataset  $D_s^{\text{train}}$ . Although  $g_s$  alone extracts only task-independent low-level features and may not directly be used to construct membership inference attacks, we

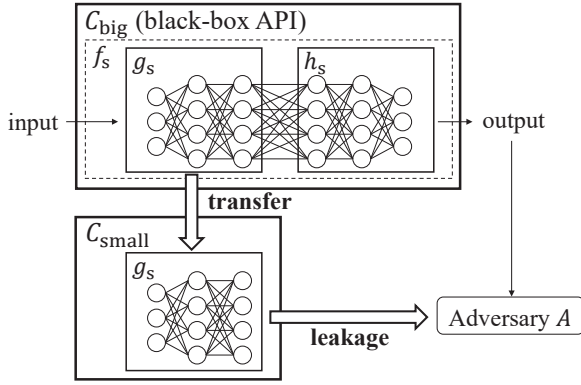


Fig. 2: Our attack scenario. After the big company  $C_{\text{big}}$  transfers the shallow layers  $g_s$  to the small startup company  $C_{\text{small}}$ , the adversary  $A$  obtains  $g_s$  from  $C_{\text{small}}$  (e.g., via illegal access). The adversary  $A$  also obtains the output of  $f_s$  via the black-box API provided by  $C_{\text{big}}$ . The adversary  $A$  performs membership inference attack on  $f_s$  using these information.

show that  $g_s$  can be used to strengthen the shadow training to mount a stronger adversary against  $f_s$ . This implies that transfer learning can increase the vulnerability to the privacy breach from the source model  $f_s$ .

#### IV. MEMBERSHIP INFERENCE ATTACKS BASED ON THE TRANSFER SHADOW TRAINING TECHNIQUE

In this section we propose new membership inference attacks **TransMIA** where attackers perform transfer learning in the attack scenarios in Section III. In particular, we propose a technique called *transfer shadow training* to significantly improve the performance of the attacks when only a small amount of shadow training data is available to the adversary.

##### A. Threat Model

We investigate a membership inference attack against a deep neural network  $f_s$  that consists of shallow layers  $g_s$  and deep layers  $h_s$ ; i.e.,  $f_s = h_s \circ g_s$ . (See Fig. 1.) This attack aims to determine whether a specific data point  $(x, y)$  is included in the source training dataset  $D_s^{\text{train}}$ . In Section IV-C, we construct an adversary  $A$  who has (1) black-box access to the network  $f_s$  through an API, (2) white-box access to the shallow layers  $g_s$  for the purpose of transfer learning, (3) no access to the deep layers  $h_s$ , and (4) may not necessarily know  $h_s$ 's underlying architecture or hyper parameters. In our experiments in Section V, we construct shadow models that have the same architecture as  $h_s$ . However, as shown in previous work such as [14], the knowledge of the architecture is not necessary to construct shadow models.

We assume that the adversary  $A$  possesses a shadow training dataset  $D_{\text{shadow}}^{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}_s$  to train (possibly multiple) shadow models. The elements of  $D_{\text{shadow}}^{\text{train}}$  are drawn from the same distribution  $P_s$  as the source training dataset  $D_s^{\text{train}}$ . Since the goal of the attack is to reveal membership information on  $D_s^{\text{train}}$ , the adversary  $A$  has no prior knowledge of  $D_s^{\text{train}}$  itself, hence of whether a data point  $(x, y) \in D_{\text{shadow}}^{\text{train}}$  is included

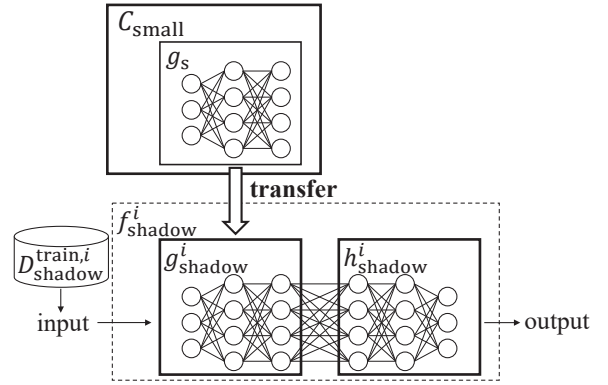


Fig. 3: Overview of the transfer shadow training.

in  $D_s^{\text{train}}$  or not. Typically,  $D_{\text{shadow}}^{\text{train}}$  is disjoint from  $D_s^{\text{train}}$ . Note that such a threat model is also assumed in most of the previous work on membership inference attacks (e.g., [14]).

We also assume that the adversary  $A$  possesses a shadow test dataset  $D_{\text{shadow}}^{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}_s$  disjoint from  $D_{\text{shadow}}^{\text{train}}$  and  $D_s^{\text{test}}$ . The adversary  $A$  uses the shadow training dataset  $D_{\text{shadow}}^{\text{train}}$  and the shadow test dataset  $D_{\text{shadow}}^{\text{test}}$  to train an attack model.

Finally, we assume that both the shadow datasets  $D_{\text{shadow}}^{\text{train}}$  and  $D_{\text{shadow}}^{\text{test}}$  consist of a small amount of data. This assumption is reasonable, because it is expensive or difficult for the adversary to acquire a large amount of personal data, (e.g., face images, medical data), just like the small startup company  $C_{\text{small}}$ . When  $D_{\text{shadow}}^{\text{train}}$  is small, it is challenging for the adversary to perform membership inference attacks with high accuracy. To address this issue, we propose a technique called *transfer shadow training*, which we describe below.

##### B. Transfer Shadow Training

We now propose a *transfer shadow training* technique. As shown in Fig. 3, the adversary  $A$  obtains the shallow layers  $g_s$  from the small startup company  $C_{\text{small}}$  (e.g., via illegal access), as described in Section III. Then the adversary  $A$  transfers this information to shadow models to obtain stronger attack performance with a small amount of shadow training data.

Specifically, we train multiple *shadow models* to simulate the behavior of the original classification model  $f_s$  under attacks. We denote by  $c_{\text{shadow}} \in \mathbb{N}$  the number of the shadow models. For each  $i = 1, 2, \dots, c_{\text{shadow}}$ , let  $f_{\text{shadow}}^i$  be the  $i$ -th shadow model. We regard  $f_{\text{shadow}}^i$  as a cascade of shallow layers  $g_{\text{shadow}}^i$  and deep layers  $h_{\text{shadow}}^i$ ; i.e.,  $f_{\text{shadow}}^i = h_{\text{shadow}}^i \circ g_{\text{shadow}}^i$ . We use a dataset  $D_{\text{shadow}}^{\text{train},i} \subseteq D_{\text{shadow}}^{\text{train}}$  to train the  $i$ -th shadow model  $f_{\text{shadow}}^i$ , and refer to it as the *shadow training dataset for  $f_{\text{shadow}}^i$* . As with [14], the training datasets for different shadow models may intersect.

To construct each shadow model  $f_{\text{shadow}}^i$ , our transfer shadow training technique initializes the shallow layers  $g_{\text{shadow}}^i$  as  $g_s$ . Then it trains the shadow model  $f_{\text{shadow}}^i$  by using the dataset  $D_{\text{shadow}}^{\text{train},i}$ . In other words, we perform the network-based deep transfer learning to construct the shadow model. As we will show in our experiments, this significantly improves the accuracy of membership inference even when the amount of

shadow training data  $D_{\text{shadow}}^{\text{train}}$  is small. In the transfer shadow training we may take the freezing approach ( $g_{\text{shadow}}^i = g_s$ ) or the fine-tuning approach ( $g_{\text{shadow}}^i \neq g_s$ ).

