

A Sample-Level Evaluation and Generative Framework for Model Inversion Attacks

Haoyang Li¹, Li Bai¹, Qingqing Ye¹, Haibo Hu^{1*}, Yaxin Xiao¹, Huadi Zheng², Jianliang Xu³

¹Department of Electrical and Electronic Engineering, The Hong Kong Polytechnic University

²Huawei Technology

³ Department of Computer Science, Hong Kong Baptist University

{hao-yang9905.li, baili.bai}@connect.polyu.hk, {qqing.ye, haibo.hu}@polyu.edu.hk, 20034165r@connect.polyu.hk, zhenghuadi@huawei.com, xujl@comp.hkbu.edu.hk

Abstract

Model Inversion (MI) attacks, which reconstruct the training dataset of neural networks, pose significant privacy concerns in machine learning. Recent MI attacks have managed to reconstruct realistic label-level private data, such as the general appearance of a target person from all training images labeled on him. Beyond label-level privacy, in this paper we show sample-level privacy, the private information of a single target sample, is also important but under-explored in the MI literature due to the limitations of existing evaluation metrics. To address this gap, this study introduces a novel metric tailored for training-sample analysis, namely, the Diversity and Distance Composite Score (DDCS), which evaluates the reconstruction fidelity of each training sample by encompassing various MI attack attributes. This, in turn, enhances the precision of sample-level privacy assessments.

Leveraging DDCS as a new evaluative lens, we observe that many training samples remain resilient against even the most advanced MI attack. As such, we further propose a transfer learning framework that augments the generative capabilities of MI attackers through the integration of entropy loss and natural gradient descent. Extensive experiments verify the effectiveness of our framework on improving state-of-the-art MI attacks over various metrics including DDCS, coverage and FID. Finally, we demonstrate that DDCS can also be useful for MI defense, by identifying samples susceptible to MI attacks in an unsupervised manner.

Introduction

With today’s explosive growth in Machine Learning (ML), all businesses are collecting personal data more frequently than ever (Li et al. 2023). As such, there is an increasing demand of privacy preservation during model training and deployment (Yan et al. 2022; Bai et al. 2024; Hu et al. 2024; Zhao et al. 2016). However, it is well known that by exploiting the confidence score of neural networks, Model Inversion (MI) attacks can partially reconstruct the training dataset from an ML model (Fredrikson, Jha, and Ristenpart 2015). Compared with membership or property inference attacks that only infer certain membership or property information as a binary classification task (Carlini et al. 2022),

an MI attack targets at the whole private information in the training dataset and reconstructs multiple private features simultaneously (Song and Namiot 2022), which poses even more severe threats than its inference counterpart. In line with previous studies (Zhang et al. 2020; Chen et al. 2021; Struppek et al. 2022), we consider white-box MI attacks, where the attacker has access to the specific parameters and architecture of the victim model.

Given the inherent difficulty for MI attacks to recover every sample within a large-scale dataset, recent works (Struppek et al. 2022; An et al. 2022) have lowered down their goal to recover one or few representatives in each class label. This trend was first motivated in the field of face recognition (Wang et al. 2021b), where each label corresponds to the identity of a person. Since these label representatives can reveal the general appearance of a person, they arise as a significant concern in recent studies (Wang et al. 2021a; Nguyen et al. 2023). However, the potential for privacy compromise varies across applications. In some domains beyond facial recognition, label representatives may convey little or even no private information, which in turn weakens such threats. For instance, in disease diagnosis, each label corresponds to a well-known disease, so label representatives merely encapsulate public knowledge about the conditions of this disease. However, MI attacks may still threaten **sample-level** privacy by reconstructing individual samples.

To address the imperative to unearth and safeguard sample-level privacy, this study critiques existing evaluation metrics of MI by first demonstrating their inadequacy in fully capturing the multifaceted characteristics of MI attacks. Furthermore, metrics based on distributional similarity are sensitive to the distribution of reconstructed samples, so that their evaluation results can be easily fooled by controlling the sample distribution (*e.g.*, producing redundant samples or approximating the distribution with fewer samples). To overcome these limitations, we introduce a novel metric, namely *Diversity and Distance Composite Score* (DDCS), that assesses the integrity of sample-level privacy in the context of MI attacks. In essence, DDCS uniquely evaluates the extent of reconstruction by the attack for each sample in the training dataset, so that it can comprehensively incorporate diverse attributes (*e.g.*, distance, coverage and distributional similarity) to indicate a good MI attack and is

*Corresponding author

robust against sample distribution manipulations. From the defense’s perspective, DDCS quantifies the susceptibility of individual samples to MI attacks by assigning each with a reconstruction distance, enabling the identification and protection of vulnerable samples in an unsupervised manner.

Leveraging Diversity and Distance Composite Score (DDCS) to evaluate the success of reconstruction, we observe that many samples in the training dataset remain un-reconstructable by even the most advanced MI attacks (Struppek et al. 2022). This observation coincides with the limited diversity and coverage of current MI methodologies, mainly because prior metrics focus primarily on label-level privacy. In response, we introduce a Generative Adversarial Network (GAN) augmentation framework designed to enhance both coverage and diversity of MI attacks, through mitigating the produced artifacts while reserving the advantages from the inversion-specific GAN (Karras et al. 2020; Chen et al. 2021). To achieve this, our framework transfers a pre-trained GAN with entropy loss regularized with natural gradient descent (Martens 2020), so as to expand the attacker’s generative ability and mitigate artifacts typically associated with deep GAN structures, such as those found in StyleGAN (Karras et al. 2020).

