

# Dual-View Inference Attack: Machine Unlearning Amplifies Privacy Exposure

Lulu Xue<sup>1</sup>, Shengshan Hu<sup>1</sup>, Linqiang Qian<sup>2</sup>, Peijin Guo<sup>1</sup>, Yechao Zhang<sup>3</sup>, Minghui Li<sup>4</sup>\*, Yanjun Zhang<sup>5</sup>, Dayong Ye<sup>6</sup>, Leo Yu Zhang<sup>7</sup>

<sup>1</sup> School of Cyber Science and Engineering, Huazhong University of Science and Technology

<sup>2</sup> Institute for Network Sciences and Cyberspace, Tsinghua University

<sup>3</sup> College of Computing and Data Science, Nanyang Technological University

<sup>4</sup> School of Software Engineering, Huazhong University of Science and Technology

<sup>5</sup> School of Computer Science, University of Technology Sydney

<sup>6</sup> Faculty of Data Science, City University of Macau

<sup>7</sup> School of Information and Communication Technology, Griffith University

{lluxue,hushengshan,gpj,minghui} @hust.edu.cn,

qlq25@mails.tsinghua.edu.cn, yech.zhang@gmail.com, Yanjun.Zhang@uts.edu.au, dayongye@outlook.com,

leo.zhang@griffith.edu.au

## Abstract

Machine unlearning is a newly popularized technique for removing specific training data from a trained model, enabling it to comply with data deletion requests. While it protects the rights of users requesting unlearning, it also introduces new privacy risks. Prior works have primarily focused on the privacy of data that has been unlearned, while the risks to retained data remain largely unexplored. To address this gap, we focus on the privacy risks of retained data and, for the first time, reveal the vulnerabilities introduced by machine unlearning under the dual-view setting, where an adversary can query both the original and the unlearned models. From an information-theoretic perspective, we introduce the concept of privacy knowledge gain and demonstrate that the dual-view setting allows adversaries to obtain more information than querying either model alone, thereby amplifying privacy leakage. To effectively demonstrate this threat, we propose DVIA, a Dual-View Inference Attack, which extracts membership information on retained data using black-box queries to both models. DVIA eliminates the need to train an attack model and employs a lightweight likelihood ratio inference module for efficient inference. Experiments across different datasets and model architectures validate the effectiveness of DVIA and highlight the privacy risks inherent in the dual-view setting.

## 1 Introduction

Machine unlearning (Bourtoule et al. 2021; Graves, Nagisetty, and Ganesh 2021) is an advanced technique designed to remove the information of specific training data points from the target model. While machine unlearning is often considered a privacy-preserving technique because it allows users to remove their private data from a trained model (Tang et al. 2023; Zhang et al. 2023; Sai et al. 2024), it also introduces new privacy risks. Several studies have examined the privacy leakage of forgotten data during the unlearning process. For example, (Hu et al. 2024a; Bertran




Figure 1: An illustration of the dual-view setting, showing that accessing two models can result in greater privacy leakage than accessing a single model alone.

et al. 2024) show that adversaries may be able to partially reconstruct deleted samples, while (Chen et al. 2021; Lu et al. 2022) focus on inferring which specific data points have been unlearned. These studies typically treat users of unlearning APIs as privacy victims, but they overlook a critical threat: **such users may also exploit the unlearning interface as an attack vector to compromise the privacy of retained data (non-forgotten data).**

This threat receives limited attention. Although (Carlini et al. 2022b) introduce the “privacy onion effect,” showing that removing certain data may inadvertently increase the privacy risk of retained samples, and (Hayes et al. 2025) point out that unlearned models may overfit more severely to the retained data, these analyses focus solely on the unlearned model and remain constrained to the traditional single-model inference setting. They fail to consider a more realistic and concerning scenario in which users who initiate unlearning requests typically access the models before and after unlearning, referred to as the original model and the unlearned model, respectively. This access is a reasonable requirement, as users need to verify whether their data are successfully removed (Xu et al. 2024a). However, it also introduces a new privacy attack surface against retained data.

\*Corresponding Author.

In this work, we uncover this previously overlooked risk dimension and define this setting as the dual-view scenario, where the two “views” correspond to access to the original and the unlearned models, as illustrated in Fig. 1. Our analysis shows that in this setting, attackers leverage the responses from both models to extract richer membership information and significantly enhance inference on retained data.