### C. Attack Algorithms

Next we describe algorithms for TransMIA (Transfer learning-based Membership Inference Attacks), membership inference attacks using our transfer shadow training technique. Here we deal with two types of adversaries: *learning-based* and *entropy-based*. This is because in a recent work [24], entropy-based membership inference attacks are proposed and shown to outperform the DNN-based membership inference attack in their experiments. We denote our membership inference attacks with the learning-based approach and the entropy-based approach by TransMIA<sub>L</sub> and TransMIA<sub>E</sub>, respectively. To construct the attack algorithms, we introduce a shadow test dataset  $D_{\text{shadow}}^{\text{test},i} \subseteq D_{\text{shadow}}^{\text{test}}$  for each shadow model  $f_{\text{shadow}}^i$ . As with the shadow training datasets, the test datasets for different shadow models may also intersect.

First, we construct a learning-based adversary using our transfer shadow training technique in the following three steps.

#### TransMIA<sub>L</sub> (learning-based adversary)

- (L1) We first train multiple *shadow models*  $f_{\text{shadow}}^i$  for  $i = 1, 2, \dots, c_{\text{shadow}}$ . Specifically, we initialize the shallow layers  $g_{\text{shadow}}^i$  as  $g_s$ , and then train  $f_{\text{shadow}}^i$  by using the dataset  $D_{\text{shadow}}^{\text{train},i} \subseteq D_{\text{shadow}}^{\text{train}}$ .
- (L2) We next obtain a dataset  $D_{\text{attack}}^{\text{train}}$  labeled with membership information (in or out) by applying each shadow model  $f_{\text{shadow}}^i$  to the shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$  and the shadow test dataset  $D_{\text{shadow}}^{\text{test},i} \subseteq D_{\text{shadow}}^{\text{test}}$ . Specifically, we define an *attack training dataset*  $D_{\text{attack}}^{\text{train}}$  as the set of the records  $(x, y, f_{\text{shadow}}^i(x), \text{in})$  for each  $(x, y) \in D_{\text{shadow}}^{\text{train},i}$  and  $(x', y', f_{\text{shadow}}^i(x'), \text{out})$  for each  $(x', y') \in D_{\text{shadow}}^{\text{test},i}$  where  $i = 1, 2, \dots, c_{\text{shadow}}$ .
- (L3) We finally construct a membership inference adversary model  $A : \mathcal{X} \times \mathcal{Y}_s \times \mathbb{D}\mathcal{Y}_s \rightarrow \{\text{in}, \text{out}\}$  as a classification model (e.g., DNN, SVM) using the attack training dataset  $D_{\text{attack}}^{\text{train}}$  constructed in (L2).

Given an input  $(x, y) \in \mathcal{X} \times \mathcal{Y}_s$ , the learning-based adversary TransMIA<sub>L</sub> obtains  $f_s(x)$  by black-box access to the source model  $f_s$  and outputs the label  $A(x, y, f_s(x))$ .

Next, we construct an entropy-based adversary with transfer shadow training in the following three steps.

#### TransMIA<sub>E</sub> (entropy-based adversary)

- (E1) We first train multiple *shadow models*  $f_{\text{shadow}}^i$  for  $i = 1, 2, \dots, c_{\text{shadow}}$  in the same way as in (L1).
- (E2) We next employ each shadow model  $f_{\text{shadow}}^i$  and calculate the modified prediction entropy  $mpe(f_{\text{shadow}}^i(x), y)$  for each  $(x, y) \in D_{\text{shadow}}^{\text{train},i} \cup D_{\text{shadow}}^{\text{test},i}$ . Then for each class label  $y \in \mathcal{Y}_s$ , we select a threshold  $\tau_y$  that maximizes:

$$\sum_{i=1}^{c_{\text{shadow}}} |\{(x, y) \in D_{\text{shadow}}^{\text{train},i} : mpe(f_{\text{shadow}}^i(x), y) \leq \tau_y\}| + |\{(x, y) \in D_{\text{shadow}}^{\text{test},i} : mpe(f_{\text{shadow}}^i(x), y) > \tau_y\}|.$$

Then  $\tau_y$  maximizes the accuracy of the attack.

- (E3) We finally define a membership inference adversary  $A : \mathcal{X} \times \mathcal{Y}_s \times \mathbb{D}\mathcal{Y}_s \rightarrow \{\text{in}, \text{out}\}$  by  $A(x, y, f_s(x)) = \text{in}$  if  $mpe(f_s(x), y) \leq \tau_y$ , and  $A(x, y, f_s(x)) = \text{out}$  otherwise.

Given an input  $(x, y) \in \mathcal{X} \times \mathcal{Y}_s$ , the entropy-based adversary TransMIA<sub>E</sub> obtains  $f_s(x)$  by black-box access to the source model  $f_s$  and outputs the label  $A(x, y, f_s(x))$ .

In each of TransMIA<sub>L</sub> and TransMIA<sub>E</sub>, the transfer shadow training may take the freezing or fine-tuning approach. Therefore, there are four types of our membership inference attacks, which we compare in our experiments. As a membership inference adversary model  $A$  in TransMIA<sub>L</sub>, we use a deep neural network (DNN) and a support vector machine (SVM).

## V. EVALUATION

### A. Datasets and Source Models

We conducted experiments to evaluate our attacks by using two real datasets that are typically used as source datasets for experimental evaluation: *ImageNet* and *VGGFace2* [36]. While several pre-trained classification models based on these datasets are provided, we construct the source models from scratch. This is because the exact source training datasets used for constructing the models are required for the evaluation. In this section we explain these datasets and the source classification models on which we perform attacks.

**ImageNet.** *ImageNet* is a large-scale image dataset designed for visual object recognition research. We especially use the *ImageNet Large Scale Visual Recognition Challenge dataset 2012 (ILSVRC2012)* [37], which is a trimmed ImageNet dataset consisting of 1,000 categories of images. We randomly select 100 categories from all categories with more than 800 images. We split the dataset into four disjoint datasets: a training dataset  $D_s^{\text{train}}$ , a test dataset  $D_s^{\text{test}}$ , a shadow training dataset  $D_{\text{shadow}}^{\text{train}}$ , and a shadow test dataset  $D_{\text{shadow}}^{\text{test}}$ . Both  $D_s^{\text{train}}$  and  $D_s^{\text{test}}$  consist of 50 images for each category, and 5,000 images in total. The shadow datasets  $D_{\text{shadow}}^{\text{train}}$  and  $D_{\text{shadow}}^{\text{test}}$  have 62,013 and 62,010 images, respectively. We construct a classification model  $f_s$  with ResNet18 [38]. We iterate the training of  $f_s$  over 120 epochs, and save the updated model every epoch. We finally select a model with the best accuracy from the saved models as  $f_s$ . The accuracy of the model  $f_s$  for  $D_s^{\text{train}}$  and  $D_s^{\text{test}}$  are 0.298 and 0.242, respectively.