To validate our methodologies, we conduct comprehensive experiments on face recognition and dog breed classification datasets under various model architectures and MI algorithms. The results affirm the comprehensiveness of DDCS, underscoring its utility in evaluating sample-level privacy and robustness against manipulations of sample distributions. Further, the experiments substantiate the efficacy of our GAN augmentation framework in enhancing the performance of leading MI attacks, as evidenced by improvements in Fréchet Inception Distance, coverage, and DDCS metrics. Additionally, we also conduct a series of ablation studies to scrutinize the influence of different experimental parameters, ensuring a thorough understanding of our framework’s dynamics and its implications for privacy-preserving ML. To summarize, our contributions in this paper are as follows:

- We introduce DDCS, a comprehensive evaluation metric for MI attacks with great potential for sample-level privacy analysis.
- We propose a GAN augmentation framework to improve MI attacks on DDCS while retaining the image quality.
- We conduct extensive experiments on computer vision tasks to demonstrate DDCS and verify the effectiveness of our framework.

The rest of the paper is organized as follows. We first review recent work on MI attacks. Later, we point out the limitations of evaluation metrics and introduce DDCS. We then propose our GAN augmentation framework and conduct extensive experiments before the conclusion.

DDCS: A New Evaluation Metric for MI

Prior studies on MI attacks have employed a range of evaluation metrics, each with its own limitations. Notably, many of these metrics overlook critical attributes of the reconstructed data, such as diversity, and are susceptible to alteration in

data distribution caused by non-essential factors, such as the presence of redundant samples. This issue is prevalent across MI attack evaluations. In the rest of this section, we delve into the impact of these limitations on the thorough assessment of MI attacks and introduce a novel metric designed to overcome these challenges.

Limitations of Existing Metrics

An ideal MI attack (Wang et al. 2021a), as the name suggests, aims to accurately restore the target training dataset \mathcal{D}_{tar} . As such, attacks which can successfully reconstruct every sample in \mathcal{D}_{tar} are more severe than those only recover sensitive labels. While recent studies have highlighted the risks of label-level privacy breaches in tasks like face recognition, where label representatives can disclose an individual’s identity and appearance (An et al. 2022), we argue that the reconstruction of sample-level data poses a far greater threat to privacy. For example, in face recognition, the variety in reconstructed samples can expose nuanced personal details, such as an individual’s preferred hairstyles, accessories, and expressions (Wang et al. 2021a). This is particularly true in fields beyond face recognition, such as disease diagnosis and dog breed classification, where labels convey commonly known information.

The oversight of sample-level privacy may stem from the inadequacy of existing evaluation metrics to fully and precisely evaluate crucial characteristics indicative of successful MI attacks. To address this gap, we draw three fundamental traits that are essential for an ideal MI attack on sample-level privacy — accuracy/distance, coverage, and distributional similarity. This discussion also explores the shortcomings of current evaluation methodologies in accurately assessing these critical attributes.

Accuracy/Distance. Accuracy and distance both encourage MI attacks focusing on label-level privacy to generate label representatives. Recall that the objective function of MI is to maximize the sample’s output confidence given the victim model (Fredrikson, Jha, and Ristenpart 2015). Thus, the optimization of MI attacks can easily overfit to the specific structure of the victim neural network, and a successful MI attack should then be independent of the classifier architecture. Therefore, previous works train a model with a different structure from the victim model and measure the accuracy of reconstructed dataset \mathcal{D}_{rec} . In this context, high classification accuracy on the reconstructed dataset \mathcal{D}_{rec} signifies a reduced distance between \mathcal{D}_{rec} and the target dataset \mathcal{D}_{tar} . Comparable to accuracy, feature distance is another metric frequently employed in MI assessments, quantifying the l_2 distance between reconstructed and target samples within the evaluator’s feature space.

Coverage. Coverage refers to the proportion of \mathcal{D}_{tar} that are reconstructed by \mathcal{D}_{rec} and can be increased by improving the diversity of \mathcal{D}_{rec} . While accuracy and feature distance are the primary metrics for evaluating MI attacks, they overlook the diversity of samples and fail to comprehensively assess the coverage of MI attacks. In the toy example of Figure 1 where \mathcal{D}_{tar} only contains five images, there is a rich attack that can successfully recover all images in \mathcal{D}_{tar} and a poor attack that recovers only one of them. If all

samples from the rich and poor attack are identical to some images in \mathcal{D}_{rec} , both attacks achieve the same average accuracy, albeit the rich attack is obviously more threatening than the poor attack.

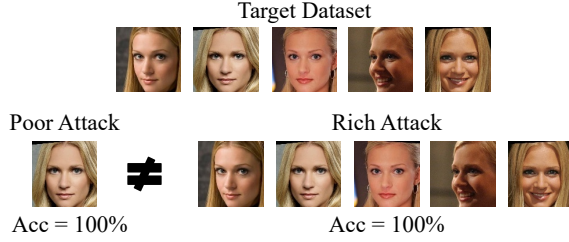


Figure 1: When each reconstructed sample is close enough to the target dataset, average accuracy or distance fail to differentiate between the rich and poor attack.