Specifically, we adopt an information-theoretic perspective to analyze the additional information an attacker can obtain when accessing both the original and unlearned models, compared to accessing a single model. We define this additional information as the *Privacy Knowledge Gain* (Gain). A positive Gain indicates a reduction in the adversary’s uncertainty about the membership status of target data, thereby implying increased privacy leakage. We further analyze the underlying cause of a positive Gain and find that it results from the asymmetric impact of unlearning on model predictions. When a set of data is removed, the model’s predictions for non-forgotten (member) data remain largely stable due to memorization, whereas its predictions for unseen (non-member) data are more likely to change.

This asymmetry is especially exploitable in the dual-view setting, where adversaries compare the outputs of the original and unlearned models to more effectively infer membership, thereby amplifying the privacy risks of machine unlearning. Building on this insight, we propose **DVIA** (Dual-View Inference Attack), the first membership inference attack for retained dataset tailored to the dual-view setting. DVIA leverages the asymmetric impact introduced by unlearning to improve inference performance. It introduces a new metric, Unlearning Confidence Difference (UCD), which quantifies the impact of unlearning on target dataset. In addition, DVIA incorporates a likelihood ratio-based inference module that requires no training and uses a fixed threshold, enabling efficient and accurate inference without auxiliary models or threshold tuning.

In summary, our contributions are as follows: 1) We reveal a previously overlooked privacy threat in machine unlearning, where users initiating unlearning requests gain a significant advantage in inferring the membership of retained data by accessing both the original and the unlearned models. 2) We formalize this threat as a dual-view setting, analyze its amplified privacy risks, and propose DVIA, the first membership inference attack specifically designed for retained data in this setting. 3) We conduct extensive experiments across diverse datasets and model architectures, showing that DVIA achieves strong performance in uncovering membership information under unlearning scenarios.

## 2 Related Works

### 2.1 Machine Unlearning

Machine unlearning (Doughan and Itani 2024; Lindstrom 2024; Yang et al. 2024) is the process of initiating a reverse learning task to eliminate the information of specific data points on the target model. Machine unlearning can be divided into exact unlearning (Yan et al. 2022; Hu et al. 2024b) and approximate unlearning (Xu et al. 2024b; Graves, Nagisetty, and Ganesh 2021). Exact unlearning in-

volves training model parameters from scratch using the retained dataset, and the retrain method, which retains the original settings for retraining, offers the best unlearning results. While exact unlearning ensures effective data removal (Liu et al. 2023; Fan et al. 2023), it comes at a high computational cost (Liu et al. 2024a; Hu et al. 2024c; Wang et al. 2024; Liu et al. 2024b). To address this challenge, many approximate unlearning algorithms have been proposed (Golatkar, Achille, and Soatto 2020; Graves, Nagisetty, and Ganesh 2021; Thudi et al. 2022; Fan et al. 2023; Liu et al. 2023; Zhang et al. 2024). Although these methods introduce some inaccuracies during the unlearning process, they demonstrate the feasibility of quickly eliminating the influence of forgotten data on model parameters.

### 2.2 Membership Inference Attacks

Membership inference attacks (MIAs) (Shokri et al. 2017; Carlini et al. 2022a; Chen et al. 2022) aim to determine whether a given sample was part of a model’s training set. Most existing approaches rely on shadow models: adversaries train multiple models on auxiliary data to simulate the target model’s behavior, then use the resulting labeled data to train a binary classifier (Shokri et al. 2017; Chen et al. 2022) or to select a decision threshold (Choquette-Choo et al. 2021; Li and Zhang 2021). These methods often require careful design and threshold tuning. Recently, (Carlini et al. 2022a) proposes a likelihood ratio-based approach that avoids explicit attack model training and complex threshold selection. However, it conducts inference on a per-sample basis and requires generating multiple shadow models and indicators for each target sample, leading to high computational overhead.

**Membership inference in machine unlearning.** Membership inference is commonly used as a tool to evaluate the effectiveness of forgetting (Liu et al. 2023; Fan et al. 2023). Recent studies (Chen et al. 2021; Hayes et al. 2025) also explore its use in exposing privacy risks introduced by unlearning, focusing on inferring deleted data. These studies typically treat users of unlearning APIs as victims, aiming to recover forgotten samples. However, prior studies (Carlini et al. 2022b; Hayes et al. 2025) suggest that unlearning may amplify privacy leakage, yet they do not propose dedicated attacks for this setting and overlook the increased information advantage that unlearning grants to adversaries.