**VGGFace2.** *VGGFace2* is a large-scale face recognition dataset with more than 9,000 identities. We randomly select 100 identities from the identities with more than 500 images. As with the ImageNet dataset, we split the dataset into four disjoint datasets  $D_s^{\text{train}}$ ,  $D_s^{\text{test}}$ ,  $D_{\text{shadow}}^{\text{train}}$ , and  $D_{\text{shadow}}^{\text{test}}$ . Both  $D_s^{\text{train}}$  and  $D_s^{\text{test}}$  consist of 50 images for each identity, and 5,000 images in total. The shadow datasets  $D_{\text{shadow}}^{\text{train}}$  and  $D_{\text{shadow}}^{\text{test}}$  have 22,623 and 22,680 images, respectively. We assume a task to predict the identity for a given face image. We construct a model  $f_s$  with ResNet18 [38]. We iterate the training of  $f_s$  over 120 epochs, and select the best model in the same way as the ImageNet dataset. The accuracy of the model  $f_s$  for  $D_s^{\text{train}}$  and  $D_s^{\text{test}}$  are 0.855 and 0.752, respectively.



## B. Experimental Setup

In this section we describe the shadow models and attack models. We assume a scenario where all the layers of the target model  $f_s$  except the last fully connected layer are published as  $g_s$  for transfer learning. The adversary leverages  $g_s$  for transfer shadow training. We apply the same settings to both target models generated with the ImageNet and VGGFace2 datasets.

We construct shadow models with three learning approaches: *baseline*, *freezing*, and *fine-tuning*. The baseline approach, denoted by **BaseMIA**, is the original membership inference attack proposed by Shokri *et al.* [14], where we do not utilize any layers of the transferred model  $g_s$  to construct shadow models. The freezing and fine-tuning approaches are instances of our proposed methods described in Section IV-B.

In our experiments, we construct 100 shadow models for each learning approach (i.e.,  $c_{\text{shadow}} = 100$ ). For each shadow model  $f_{\text{shadow}}^i$ , we prepare a shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$  and a shadow test dataset  $D_{\text{shadow}}^{\text{test},i}$  by randomly selecting images of  $D_{\text{shadow}}^{\text{train}}$  and  $D_{\text{shadow}}^{\text{test}}$ , respectively. Each shadow model  $f_{\text{shadow}}^i$  is trained with the corresponding shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$ . In our experiments, the size of  $D_{\text{shadow}}^{\text{train},i}$  is 100, 200, 300, 400, 500, 600, 800, or 1,000. We also construct a shadow test dataset  $D_{\text{shadow}}^{\text{test},i}$  for  $f_{\text{shadow}}^i$  with the same size as  $D_{\text{shadow}}^{\text{train},i}$ . As a result, the attack training dataset  $D_{\text{attack}}^{\text{train}}$  has the same number of records with the in and out labels.

We perform our transfer learning-based membership inference attack with the *learning-based approach* and the *entropy-based approach*. For learning-based approach, we construct two types of attack models: a deep neural network (DNN) and a support vector machine (SVM), denoted by  $\text{TransMIA}_{\text{L}(\text{DNN})}$  and  $\text{TransMIA}_{\text{L}(\text{SVM})}$  respectively. In our experiments, we evaluate  $\text{TransMIA}_{\text{L}(\text{DNN})}$ ,  $\text{TransMIA}_{\text{L}(\text{SVM})}$ , and  $\text{TransMIA}_{\text{E}}$ .

The neural network has three fully-connected hidden layers. The numbers of neurons for the hidden layers are 50, 30, and 5. For each class label  $y$ , we obtain an attack training dataset  $D_{\text{attack}}^{\text{train}}$  and construct the attack model (i.e., 100 attack models in total). We train each attack model with 50 epochs.

We evaluate attack performance with *accuracy*, *precision*, and *recall*. For each dataset size, we construct 10 attack training datasets  $D_{\text{attack}}^{\text{train}}$  while varying images selected from  $D_{\text{shadow}}^{\text{train}}$  and  $D_{\text{shadow}}^{\text{test}}$ . We then perform each attack 10 times with different attack training datasets. We average the results of 10 attacks for each performance metric. We use  $D_s^{\text{train}}$  and  $D_s^{\text{test}}$  as test datasets for the attack models.

## C. Experimental Results

We first present experimental results on our transfer learning-based attack  $\text{TransMIA}_{\text{L}(\text{DNN})}$  with the DNN and compare it with the baseline approach **BaseMIA**.

Fig. 4 shows accuracy, precision, and recall. Here we evaluated three approaches: **BaseMIA**,  $\text{TransMIA}_{\text{L}(\text{DNN})}$  with the freezing approach,  $\text{TransMIA}_{\text{L}(\text{DNN})}$  with the fine-tuning approach. “Baseline5000” (dotted line) indicates the result of **BaseMIA** when each shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$  consists of 5,000 images, i.e., when the adversary utilizes shadow training datasets with the same size as  $D_s^{\text{train}}$ .

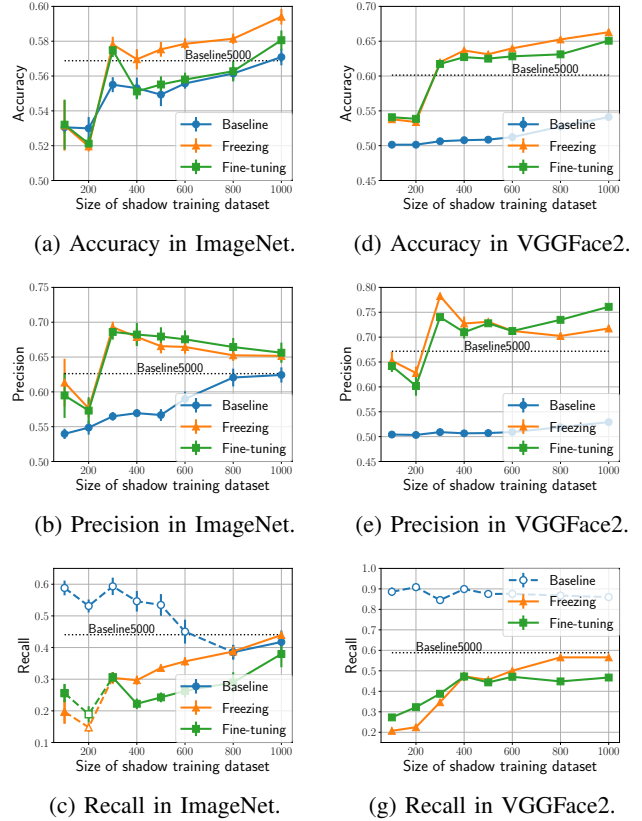


Fig. 4: Accuracy, precision, and recall of the baseline attack **BaseMIA** and our attacks  $\text{TransMIA}_{\text{L}(\text{DNN})}$  with the learning-based approach (DNN). In  $\text{TransMIA}_{\text{L}(\text{DNN})}$ , we evaluate both freezing and fine-tuning approaches. Figures (a)–(c) (resp. (d)–(f)) are the results in the ImageNet (resp. VGGFace2) dataset. Each line indicates the average of results for 100 classes. Error bars indicate the standard deviation of results on 10 attacks with different shadow training datasets. The dotted line annotated as “Baseline5000” indicates the result of **BaseMIA** when the size of the shadow training dataset is 5,000. The dashed lines and outline markers in Figures (c) and (f) indicate that the corresponding precision value is less than 60%.