Distributional Similarity. Being aware of the limitation in accuracy and distance, recent works (Chen et al. 2021; Wang et al. 2021a; Struppek et al. 2022) have incorporated distributional similarity, e.g., Fréchet Inception Distance (FID) (Heusel et al. 2017), as an additional metric by treating \mathcal{D}_{tar} and \mathcal{D}_{rec} as two distinct image distributions. Since the diversity of \mathcal{D}_{rec} will affect its distribution, the problem of Figure 1 can be alleviated with distributional similarity. However, the distribution of images can be easily affected by factors that are trivial in privacy leakage, leading to misjudgment of MI attacks. One example is the bias caused by redundant samples that MI attacks can easily generate, which leads to a distribution shift between \mathcal{D}_{tar} and \mathcal{D}_{rec} , thereby causing a worse FID that is not true. As in Figure 2, there is an attack that can restore most of the samples without any duplication (i.e., strong and concise attack), and an attack that can restore the same diverse samples with some duplications (i.e., strong but verbose attack). The former attack can achieve a better FID, but it is premature to conclude that the former attack is more powerful than the latter, because they both successfully restore four images in \mathcal{D}_{tar} . This phenomenon arises because distributions can be approximated using varying sample batches, allowing for the manipulation of metrics by altering the sample distribution. We will empirically show how FID is susceptible to such manipulations in the experiment, and therefore these factors do not fundamentally characterize a good MI attack.

Diversity and Distance Composite Score

We believe the core reason for the above limitations is that these metrics overly concentrate on the distribution of reconstructed samples (\mathcal{D}_{rec} -oriented), which leads to an incomplete assessment of privacy leakage in the target samples \mathcal{D}_{tar} . This approach diverges from the objectives of ideal MI attacks, as outlined previously. Consequently, an MI algorithm could generate high-fidelity label representatives and manipulate the distribution of \mathcal{D}_{rec} to perform favorably according to these metrics.

Different from \mathcal{D}_{rec} -oriented metrics introduced in previous works, we truly measure the degree to which \mathcal{D}_{tar} is




Figure 2: Redundant samples change the data distribution, thus affecting the FID calculation, although both attacks successfully recover four images from the target dataset.

recovered by \mathcal{D}_{rec} , and propose a more robust and comprehensive evaluation metric, namely *Diversity and Distance Composite Score* (DDCS). We start from the ideal MI attack, where each sample from \mathcal{D}_{tar} has its one-to-one matching in \mathcal{D}_{rec} (i.e., this attack achieves perfect diversity and distance regarding \mathcal{D}_{tar}), and any other attack that reduces \mathcal{D}_{rec} 's diversity and distance to \mathcal{D}_{tar} is a suboptimal attack, which inspires us how to calculate DDCS, as shown in Algorithm 1. One worth-noting term used in DDCS is the *Reconstruction Distance* d_{tar}^i for every target sample $x_{tar}^i \in \mathcal{D}_{tar}$. Specifically, given a distance metric $d(\cdot, \cdot)$, if x_{rec}^j 's closest target sample is x_{tar}^i with the distance as $d(x_{rec}^j, x_{tar}^i)$, x_{tar}^i is said to be reconstructed by x_{rec}^j with a reconstruction distance of $d(x_{rec}^j, x_{tar}^i)$ (Line 3). $d(x_{rec}^j, x_{tar}^i)$ is then recorded into a set, denoted by S_{tar}^i of x_{tar}^i for further handling.

Algorithm 1: DDCS Computation

Input: Target dataset \mathcal{D}_{tar} , reconstructed dataset \mathcal{D}_{rec} , constant c

Output: $DDCS_{avg}$, $DDCS_{best}$

- 1: Initialize $S_{tar}^i = \emptyset$ for $x_{tar}^i \in \mathcal{D}_{tar}$
 - 2: **for** $x_{rec}^j \in \mathcal{D}_{rec}$ **do**
 - 3: Record $d_{tar}^i = d(x_{tar}^i, x_{rec}^j)$ into S_{tar}^i , where x_{rec}^j and x_{tar}^i achieve the smallest d_{tar}^i for $x_{tar}^i \in \mathcal{D}_{tar}$
 - 4: **end for**
 - 5: Initialize $DDCS_{avg} = 0$, $DDCS_{best} = 0$
 - 6: **for** $x_{tar}^i \in \mathcal{D}_{tar}$ **do**
 - 7: $DDCS_{avg} += mdf[\text{avg}(S_{tar}^i) + c]$
 - 8: $DDCS_{best} += mdf[\min(S_{tar}^i) + c]$
 - 9: **end for**
 - 10: Normalize $DDCS_{avg}$ and $DDCS_{best}$ with $|\mathcal{D}_{tar}|$
-

In Lines 7 and 8, we propose two variants of DDCS tailored to distinct scenarios. To ensure that higher DDCS values indicate better outcomes, we employ monotonically decreasing functions, $mdf[\cdot]$, in its calculation, specifically utilizing the simple reciprocal function on $\text{avg}(S_{tar}^i) + c$ and $\min(S_{tar}^i) + c$ in this paper. $DDCS_{avg}$ quantifies the overall privacy risk of \mathcal{D}_{tar} against \mathcal{D}_{rec} , calculated as the sum of the reciprocals of the averages for each S_{tar}^i , i.e., $\text{avg}(S_{tar}^i)$. $DDCS_{best}$ assesses the minimum reconstruction distance across all S_{tar}^i , i.e., $\min(S_{tar}^i)$, offering insight into the potential maximum effectiveness of an attack on \mathcal{D}_{tar} .

If x_{tar}^i is not recovered by any reconstructed sample, we set their $mdf[\text{avg}(S_{tar}^i) + c]$ and $mdf[\text{min}(S_{tar}^i) + c]$ to 0. Besides, we reserve a constance c to regularize the range of DDCCS when d_{tar}^i is so small. As a final note, both $\text{DDCCS}_{\text{avg}}$ and $\text{DDCCS}_{\text{best}}$ are normalized by the total number of samples $|\mathcal{D}_{tar}|$ in the target dataset (Line 1).