## 3 Privacy Analysis

### 3.1 Threat Model

We consider a user invoking the unlearning API as a curious adversary aiming to infer membership. The adversary requests the removal of a forgotten dataset  $D_f$  and queries both the original model  $\theta^o$  and the unlearned model  $\theta^u$ , forming a dual-view setting. This setting is realistic, since users can often access both models to confirm data removal (Xu et al. 2024a; Xue et al. 2025). Using this access, the adversary aims to infer the membership status of samples in a target dataset  $D_t$ , where samples from the retained set  $D_{\text{retain}}$  are considered members and those from the validation set  $D_{\text{val}}$  are non-members. We refer to this as the *non-*

forgetting inference scenario. Following (Chen et al. 2021; Hayes et al. 2025), we assume a black-box setting where the adversary knows the data distribution  $\mathbb{D}$ , as well as the training and unlearning procedures and the model architecture, but not the model parameters.

### 3.2 Privacy Knowledge Gain

To analyze the privacy risks introduced by this dual-view, we adopt an information-theoretic perspective and define the adversary’s improved privacy knowledge as *privacy knowledge gain*. The specific definition is as follows.

**Definition 1** (Privacy Knowledge Gain). *Let  $s_o(D_t)$  and  $s_u(D_t)$  denote the observable behaviors of the original and unlearned models on the target dataset  $D_t$ , and let  $s(D_t) \in \{s_o(D_t), s_u(D_t)\}$  denote a single-view observation. The privacy knowledge gain measures the additional membership information revealed by dual-view access, defined as:*

$$\text{Gain}(D_t) = I(M_{D_t}; s_o(D_t), s_u(D_t)) - I(M_{D_t}; s(D_t)),$$

where  $M_{D_t} \in \{0, 1\}^{|D_t|}$  denotes the membership status of each sample in  $D_t$ , with 1 indicating member data.  $I(\cdot; \cdot)$  denotes the mutual information (Shannon 1948), which measures the reduction in uncertainty of one variable given knowledge of another.

When  $\text{Gain}(D_t) > 0$ , it indicates that the adversary has acquired additional information about the membership status of samples in  $D_t$ , reducing uncertainty and enabling more effective privacy inference. Our goal is to analyze the conditions under which  $\text{Gain} > 0$  holds in the dual-view setting, in order to reveal how this setting amplifies privacy leakage risks. Specifically, we investigate the necessary and sufficient conditions for  $\text{Gain}(D_t) > 0$  under machine unlearning, and validate them with practical examples. To facilitate this analysis, we first define the behavioral impact of machine unlearning on the  $D_t$  as follows:

**Definition 2.** *The behavioral impact of machine unlearning on the target dataset  $D_t$  is defined as*

$$\delta(D_t) = 1 - \text{sim}(s_o(D_t), s_u(D_t)),$$

where  $s_o(D_t)$  and  $s_u(D_t)$  denote the observable behaviors of the original and unlearned models on  $D_t$ , respectively, and  $\text{sim}(\cdot, \cdot)$  is a similarity function with larger values indicating higher behavioral closeness.

Intuitively, if  $\delta(D_t)$  is independent of the membership status of  $D_t$ , this implies that the behavioral change between the two models does not statistically differ for members and non-members. Therefore, the unlearned model does not expose any additional membership information beyond what is already present in the original model. Conversely, if  $\delta(D_t)$  exhibits statistical differences conditioned on membership status, then it must encode additional information about whether a sample belongs to the retained dataset. This additional signal can be exploited by an adversary to improve membership inference, thereby resulting in a positive privacy gain. Based on this definition, we formalize the notion of privacy knowledge gain under membership conditions in the following theorem.

**Theorem 1.** *For the target dataset  $D_t$ , The privacy knowledge gain satisfies:*

$$\begin{aligned} \text{Gain}(D_t) > 0 \quad &\text{if and only if} \\ \mathbb{P}(\delta(D_t) \mid M_{D_t} = 1) &\neq \mathbb{P}(\delta(D_t) \mid M_{D_t} = 0), \end{aligned}$$

where  $\delta(D_t)$  is the behavioral impact for  $D_t$ ,  $M_{D_t}$  is the membership status of  $D_t$ , and  $\mathbb{P}(\cdot \mid \cdot)$  denotes conditional probability.