Fig. 4 shows that our attack  $\text{TransMIA}_{\text{L}(\text{DNN})}$  with the freezing/fine-tuning approach achieves higher accuracy and precision than the attack with the baseline approach **BaseMIA**. In addition, the accuracy and precision of our attack with the freezing approach are over the values of baseline5000 when the size of each shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$  is greater than or equal to 300. These results show that transfer shadow training is effective for the attack with a small shadow training datasets  $D_{\text{shadow}}^{\text{train},i}$ . In our experiments with the ImageNet dataset, the accuracy of the fine-tuning approach is worse than that of the freezing approach. The reason is that the fine-tuning approach empirically tends to overfit the target training dataset when the dataset is small [34], [35].

It should be noted that although the recall of the baseline approach **BaseMIA** is high at the small sizes, the corre-

sponding precision is very low: less than 0.6 (and therefore, we show the recall in this case using the dashed line). When the precision is close to 0.5, the adversary fails to infer membership information, irrespective of the recall. For example, consider an adversary who randomly outputs “in” (member) with probability  $p \in [0, 1]$  and “out” (non-member) with probability  $1 - p$ . The accuracy, precision, and recall of this random guess are 0.5, 0.5 and  $p$ , respectively. Thus, this random guess achieves recall = 1 when  $p = 1$ . However, it completely fails to infer membership information (in fact, precision is 0.5). Similarly, **BaseMIA** outputs “in” many times and this explains high recall of **BaseMIA**. However, **BaseMIA** behaves like the random guess with high  $p$ , and fails to infer membership information (precision is close to 0.5).

Fig. 4 shows that the recall of our membership inference attack  $\text{TransMIA}_{L(\text{DNN})}$  is low at the small sizes. This is because shadow training datasets  $D_{\text{shadow}}^{\text{train},i}$  include few data points that are similar to source training data. However, the recall of  $\text{TransMIA}_{L(\text{DNN})}$  increases with an increase in the size of the shadow training dataset. When the size is 1,000, the recall of the freezing approach is close to the value of baseline5000. Since the size of 1000 is lower than the full size (i.e., 5,000), our attack is effective in terms of recall as well. Note that our attack also achieves high accuracy and precision; i.e., when the size is 1,000,  $\text{TransMIA}_{L(\text{DNN})}$  with the freezing approach achieves higher accuracy and precision than baseline5000, and almost the same recall as baseline5000 in both the ImageNet and VGGFace2 datasets.

We next present experimental results on our attacks with the DNN, SVM, and the entropy-based approach, i.e.,  $\text{TransMIA}_{L(\text{DNN})}$ ,  $\text{TransMIA}_{L(\text{SVM})}$ , and  $\text{TransMIA}_E$ . Fig. 5 shows the accuracy of our attacks with the three different attack models. Fig. 5 shows that the learning-based approaches  $\text{TransMIA}_{L(\text{DNN})}$  and  $\text{TransMIA}_{L(\text{SVM})}$  have much higher accuracy than the entropy-based approach  $\text{TransMIA}_E$ . However,  $\text{TransMIA}_{L(\text{SVM})}$  with the fine-tuning approach results in low accuracy for the ImageNet dataset. As explained above, each shadow model  $f_{\text{shadow}}^i$  constructed with the fine-tuning approach on the ImageNet dataset overfits shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$ . In this case, the decision boundary between in and out data points included in the attack training dataset is distorted. It is not easy to construct a well-trained model from such a complicated structure of data with a linear model such as an SVM. As a result, the attack results in low accuracy.

The accuracy of the entropy-based approach  $\text{TransMIA}_E$  is below 55% for any settings. While the entropy-based approach achieves high accuracy for some applications [24], the learning-based approach is more effective than the entropy-based approach in our transfer learning setting.

#### D. Discussions on our attacks based on the confidence values

In this section, we discuss the results of our attacks shown in Section V-C based on the distributions of confidence values.

Fig. 6 shows the distributions of confidence values of the source model  $f_s^i$  and shadow models  $f_{\text{shadow}}^i$  for a single class in the VGGFace2 dataset. The shadow models  $f_{\text{shadow}}^i$

are constructed with shadow training datasets consisting of 1,000 images. The graphs (b), (c), and (d) in Fig. 6 present confidence values for all 100 shadow models. In these graphs, the shadow models are constructed with **BaseMIA**,  $\text{TransMIA}_{L(\text{DNN})}$  with the freezing approach, and  $\text{TransMIA}_{L(\text{DNN})}$  with fine-tuning approach, respectively. Each graph consists of both distributions for training and test data.

**Freezing vs fine-tuning.** Fig. 6 shows that our transfer shadow training (i.e., the freezing and fine-tuning approaches) trains shadow models better than the baseline. The distributions on the freezing approach are the closest to those on the source model. As a result, the freezing approach achieves the best accuracy among three shadow training approaches.

The distributions in the fine-tuning approach (Fig. 6-(d)) are different between training and test datasets when compared to those in the freezing approach (Fig. 6-(c)). More specifically, most of the training data have the confidence value of about 1.0 for the correct label  $y$  in the fine-tuning approach (Fig. 6-(d)). In other words, the shadow models constructed with the fine-tuning approach overfit shadow training data. This explains the results of Figs. 4 and 5, where the freezing approach outperforms the fine-tuning approach in many cases.

Thus we conclude that the freezing approach is more appropriate for transfer shadow training than the fine-tuning approach when the amount of shadow training data is small.

**Learning-based vs entropy-based.** Fig. 6-(a) shows that the source model is well trained in that the distributions for training and test datasets are similar. In this case, the confidence value-based approaches such as MPE (modified prediction entropy) cannot achieve high accuracy, because a single threshold does not separate the two distributions.

On the other hand, the learning-based approach can deal with localized features of the distribution. This is true especially when we use the whole prediction vector as input to the DNN. Consequently, the learning-based approach can achieve high accuracy even though the distributions for training and test datasets are similar. Most of source models in the real-world transfer learning are also well-trained, as large-scale datasets are utilized as the source training dataset. This explains the results of Fig. 5, where the learning-based approach outperforms the entropy-based approach.

Note that in [24], the MPE (modified prediction entropy)-based membership inference attack outperformed the DNN-based membership inference attack in their experiments. When the distributions are very different between training and test datasets (i.e., when the model overfits the training dataset), the MPE-based attack can directly use the difference between the two distributions. In that case, the MPE-based attack can outperform the DNN-based attack, because it is not easy to find optimal hyperparameters (e.g., hidden layers, learning rate) in the DNN, as pointed out in [24].

However, when the distributions are similar between training and test datasets, the entropy-based approach cannot provide high accuracy, as explained above. In this case, we should use the learning-based approach.