DDCCS mitigates those limitations discussed before from two aspects. Firstly, this metric addresses the privacy for **every** sample in the target dataset (\mathcal{D}_{tar} -oriented). Thus, DDCCS encourages MI attacks to focus on sample-level privacy instead of label representatives, with high robustness against distribution change of \mathcal{D}_{rec} as in Figure 2. Secondly, DDCCS takes multiple important characteristics into account, so that an MI attack can win this metric by truly attacking the privacy of target samples. For example, if an attack reaches higher coverage, more samples from \mathcal{D}_{tar} will be reconstructed, increasing DDCCS value and mitigating the problem in Figure 1. On the other hand, when diversity remains unchanged, producing better reconstructed sample to reduce the reconstruction distance can also improve DDCCS. Besides, reconstruction distance of DDCCS can also be harnessed beyond attacks, such as measuring per-example vulnerability for \mathcal{D}_{tar} to selectively protect those vulnerable samples, which will be discussed in detail in the experiment. This application scenario has been recently studied in the context of membership inference attacks (Carlini et al. 2022), but we are the first to study it in MI scenarios.

Enhancing MI Attacks over DDCCS

Based on the \mathcal{D}_{tar} -oriented evaluation by DDCCS, we evaluate the actual extent of privacy intrusion by attackers into \mathcal{D}_{tar} , as depicted in Figure 3. This figure illustrates the percentage of per-label samples in \mathcal{D}_{tar} that are reconstructable by at least one counterpart in \mathcal{D}_{rec} , reflecting the attack’s coverage. Despite the focus of prior research on producing high-quality label representatives, we note that even the most advanced MI attacks (Struppek et al. 2022) fall short in achieving comprehensive coverage. This revelation significantly undermines the perceived threat level of MI attacks, as the ML training party can easily identify these samples using DDCCS and protect them. This insight underscores the need for us to enhance the generative capabilities of attackers in existing MI attack methodologies.

Difficulties on Training with Entropy Loss

Enhancing the generative capacity of a GAN for MI attacks involves refining the training methodology of the generative network, allowing the generator G to leverage knowledge from the victim model V . To facilitate more effective utilization of V ’s knowledge by the attacker, KEDMI (Chen et al. 2021) addresses this by incorporating an additional entropy loss \mathcal{L}_{adv} , into the standard DCGAN training regimen of G . \mathcal{L}_{adv} is calculated based on the information entropy of V ’s confidence score for a set of generated images $G(z)$. Minimizing \mathcal{L}_{adv} aims to enhance the similarity between $G(z)$ and the private training datasets of V . However, this approach introduces two notable challenges:

P1. The addition of \mathcal{L}_{adv} to the training deviates from the

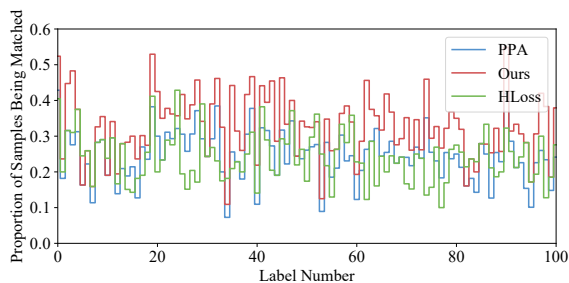


Figure 3: Visualization of DDCCS indicating the proportion of samples for the first 100 labels of VGG16BN-UMDFaces that are matched to \mathcal{D}_{rec} with three attacks (PPA, HLoss and Ours). As explained in the section of DDCCS, a target sample x_{tar}^i will be matched if its set S_{tar}^i is non-empty after the execution of Algorithm 1.

GAN’s original objective of maximizing image quality, consequently compromising it.

P2. Training a GAN from the scratch with \mathcal{L}_{adv} is redundant under the assumption of publicly available pre-trained GANs.

Compared to P2, P1 poses a more significant challenge, as evidenced by recent findings (Struppek et al. 2022) where the KEDMI approach underperforms traditional methods when integrated with StyleGAN. We attribute P1 to the optimization of entropy loss under a loosely enforced image quality constraint, which inadvertently leads to the generation of artifacts rather than the intended private information of V . This issue is exemplified in Figures 4 (a) and (b), where, despite the intention for entropy loss to enhance V ’s confidence in the generated images, it instead prompts G to produce artifacts rather than accurate representations of target dataset \mathcal{D}_{tar} . In this figure, FID assesses image quality, different from its use in MI attacks where it evaluates attack performance.

Natural Gradient for Entropy Loss

Our methodology for refining entropy loss draws upon the manifold hypothesis, which states that the distribution of natural images, considered as a high-dimensional manifold, is representable through latent codes on a lower-dimensional manifold (Goodfellow, Bengio, and Courville 2016; Shao, Kumar, and Thomas Fletcher 2018). In other words, artifacts caused by entropy loss are interpreted as the result of generated images deviating from their native manifold. Therefore, we propose to mitigate the emergence of artifacts by maintaining the generated images within their original manifold throughout the training of \mathcal{L}_{adv} -assisted GAN.