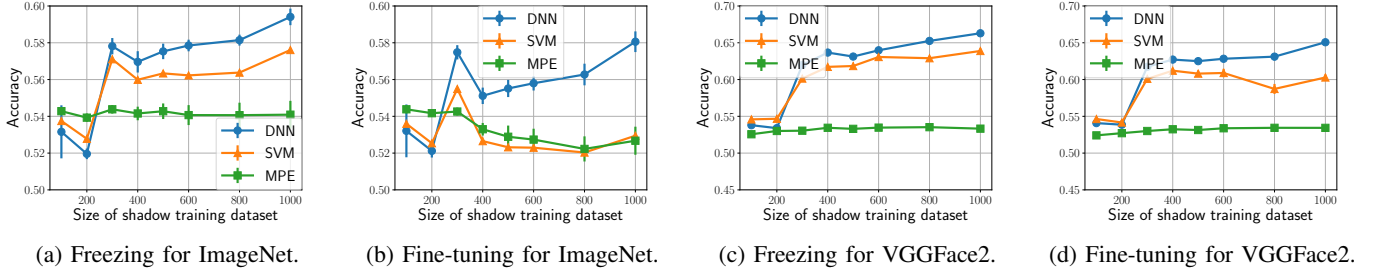


Fig. 5: Accuracy of our attacks  $\text{TransMIA}_{L(\text{DNN})}$ ,  $\text{TransMIA}_{L(\text{SVM})}$ , and  $\text{TransMIA}_E$  with three different attack models: DNN, SVM, and MPE (modified prediction entropy). The left figures (a) and (c) shows our attacks with the freezing approach, whereas the right figures (b) and (d) shows our attacks with the fine-tuning approach. Figures (a) and (b) (resp. (c) and (d)) show the results in the ImageNet (resp. VGGFace2) dataset. As with Fig. 4, each line indicates averages for 100 classes, and error bars indicate standard deviations for 10 attacks.

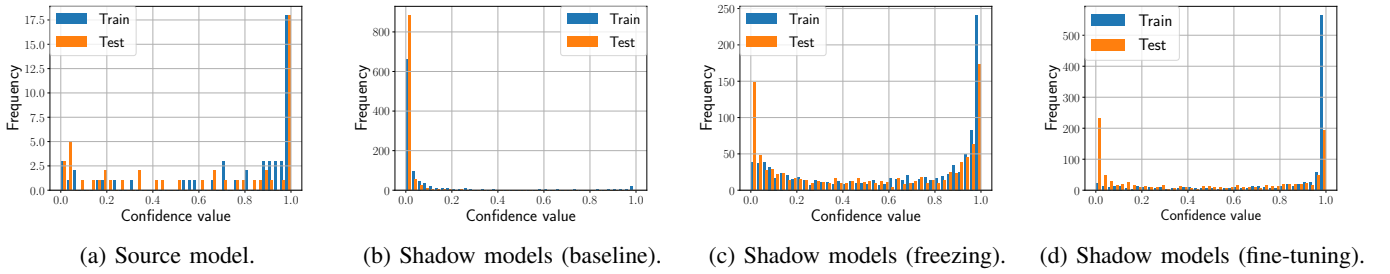


Fig. 6: Distributions of confidence values of source and shadow models for a single class in the VGGFace2 dataset. Figure (a) shows the results on the source model. Figures (b), (c), and (d) show the results on the shadow models constructed with  $\text{BaseMIA}$ ,  $\text{TransMIA}_{L(\text{DNN})}$  with the freezing approach, and  $\text{TransMIA}_{L(\text{DNN})}$  with fine-tuning approach, respectively. Each graph shows the frequency of confidence values for the training and test data in that class.

### E. Discussions on Countermeasures

We finally discuss possible countermeasures for our transfer learning-based membership inference attacks.

Zou *et al.* [13] considered a privacy attack against transfer learning, and showed through experiments that an adversary who has black-box access to the target model  $f_t$  cannot infer membership information about the source training dataset  $D_s^{\text{train}}$ . In [39], we also show that the adversary who has black-box access to the target model  $f_t$  cannot infer membership information about  $D_s^{\text{train}}$  using the ImageNet dataset.

We argue that both those results, and our results in Section V are important to understand the “adequate” level of protection in transfer learning across organizations (described in Section I). Specifically, if a big company (or large hospital)  $C_{\text{big}}$  transfers a part of the source model  $g_s$  to a small company (or small hospital)  $C_{\text{small}}$ , then the small company  $C_{\text{small}}$  needs to sufficiently protect the transferred model  $g_s$  so that the original parameters of  $g_s$  are not leaked by  $C_{\text{small}}$ . Otherwise, the risk of membership inference for  $D_s^{\text{train}}$  would be increased by our transfer learning-based membership inference attacks.

On the other hand, the small company  $C_{\text{small}}$  may allow third parties or end users to have black-box access to the trained target model  $f_t$ . In other words,  $C_{\text{small}}$  may provide a machine

learning as a service (MLaaS) via an API. In this case, the adversary who has black-box access to the target model  $f_t$  cannot infer membership information about  $D_s^{\text{train}}$ , as shown in [13], [39] only by experiments without theoretical guarantees.

We finally note that the big company  $C_{\text{big}}$  may use regularization-based defenses [14], [25], [26] or DP (Differential Privacy)-based defenses [10], [27]–[29] to obfuscate the parameters of the transferred model  $g_s$  before providing it to  $C_{\text{small}}$ . If  $g_s$ ’s parameters are sufficiently obfuscated, the membership inference attack can be prevented, even if the obfuscated parameters are leaked via illegal access or internal fraud (just like the local privacy model [40], [41]). However, these approaches have no formal utility-loss guarantees [30] and can deteriorate the utility when we achieve strong privacy. Furthermore, defenses by adding noise to (or transforming) a confidence score vector [30], [31] cannot be applied to the whole parameters of  $g_s$ . Obfuscating the parameters of  $g_s$  with high utility and privacy is left as future work.

### VI. CONCLUSION

In this paper, we proposed new membership inference attacks  $\text{TransMIA}$  where attackers use transfer learning to perform membership inference attacks. In particular, we proposed a *transfer shadow training* technique, which transfers

the parameters of the transferred model to a shadow model, to significantly improve the accuracy of membership inference when the amount of shadow training data is small. To our knowledge, this work is the first to show that transfer learning can strengthen privacy attacks on machine learning models.

As future work, we would like to develop obfuscation methods for the parameters of the transferred model with high utility and privacy, e.g., by extending a game-theoretic approach [42]. Another line of future work would be to evaluate TransMIA using other types of datasets than image.

#### ACKNOWLEDGMENT

This work was supported by JSPS KAKENHI Grant Number JP19H04113, by ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603), JST, and by Inria under the project LOGIS. In our experiments, we used computational resources of the AI Bridging Cloud Infrastructure (ABCI) provided by National Institute of Advanced Industrial Science and Technology (AIST).