To tackle the above P1, we incorporate Natural Gradient Descent (NGD) (Martens 2020) into our methodology, aiming to preserve the manifold structure of generated images while training with entropy loss. NGD, an approximate second-order optimization technique, is conceptualized as the optimization on the Riemannian manifold (Boumal 2023). Its implementation relies on the inverse of Fisher Information Matrix (FIM), denoted as \mathcal{F}^{-1} , an intrinsic distance metric that facilitates the preservation

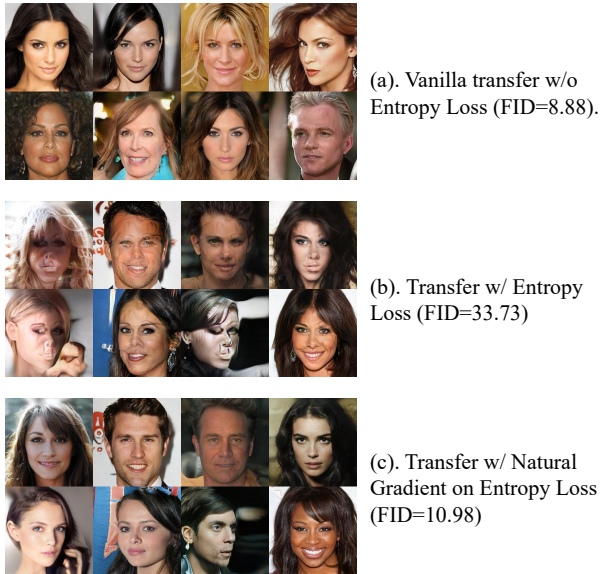


Figure 4: Snapshots and image quality of generated images for three different approaches. Image quality, evaluated by FIDs, are calculated with the same random seed and training configurations. Snapshots are generated using the same and fixed latent codes.

of variable on the manifold through successive optimization steps. To approximate \mathcal{F} , (Wang and Ponce 2021) suggests utilizing the Hessian of a squared LPIPS distance function $H_{G(z)}(d^2(G(z), G(z_0)))$ as the metric tensor for the manifold of deep generative image models. This insight underpins our proposed NGD-based strategy for optimizing entropy loss:

$$P\left(\frac{\partial \mathcal{L}_{adv}}{\partial G}\right) = H_{G(z)}^{-1}(d^2(G(z), G(z_0))) \frac{\partial \mathcal{L}_{adv}}{\partial G(z)} \cdot \frac{\partial G(z)}{\partial G} \quad (1)$$

In each optimization step, we project the original gradient $\frac{\partial \mathcal{L}_{adv}}{\partial G}$ onto its natural gradient $P\left(\frac{\partial \mathcal{L}_{adv}}{\partial G}\right)$ using the inverse Hessian matrix of the squared distance function $H_{G(z)}^{-1}(d^2(G(z), G(z_0)))$ for a given batch of generated images $G(z)$ and their entropy loss \mathcal{L}_{adv} . It is worth noting that alternative methodologies for specifying the metric tensor of the image manifold, such as employing the Jacobian of the generator G (Shao, Kumar, and Thomas Fletcher 2018), are also viable within this framework.

Figure 4 (c) illustrates the efficacy of Equation 1 in artifact mitigation and improvement of image quality, where FID is reduced from 33.73 to 10.98 and the artifacts are enhanced with realistic features.

To reduce the computational cost in P2, we adopt a transfer learning approach for our training regimen. Utilizing a pre-trained GAN, we apply our training framework on an auxiliary dataset for a limited number of epochs. In line with (Wang and Ponce 2021), we employ the Hessian Vector Product (HVP) to efficiently approximate the Hessian, pre-compute a batch of HVPs, and then use them throughout the training process. Notably, the image domains produced by a GAN trained with entropy loss diverge from those generated

by a conventional (vanilla) GAN, leading to a restricted and less diverse \mathcal{D}_{rec} when using a single GAN model (Struppek et al. 2022; Wang et al. 2021a). To counteract this limitation, we propose to augment a vanilla GAN with an \mathcal{L}_{adv} -enhanced GAN for a more diverse \mathcal{D}_{rec} .

To summarize, given a pre-trained GAN and victim model, we calculate the standard StyleGAN training loss as well as the entropy loss at each training step. The gradients from the loss of standard training and entropy loss will be combined to update the generator G , while the gradient from the entropy loss will first be projected using Equation 1 before applying gradient descent. Finally, the optimized GAN, trained with this methodology, will produce the reconstructed samples following the standard MI recovery strategy.

Experiments

Experimental Setup

Datasets. We conduct experiments¹ on four high-resolution image datasets for two image classification tasks, namely, face recognition and dog breed classification. For face recognition, we use UMDFaces (Bansal et al. 2017) face dataset for training victim models and CelebA-HQ (Karras et al. 2018; Liu et al. 2015) dataset as the attacker’s auxiliary dataset. For dog breed classification, Stanford Dogs (Khosla et al. 2011) dataset is selected as the training dataset of victim models and Tsinghua Dogs (Zou et al. 2020) for the auxiliary datasets. Similar to previous works (Chen et al. 2021; Zhang et al. 2020), we take a subset of 1000 identities from UMDFaces that contains 50981 training samples and 3000 test samples as \mathcal{D}_{tar} to train the victim models. Following (Struppek et al. 2022), we use all labels for Stanford Dogs and split them into 18780 training samples and 1800 test samples. To prevent strong assumptions about attacks, we limit the attacker’s sample size in the auxiliary dataset to 30000 for both tasks.

Models. We train 4 types of victim models: (1) VGG16 (Simonyan and Zisserman 2015) with Batch Normalization (Ioffe and Szegedy 2015) (VGG16BN), (2) ResNet50 (He et al. 2016), (3) Improved ResNet50 with Squeeze-and-Excitation (Hu, Shen, and Sun 2018) blocks (IR50-SE), and (4) AlexNet (Krizhevsky 2014). We resume pre-trained StyleGAN2 (Karras et al. 2020) of resolution 256×256 trained on two datasets, CelebA-HQ for face identification task and LSUN-Dogs (Yu et al. 2015) for dog breed classification task.

Competing MI Attacks. Our GAN augmentation framework (Ours) is compared against two leading MI attack strategies: the algorithms of Plug & Play Attack (Struppek et al. 2022) (PPA) and training with entropy loss of KEDMI (Chen et al. 2021) (HLoss). Given that KEDMI’s algorithms, aside from HLoss, are tailored specifically for DCGAN (Struppek et al. 2022), and PPA is recognized as a versatile, open-source framework for incorporating StyleGAN in MI attacks (Struppek et al. 2022; Qi et al. 2023), we adopt PPA’s StyleGAN-based architecture to implement

¹https://github.com/haoyangli/ASTAPLE/DDCS_MI.git

both HLoss and our proposed approach. Note that while several MI studies (Zhang et al. 2020; An et al. 2022; Wang et al. 2021a; Nguyen et al. 2023; Qi et al. 2023) focus on the latter phase of MI attacks (i.e., fine-tuning GAN latent codes), our approach primarily enhances the GAN transfer learning phase of MI Attacks. Thus, our approach is compatible with these works, as well as future MI attack strategies centered on latent code optimization.

Competing Evaluation metrics. We evaluate the attacks in terms of $DDCS_{avg}$ and $DDCS_{best}$ as introduced in Algorithm 1, as well as other standard metrics including accuracy, feature distance, FID, and coverage from (Naeem et al. 2020; Wang et al. 2021a). Additionally, an Inception-v3 (Szegedy et al. 2016) model is trained on the target dataset as the evaluator on top-1 accuracy (Acc@1), top-5 accuracy (Acc@5) and feature distance (Dist). A publicly available pre-trained Inception-v3 model is used to calculate the coverage and FID between the reconstructed samples and target samples. Furthermore, to measure LPIPS distance, a default VGG network setting is adopted.

DDCS Reveals Per-example Vulnerability Against MI Attacks

Previous works fail to pinpoint which sample from \mathcal{D}_{tar} is the most similar to a given reconstructed sample. As a result, in the visualization, one must manually select the most similar sample from \mathcal{D}_{tar} , which is not only best guaranteed but also hard to interpret. On the other hand, **DDCS can indicate how vulnerable each sample is under the MI attack in an unsupervised manner.** Specifically, recall that each sample from \mathcal{D}_{tar} is assigned a reconstruction distance to indicate its degree of restoration by \mathcal{D}_{rec} in either worst case or average case. Therefore, we can easily build the reconstruction pairs between \mathcal{D}_{tar} and \mathcal{D}_{rec} by searching for samples in \mathcal{D}_{rec} that can achieve the smallest reconstruction distance to a given sample in \mathcal{D}_{tar} .

In Figure 5, we present several samples from \mathcal{D}_{tar} that are successfully matched by samples from \mathcal{D}_{rec} together with their reconstruction distance. Samples that are not matched by any samples from \mathcal{D}_{rec} are shown in the rightmost column. Results are averaged from 10 independent runs to account for randomness. These matched samples can be regarded as samples that are vulnerable to MI attacks and are more likely to leak their private information. The data or model owner can thus focus on protecting these samples, by perturbing their values or removing them from the training. Besides, we find that most samples in \mathcal{D}_{tar} are unmatched, and these samples usually have more complex features than paired samples. As such, future works could target at those difficult-to-attack samples to achieve a better coverage of MI attacks.

Comparison of Various Evaluation Metrics

In this subsection, we experimentally study the evaluation and robustness of DDCS against limitations discussed previously. Specifically, we build a customized \mathcal{D}_{rec} to simulate an MI attack that are very close to the ideal MI attack and can reconstruct \mathcal{D}_{tar} very well. As such, we can then



Figure 5: Visualization of reconstruction pairs between target (left) and reconstructed (right) samples for VGG16BN-UMDFaces with reconstruction distance attached on their bottom right corner. Samples on the right most column have no reconstruction pairs.

observe how DDCS successfully evaluates this customized attack and how existing metrics fail to achieve it. Two customized datasets, denoted as D1 and D2, are constructed based on UMDFaces and the settings as follows:

D1. This dataset represents a \mathcal{D}_{rec} of perfect distance to \mathcal{D}_{tar} but poor diversity. To construct D1, we sub-sample a fixed number of images for every label, so that those samples have their identical matches in original \mathcal{D}_{tar} and the diversity of D1 can be manually controlled.

D2. This dataset represents a \mathcal{D}_{rec} of perfect distance and diversity to \mathcal{D}_{tar} but with a modified sample distribution different from \mathcal{D}_{tar} . To construct D2, we create a fixed number of redundant samples from \mathcal{D}_{tar} to result in different distribution between \mathcal{D}_{tar} and \mathcal{D}_{rec} .

In D1, we simulate the condition of varied diversity in \mathcal{D}_{rec} , and thus DDCS is compared with accuracy in this case. As for D2, we test the metrics' robustness against distribution change of \mathcal{D}_{rec} , a common phenomenon in MI attacks, where we choose FID as the competing metric for DDCS. Since both customized datasets can achieve perfect distance to \mathcal{D}_{tar} , we regularize the range of DDCS by setting c in Algorithm 1 to 1.0.