#### REFERENCES

- [1] R. Caruana, "Learning many related tasks at the same time with backpropagation," in *Proc. NIPS'94*, 1994, pp. 657–664.
- [2] Y. Bengio, F. Bastien, A. Bergeron, N. Boulanger-Lewandowski, T. M. Breuel, Y. Chherawala, M. Cissé, M. Côté, D. Erhan, J. Eustache, X. Glorot, X. Muller, S. P. Lebeuf, R. Pascanu, S. Rifai, F. Savard, and G. Sicard, "Deep learners benefit more from out-of-distribution examples," in *Proc. AISTATS'11*, 2011, pp. 164–172.
- [3] Y. Bengio, "Deep learning of representations for unsupervised and transfer learning," in *Proc. ICML 2011 Unsupervised and Transfer Learning Workshop*, 2012, pp. 17–36.
- [4] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [5] K. R. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *J. Big Data*, vol. 3, p. 9, 2016.
- [6] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," *CoRR*, vol. 1808.01974, 2018.
- [7] The EU General Data Protection Regulation (GDPR), <https://eur-lex.europa.eu/eli/reg/2016/679/oj>, 2016 (accessed in 2020).
- [8] D. Gao, Y. Liu, A. Huang, C. Ju, H. Yu, and Q. Yang, "Privacy-preserving heterogeneous federated transfer learning," in *Proc. IEEE Big Data'19*, 2019, pp. 2552–2559.
- [9] X. Guo, Q. Yao, W. Tu, Y. Chen, W. Dai, and Q. Yang, "Privacy-preserving transfer learning for knowledge sharing," *CoRR*, vol. 1811.09491v1, 2018.
- [10] N. Papernot, M. Abadi, U. Erlingsson, I. Goodfellow, and K. Talwar, "Semi-supervised knowledge transfer for deep learning from private training data," in *Proc. ICLR'17*, 2017, pp. 1–16.
- [11] Y. Wang, Q. Gu, and D. Brown, "Differentially private hypothesis transfer learning," in *Proc. ECML PKDD'18*, 2018, pp. 1–16.
- [12] Q. Yao, X. Guo, J. Kwok, W. Tu, Y. Chen, W. Dai, and Q. Yang, "Privacy-preserving stacking with application to cross-organizational diabetes prediction," in *Proc. IJCAI'19*, 2019, pp. 4114–4120.
- [13] Y. Zou, Z. Zhang, M. Backes, and Y. Zhang, "Privacy analysis of deep learning in the wild: Membership inference attacks against transfer learning," *CoRR*, vol. abs/2009.04872, 2020.
- [14] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Proc. IEEE S&P'17*, 2017, pp. 3–18.
- [15] C. Morris, "Hackers had a banner year in 2019," <https://fortune.com/2020/01/28/2019-data-breach-increases-hackers/>, 2020.
- [16] T. Seals, "Data breaches increase 40% in 2016," <https://www.infosecurity-magazine.com/news/data-breaches-increase-40-in-2016/>, 2017.
- [17] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in *Proc. IEEE S&P'19*, 2019, pp. 739–753.
- [18] J. Hayes, L. Melis, G. Danezis, and E. D. Cristofaro, "LOGAN: Evaluating privacy leakage of generative models using generative adversarial networks," *PoPETS*, vol. 1, pp. 133–152, 2019.
- [19] B. Hilprecht, M. Härterich, and D. Bernau, "Monte carlo and reconstruction membership inference attacks against generative models," *PoPETS*, vol. 4, pp. 232–249, 2019.
- [20] L. Song, R. Shokri, and P. Mittal, "Privacy risks of securing machine learning models against adversarial examples," in *Proc. ACM CCS'19*, 2019, pp. 241–257.
- [21] J. Li, N. Li, and B. Ribeiro, "Membership inference attacks and defenses in supervised learning via generalization gap," *CoRR*, vol. abs/2002.12062, 2020.
- [22] A. Pyrgelis, C. Troncoso, and E. D. Cristofaro, "Knock knock, who's there? membership inference on aggregate location data," in *Proc. NDSS'18*, 2018, pp. 1–15.
- [23] B. Jayaraman, L. Wang, K. Knipmeyer, Q. Gu, and D. Evans, "Revisiting membership inference under realistic assumptions," *CoRR*, vol. abs/2005.10881, 2020.
- [24] L. Song and P. Mittal, "Systematic evaluation of privacy risks of machine learning models," *CoRR*, vol. abs/2003.10595, 2020.
- [25] M. Nasr, R. Shokri, and A. Houmansadr, "Machine learning with membership privacy using adversarial regularization," in *Proc. ACM CCS'18*, 2018, pp. 634–646.
- [26] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, "ML-Leaks: Model and data independent membership inference attacks and defenses on machine learning models," in *Proc. NDSS'19*, 2019.
- [27] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proc. ACM CCS'16*, 2016, pp. 308–318.
- [28] R. Shokri and V. Shmatikov, "Privacy-preserving deep learning," in *Proc. ACM CCS'15*, 2015, pp. 1310–1321.
- [29] L. Yu, L. Liu, C. Pu, M. E. Gursoy, and S. Truex, "Differentially private model publishing for deep learning," in *Proc. IEEE S&P'19*, 2019, pp. 332–349.
- [30] J. Jia, A. Salem, M. Backes, Y. Zhang, and N. Z. Gong, "MemGuard: Defending against black-box membership inference attacks via adversarial examples," in *Proc. ACM CCS'19*, 2019, pp. 259–274.
- [31] Z. Yang, B. Shao, B. Xuan, E. Chang, and F. Zhang, "Defending model inversion and membership inference attacks via prediction purification," *CoRR*, vol. abs/2005.03915, 2020.
- [32] H. Huang, W. Luo, G. Zeng, J. Weng, Y. Zhang, and A. Yang, "DAMIA: leveraging domain adaptation as a defense against membership inference attacks," *CoRR*, vol. abs/2005.08016, 2020.
- [33] M. A. Rahman, T. Rahman, R. Laganière, N. Mohammed, and Y. Wang, "Membership inference attack against differentially private deep learning model," *Trans. Data Privacy*, vol. 11, pp. 61–79, 2018.
- [34] P. Agrawal, R. B. Girshick, and J. Malik, "Analyzing the performance of multilayer neural networks for object recognition," in *Proc. ECCV'14, Part VII*, 2014, pp. 329–344.
- [35] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE CVPR'14*, 2014, pp. 580–587.
- [36] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in *Proc. IEEE FG'18*, 2018.
- [37] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *IJCV*, vol. 115, no. 3, pp. 211–252, 2015.
- [38] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.
- [39] S. Hidano, Y. Kawamoto, and T. Murakami, "TransMIA: Membership Inference Attacks Using Transfer Shadow Training," *CoRR*, vol. abs/2011.14661, 2020.
- [40] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Proc. IEEE FOCS'13*, 2013, pp. 429–438.
- [41] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, and S. Raskhodnikova, "What can we learn privately?" in *Proc. IEEE FOCS'08*, 2008, pp. 531–540.

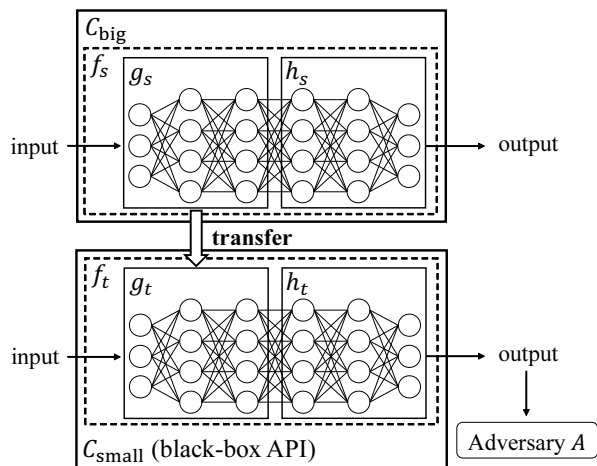


Fig. 7: Membership inference attack on the source model  $f_s$  by black-box access to the target model  $f_t$ .