Figure 6 and 7 are the results for D1 and D2 respectively. In Figure 6, as more images are sampled, DDCS successfully captures the increasing diversity of \mathcal{D}_{rec} and thus increases. On the other hand, since \mathcal{D}_{rec} has a very small distance from the samples of \mathcal{D}_{tar} , the accuracy approaches 100% but it ignores the diversity of \mathcal{D}_{rec} . In Figure 7, since FID is sensitive to the changes of sample distribution, it always grows with the addition of additional redundant sam-

Model	Attack	↑ Acc@1	↑ Acc@5	↓ Dist	↓ FID	↑ Coverage	↑ DDCS _{avg}	↑ DDCS _{best}
VGG16BN	HLoss	49.64%	69.10%	15978.52	64.00	0.1601	0.4162	0.4308
	PPA	97.67%	99.89%	7044.45	48.35	0.3213	0.5046	0.5240
	Ours	97.67%	99.86%	7065.28	47.64	0.3907	0.6772	0.7114
ResNet50	HLoss	35.15%	57.01%	18086.90	51.48	0.1727	0.4118	0.4243
	PPA	81.77%	94.84%	10605.46	47.10	0.3234	0.4901	0.5088
	Ours	81.65%	95.01%	10642.02	45.83	0.3951	0.6632	0.6961
IR50-SE	HLoss	31.58%	49.47%	18323.38	79.28	0.1303	0.4238	0.4364
	PPA	85.65%	95.64%	9728.30	47.55	0.3186	0.5146	0.5339
	Ours	86.12%	95.92%	9668.79	46.18	0.3901	0.7017	0.7367
AlexNet	HLoss	16.68%	36.07%	19990.63	62.06	0.1543	0.3901	0.4031
	PPA	53.96%	78.54%	13710.45	45.62	0.2979	0.4600	0.4762
	Ours	53.97%	78.43%	13751.00	44.11	0.3696	0.6193	0.6481

Table 1: Comparison on UMDFaces dataset between our approach (Ours), HLoss and PPA across different metrics. ↑ and ↓ mean the higher the better and the lower the better, respectively. The best values for each metric and each model are in bold.

ples. Thanks to the \mathcal{D}_{rec} -oriented evaluation, DDCS is robust against the varied distribution and stays in a high and stable range.

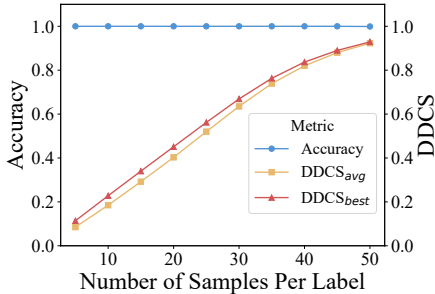


Figure 6: Evaluation results for various metrics on a customized UMDFaces dataset, in which the diversity is controlled with different number of samples per label.

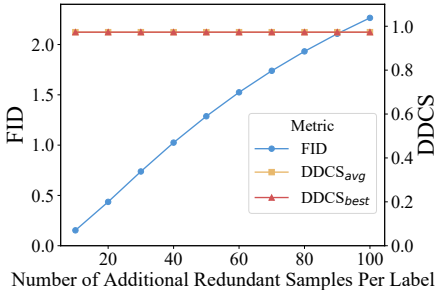


Figure 7: Evaluation results for various metrics on a customized UMDFaces dataset, in which the distribution is controlled with different number of redundant samples per label.

Comparison Results of MI Attacks

Table 1 shows the comparison results on UMDFaces between Ours, HLoss and PPA. We keep two decimal places for numbers greater than 10, and four decimal places for the rest.

Per-label Reconstruction of MI Attacks

In this subsection, we study the reconstruction degree of MI attacks in a label-wise manner through visualizing the

reconstruction distance from DDCS. Among three attacks discussed in UMDFaces dataset and VGG16BN model scenario, Figure 3 plots the proportion of target samples that have at least one matching to \mathcal{D}_{rec} , and Figure 8 illustrates the per-label reconstruction distance by averaging the cumulative distance of these matched target samples. We observe that our method significantly enhances the proportion of those samples that are reconstructed by \mathcal{D}_{rec} , while achieving a similar reconstruction distance to PPA, thereby improving DDCS. On the other hand, suffering from entropy loss with soft-constraint that deteriorates both the image quality and generative power, HLoss achieves a much higher average reconstruction distance and thus a lower DDCS, though it can cover some labels in Figure 3.

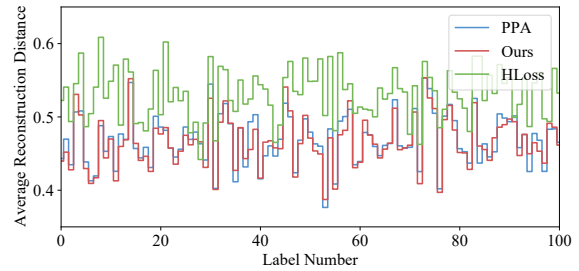


Figure 8: Average reconstruction distance for matched samples in each label with VGG16BN-UMDFaces and three attacks (PPA, HLoss and Ours).

Conclusion

In this paper, we point out several limitations in the existing evaluation metrics of MI attacks and propose a novel metric, Diversity and Distance Composite Score (DDCS), to alleviate those limitations and encourage a more comprehensive evaluation of MI attacks. Furthermore, to enhance existing MI attacks, we further propose a GAN augmentation framework with transfer learning for state-of-the-art MI attacks. Overall, we emphasize the importance of informing the academic community about potential threat models of MI attacks and introducing new perspectives on privacy measurement, to foster the development of more robust privacy-preserving ML algorithms.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No: 92270123, 62072390, and 62372122), and the Research Grants Council, Hong Kong SAR, China (Grant No: 15203120, 15226221, 15209922, 15210023, and C2004-21GF). We appreciate Mr Andy Cheung Yat-ming for constructive discussions in the natural gradient descent part and anonymous reviewers for their constructive comments.

References

- An, S.; Tao, G.; Xu, Q.; Liu, Y.; Shen, G.; Yao, Y.; Xu, J.; and Zhang, X. 2022. Mirror: Model inversion for deep learning network with high fidelity. In *Proceedings of the 29th Network and Distributed System Security Symposium*.
- Bai, L.; Hu, H.; Ye, Q.; Li, H.; Wang, L.; and Xu, J. 2024. Membership Inference Attacks and Defenses in Federated Learning: A Survey. *ACM Computing Surveys*, 57(4): 1–35.
- Bansal, A.; Nanduri, A.; Castillo, C. D.; Ranjan, R.; and Chellappa, R. 2017. Umdfaces: An annotated face dataset for training deep networks. In *2017 IEEE international joint conference on biometrics (IJCB)*, 464–473. IEEE.
- Boumal, N. 2023. *An introduction to optimization on smooth manifolds*. Cambridge University Press.
- Carlini, N.; Chien, S.; Nasr, M.; Song, S.; Terzis, A.; and Tramèr, F. 2022. Membership Inference Attacks From First Principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, 1897–1914.
- Chen, S.; Kahla, M.; Jia, R.; and Qi, G.-J. 2021. Knowledge-Enriched Distributional Model Inversion Attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 16178–16187.
- Fredrikson, M.; Jha, S.; and Ristenpart, T. 2015. Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In Ray, I.; Li, N.; and Kruegel, C., eds., *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, 1322–1333. ACM.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-Excitation Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hu, L.; Yan, H.; Peng, Y.; Hu, H.; Wang, S.; and Li, J. 2024. VAE-Based Membership Cleanser Against Membership Inference Attacks. *IEEE Transactions on Dependable and Secure Computing*.
- Ioffe, S.; and Szegedy, C. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 448–456. Lille, France: PMLR.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Khosla, A.; Jayadevaprakash, N.; Yao, B.; and Fei-Fei, L. 2011. Novel Dataset for Fine-Grained Image Categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*. Colorado Springs, CO.
- Krizhevsky, A. 2014. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*.
- Li, H.; Ye, Q.; Hu, H.; Li, J.; Wang, L.; Fang, C.; and Shi, J. 2023. 3dfed: Adaptive and extensible framework for covert backdoor attack in federated learning. In *2023 IEEE Symposium on Security and Privacy (SP)*, 1893–1907. IEEE.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Martens, J. 2020. New Insights and Perspectives on the Natural Gradient Method. *J. Mach. Learn. Res.*, 21(1).
- Naeem, M. F.; Oh, S. J.; Uh, Y.; Choi, Y.; and Yoo, J. 2020. Reliable fidelity and diversity metrics for generative models. In *International Conference on Machine Learning*, 7176–7185. PMLR.
- Nguyen, N.-B.; Chandrasegaran, K.; Abdollahzadeh, M.; and Cheung, N.-M. 2023. Re-Thinking Model Inversion Attacks Against Deep Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16384–16393.
- Qi, G.; Chen, Y.; Mao, X.; Hui, B.; Li, X.; Zhang, R.; and Xue, H. 2023. Model Inversion Attack via Dynamic Memory Learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, 5614–5622.
- Shao, H.; Kumar, A.; and Thomas Fletcher, P. 2018. The Riemannian Geometry of Deep Generative Models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.

Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Song, J.; and Namiot, D. 2022. A Survey of the Implementations of Model Inversion Attacks. In *International Conference on Distributed Computer and Communication Networks*, 3–16. Springer.

Struppek, L.; Hintersdorf, D.; De Almeida Correia, A.; Adler, A.; and Kersting, K. 2022. Plug & Play Attacks: Towards Robust and Flexible Model Inversion Attacks. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 20522–20545. PMLR.

Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Wang, B.; and Ponce, C. R. 2021. The geometry of deep generative image models and its applications. *arXiv preprint arXiv:2101.06006*.

Wang, K.-C.; FU, Y.; Li, K.; Khisti, A.; Zemel, R.; and Makhzani, A. 2021a. Variational Model Inversion Attacks. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 9706–9719. Curran Associates, Inc.

Wang, Q.; Zhang, P.; Xiong, H.; and Zhao, J. 2021b. Face.evoLve: A High-Performance Face Recognition Library. *arXiv preprint arXiv:2107.08621*.

Yan, H.; Li, S.; Wang, Y.; Zhang, Y.; Sharif, K.; Hu, H.; and Li, Y. 2022. Membership inference attacks against deep learning models via logits distribution. *IEEE Transactions on Dependable and Secure Computing*, 20(5): 3799–3808.

Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.

Zhang, Y.; Jia, R.; Pei, H.; Wang, W.; Li, B.; and Song, D. 2020. The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Zhao, Z.-Q.; Cheung, Y.-m.; Hu, H.; and Wu, X. 2016. Corrupted and occluded face recognition via cooperative sparse representation. *Pattern Recognition*, 56: 77–87.

Zou, D.-N.; Zhang, S.-H.; Mu, T.-J.; and Zhang, M. 2020. A new dataset of dog breed images and a benchmark for finegrained classification. *Computational Visual Media*, 6: 477–487.