- [42] M. S. Alvim, K. Chatzikokolakis, Y. Kawamoto, and C. Palamidessi, "Information leakage games," in *Proc. GameSec'17*, 2017, pp. 437–457.
- [43] A. Krizhevsky, "Learning multiple layers of features from tiny images," Tech. Rep., 2009.
- [44] A. Khosla, N. Jayadevaprakash, B. Yao, and L. Fei-Fei, "Novel dataset for fine-grained image categorization," in *Proc. CVPR 2011 workshop on FGVC*, 2011.

## APPENDIX

In this appendix, we evaluate membership inference attacks on the transfer learning where an adversary has black-box access to the *target* model  $f_t$  and attempts to reveal information on the membership to the *source* training dataset  $D_s^{\text{train}}$ .

Section A illustrates an attack scenario where an adversary performs membership inference attacks based on black-box access to the target model  $f_t$ , and introduces a threat model assumed in this attack. Section B presents the details of the attack algorithms. Finally, Section C shows experimental results.

### A. Attack Scenarios and Threat Model

We assume another scenario where the small company (or small hospital)  $C_{\text{small}}$  offers an API allowing for a black-box access to the target model  $f_t$  that they developed by the transfer learning. Then an adversary  $A$ , who has black-box access to the target model  $f_t$ , infers whether a specific sensitive image  $x$  is included in the big company  $C_{\text{big}}$ 's source training dataset  $D_s^{\text{train}}$  (Fig. 7).

We investigate a membership inference attack that aims to determine whether a specific data point  $(x, y)$  is included in the source training dataset  $D_s^{\text{train}}$ . The adversary  $A$  has black-box access to the target network  $f_t$  through an API. Note that we assume that the adversary  $A$  does not obtain the output of the source model  $f_s$  upon an arbitrary input (i.e.,  $A$  does not have black-box access to  $f_s$ , and have black-box access to only  $f_t$ ), because we want to investigate whether the membership information about  $D_s^{\text{train}}$  is leaked by only the output of  $f_t$ .

As with the previous work on the membership inference attacks (e.g., [14]), we assume that the adversary  $A$  possesses

a shadow training dataset  $D_{\text{shadow}}^{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}_s$  and a shadow test dataset  $D_{\text{shadow}}^{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}_s$  whose elements were drawn from the same distribution  $P_s$  as the source training dataset  $D_s^{\text{train}}$ . Since the goal of the attack is to reveal information on the membership to the source training dataset  $D_s^{\text{train}}$ , the adversary  $A$  has no prior knowledge of  $D_s^{\text{train}}$  itself, hence of whether a data point  $(x, y) \in D_{\text{shadow}}^{\text{train}} \cup D_{\text{shadow}}^{\text{test}}$  is included in  $D_s^{\text{train}}$  or not. To consider the strongest possible adversary, we assume that  $A$  possesses the target training dataset  $D_t^{\text{train}}$  and the target test dataset  $D_t^{\text{test}}$ . We also assume that  $A$  knows whether the small company  $C_{\text{small}}$  constructs a target model  $f_t$  with the freezing approach or the fine-tuning approach.

### B. Attack Algorithms

Next we present our membership inference attacks against the target model  $f_t$ . Here we deal with the learning-based adversary, which is shown to outperform the entropy-based adversary in Section V. To construct the attack algorithms, we introduce a shadow training dataset  $D_{\text{shadow}}^{\text{train},i} \subseteq D_{\text{shadow}}^{\text{train}}$  and a shadow test dataset  $D_{\text{shadow}}^{\text{test},i} \subseteq D_{\text{shadow}}^{\text{test}}$  for each shadow model  $f_{\text{shadow}}^i$ . As with the previous work (e.g., [14]), these shadow training datasets (resp. shadow test datasets) for different shadow models may intersect.

We construct a learning-based adversary with black-box access to the target model  $f_t$  in the following four steps.

#### $\text{MIA}_{\text{L}}^{f_t}$ (learning-based adversary accessing $f_t$ )

- (TL1) We first construct multiple *shadow source models*  $f_{s,\text{shadow}}^i$  to simulate the behavior of the source model  $f_s$  as follows. For each  $i = 1, 2, \dots, c_{\text{shadow}}$ , we train a shadow source model  $f_{s,\text{shadow}}^i$  using the shadow training dataset  $D_{\text{shadow}}^{\text{train},i}$ . Let  $g_{s,\text{shadow}}^i$  be the shallow layers obtained by removing the last fully connected layer from  $f_{s,\text{shadow}}^i$ .
- (TL2) We next construct multiple *shadow target models*  $f_{t,\text{shadow}}^i$  by transfer learning to simulate the behavior of the target model  $f_t$  as follows. For each  $i = 1, 2, \dots, c_{\text{shadow}}$ , we build a shadow target model  $f_{t,\text{shadow}}^i$  by initializing its shallow layers  $g_{t,\text{shadow}}^i$  as the network  $g_{s,\text{shadow}}^i$  obtained in (TL1) and then training  $f_{t,\text{shadow}}^i$  with the target training dataset  $D_t^{\text{train}}$ .
- (TL3) We then obtain a dataset  $D_{\text{attack}}^{\text{train}}$  labeled with membership information (in or out) by applying each shadow target model  $f_{t,\text{shadow}}^i$  to the shadow training dataset  $D_{\text{shadow}}^{\text{train}}$  and the shadow test dataset  $D_{\text{shadow}}^{\text{test}}$ . Specifically, we define  $D_{\text{attack}}^{\text{train}}$  as the set of the records  $(x, y, f_{t,\text{shadow}}^i(x), \text{in})$  for each  $(x, y) \in D_{\text{shadow}}^{\text{train}}$  and  $(x', y', f_{t,\text{shadow}}^i(x'), \text{out})$  for each  $(x', y') \in D_{\text{shadow}}^{\text{test}}$  where  $i = 1, 2, \dots, c_{\text{shadow}}$ . We call  $D_{\text{attack}}^{\text{train}}$  an *attack training dataset*.
- (TL4) We finally construct a membership inference adversary  $A : \mathcal{X} \times \mathcal{Y}_s \times \mathbb{D}\mathcal{Y}_t \rightarrow \{\text{in}, \text{out}\}$  as a classification model (e.g., deep neural network, support vector machine) using the attack training dataset  $D_{\text{attack}}^{\text{train}}$  constructed in (TL3).

Given an input  $(x, y) \in \mathcal{X} \times \mathcal{Y}_s$ , the learning-based adversary  $\text{MIA}_L^f$  obtains  $f_t(x)$  by black-box access to the target model  $f_t$  and outputs the label  $A(x, y, f_t(x))$ .

Since the adversary  $A$  knows whether  $C_{\text{small}}$  constructs a target model  $f_t$  with the freezing approach or the fine-tuning approach (as described in Section A),  $A$  uses the same approach as  $C_{\text{small}}$ . That is, if  $C_{\text{small}}$  uses the freezing (resp. fine-tuning) approach, then  $A$  also uses the freezing (resp. fine-tuning) approach. As a membership inference adversary model  $A$  in  $\text{MIA}_L^f$ , we use a deep neural network (DNN), because it outperformed a support vector machine (SVM) in Section V.

### C. Experimental Evaluation

1) *Datasets and Target Models*: We construct a target model  $f_t$  with the source model  $f_s$  pre-trained in the experiments of Section V. We especially reuse the model constructed from the *ImageNet* dataset as  $f_s$ .

As target training datasets  $D_t^{\text{train}}$ , we use two real datasets: *CIFAR-100* [43] and *Stanford Dogs (Dogs)* [44]. In this section we explain these datasets and the target classification models  $f_t$ .

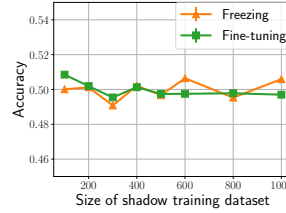
**CIFAR-100.** *CIFAR-100* is an image dataset that has 100 classes each consisting of 600 images. We randomly select the same number of images from each class, and construct 8 target training datasets with different sizes. We set the size of a target training dataset  $D_t^{\text{train}}$  to 100, 200, 300, 400, 500, 600, 800, or 1,000. We also construct a target test dataset  $D_t^{\text{test}}$  that corresponds to each target training dataset  $D_t^{\text{train}}$  and consists of the same number of images as  $D_t^{\text{train}}$ . We train a target model  $f_t$  from each target training dataset with the freezing and fine-tuning approaches. In the training of each model, we use all the layers of  $f_s$  except the last fully connected layer as the transferred source model  $g_s$ . We iterate the training over 120 epochs, and select a model with the best accuracy as  $f_t$ . Table I shows the accuracy of the model  $f_t$  for the target training and test datasets.

**Stanford Dogs (Dogs).** *Dogs* is a dataset that has 120 breeds of dog images. We randomly select 100 breeds from all the breeds. In the same manner as for the *CIFAR-100* dataset, we construct 8 target training and test datasets, and then train a target model  $f_t$  with each training dataset. Table I also shows the accuracy of each target training model in this dataset.

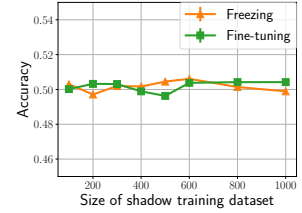
2) *Experimental Setup*: We assume a scenario where an adversary  $A$  possesses the target training dataset  $D_t^{\text{train}}$  as described in Section A. In this scenario, the adversary  $A$  trains a shadow source model  $f_{s,\text{shadow}}^i$  to obtain its shallow layers  $g_{s,\text{shadow}}^i$ . Then  $A$  constructs a shadow target model  $f_{t,\text{shadow}}^i$  from  $g_{s,\text{shadow}}^i$  by transfer learning using  $D_t^{\text{train}}$  as explained in Section B. We apply the same settings to both the target models generated with the *CIFAR-100* and *Dogs* datasets.

In our experiments, we construct 100 shadow target models  $f_{t,\text{shadow}}^i$  for each target model  $f_t$  by the *freezing* or *fine-tuning* approach. Then we produce an attack training dataset  $D_{\text{attack}}^{\text{train}}$

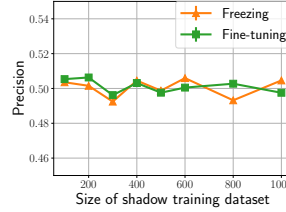
<sup>1</sup>This is the same as the shadow source model  $f_{s,\text{shadow}}^i$  constructed in the experiments of Section V.



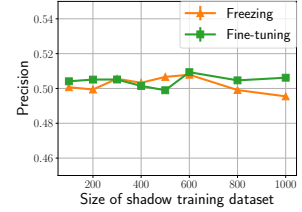
(a) Accuracy in CIFAR-100.



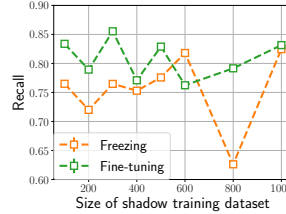
(d) Accuracy in Dogs.



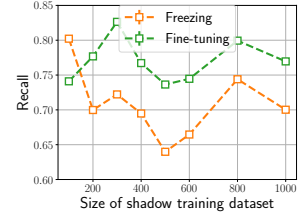
(b) Precision in CIFAR-100.



(e) Precision in Dogs.



(c) Recall in CIFAR-100.



(f) Recall in Dogs.

Fig. 8: Accuracy, precision, and recall of our attacks  $\text{TransMIA}_{L(\text{DNN})}$  with the learning-based approach (DNN). In  $\text{TransMIA}_{L(\text{DNN})}$ , we evaluate both freezing and fine-tuning approaches. Figures (a)–(c) (resp. (d)–(f)) are the results in the *CIFAR-100* (resp. *Dogs*) dataset. Each line indicates the average of results for 100 classes. The dashed lines and outline markers in Figures (c) and (f) indicate that the corresponding precision value is less than 60%.

as described in Section B. For each class label  $y$ , we construct an attack model, i.e., 100 attack models in total.

We perform membership inference attacks with the *learning-based approach* using a DNN. We denote by  $\text{MIA}_{L(\text{DNN})}^f$  the attack with the DNN. The DNN has three fully-connected hidden layers with 50, 30, and 5 neurons, and is trained with 50 epochs. We evaluate attack performance with *accuracy*, *precision*, and *recall*.

3) *Experimental Results*: We present experimental results on the learning-based membership inference attack  $\text{MIA}_{L(\text{DNN})}^f$  with a DNN. Fig. 8 shows the accuracy, precision, and recall of the attacks with the freezing and fine-tuning approaches. Those graphs indicate that the accuracy and precision of  $\text{MIA}_{L(\text{DNN})}^f$  are much lower than those of  $\text{TransMIA}_{L(\text{DNN})}$ , and those values are very close to 0.5. As mentioned in Section V-C, when the accuracy and precision are close to 0.5, the adversary fails to infer membership information, even though the corresponding recall is high.

From these results, we argue that even if the small company

TABLE I: Accuracy of target models in the CIFAR-100 and Dogs datasets.

Size	CIFAR-100				Dogs			
	Freezing		Fine-tuning		Freezing		Fine-tuning	
	Train	Test	Train	Test	Train	Test	Train	Test
100	0.810	0.080	0.500	0.070	0.360	0.040	0.770	0.050
200	0.635	0.105	0.610	0.095	0.505	0.055	0.765	0.055
300	0.473	0.107	0.707	0.090	0.507	0.057	0.407	0.053
400	0.545	0.115	0.703	0.105	0.525	0.065	0.765	0.078
500	0.450	0.118	0.602	0.110	0.392	0.062	0.646	0.074
600	0.555	0.133	0.850	0.128	0.183	0.057	0.280	0.058
800	0.334	0.146	0.914	0.159	0.276	0.075	0.330	0.075
1,000	0.494	0.165	0.920	0.183	0.316	0.065	0.891	0.091

$C_{\text{small}}$  allows the adversary to have black-box access to the trained target model  $f_t$ , the adversary fails to infer membership information about  $D_s^{\text{train}}$ . It is left as future work to provide a formal proof that the membership information of the source model  $f_s$  is not leaked by the output of the trained target model  $f_t$ .