The proof is provided in the appendix. According to Theorem 1, the privacy knowledge of the target dataset  $D_t$  is amplified by machine unlearning if the observable behavioral impact of unlearning on member and non-member data differs. To quantify this behavioral impact, we adopt the influence function proposed in (Koh and Liang 2017), which estimates how adding a training point  $\mathbf{z}$  affects the model’s prediction on a given test point  $\mathbf{z}_{\text{test}}$ . The influence is computed as:

$$\mathcal{I}_{\text{IF}}(\mathbf{z}, \mathbf{z}_{\text{test}}) = -\nabla_{\theta} \mathcal{L}(\mathbf{z}_{\text{test}}, \theta)^{\top} H_{\theta}^{-1} \nabla_{\theta} \mathcal{L}(\mathbf{z}, \theta), \quad (1)$$

where  $\theta$  denotes the model parameters,  $\mathcal{L}$  is the loss function, and  $H_{\theta}$  is the Hessian matrix of second-order derivatives of  $\mathcal{L}$  with respect to  $\theta$ . In the scenario of removing data points from the training set, we use  $-\mathcal{I}_{\text{IF}}$  to measure the impact of unlearning. Specifically,  $-\mathcal{I}_{\text{IF}}(D_f, \mathbf{z}_{\text{test}})$  measures the behavioral change of the model on  $\mathbf{z}_{\text{test}}$  resulting from unlearning the dataset  $D_f$ , where  $\mathbf{z}_{\text{test}}$  denotes an arbitrary sample from the target dataset  $D_t$ . Due to the high computational cost of directly calculating the Hessian matrix, we approximate it using the woodfisher information matrix (Singh and Alistarh 2020) in this paper.

Based on the influence function, we perform a straightforward experiment by training a convolutional neural network model, ALLCNN (Springenberg et al. 2014), on the CIFAR10 dataset and subsequently unlearning 5% of the data points. We sample 100 points from the non-forgotten set and validation set respectively to calculate their influence scores, and display the top 3 images with the largest absolute influence scores, as shown in Figure 2. It can be observed that machine unlearning typically has a smaller impact on non-forgotten data points, while exerting a larger influence on data the model has never encountered. This is because member samples have been well memorized by the model; even after the unlearning process, the model’s predictions on these samples remain stable, resulting in lower absolute impact scores. In contrast, for unseen data, the removal of part of the training set leads to the loss of certain feature representations, weakening the model’s generalization ability and causing greater shifts in its predictions on these samples. This asymmetry in influence leads to a statistical divergence in the model behavior shifts between members and non-members, which ultimately amplifies the membership privacy risk.

## 4 Dual-View Inference Attack

This section introduces **DVIA** (Dual-View Inference Attack), the attack that exploits the amplified privacy risks enabled by access to both the original and unlearned models.

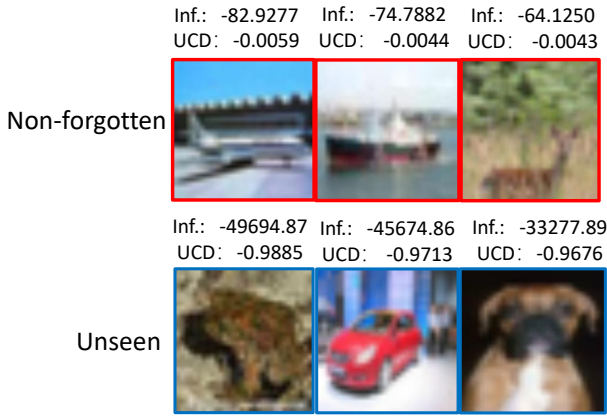


Figure 2: A depiction of the unlearning influence on the target model’s non-forgotten and unseen data. “Inf.” refers to the influence score, and UCD is the metric we propose to measure the behavioral impact in black-box settings.

DVIA comprises two components: the Unlearning Confidence Difference (UCD) as the attack metric, and a likelihood ratio inference module for membership prediction.

#### 4.1 Unlearning Confidence Difference


In this section, we design a behavioral impact–based attack metric under the DVIA framework. Our core insight is that machine unlearning induces notable differences in the model’s predictive behavior for member versus non-member data, as analyzed in Section 3.2. This paper focuses on standard image classification tasks, where model behavior is primarily reflected in the confidence scores assigned to image labels. Therefore, we propose the Unlearning Confidence Difference (UCD) metric, which focuses on changes in the model’s confidence when predicting the label of a given input. By leveraging dual queries to both the original and unlearned models, UCD reveals the privacy risks posed by machine unlearning on data that has not been forgotten. The details of the UCD metric are provided in Definition 3.

**Definition 3.** *The unlearning confidence difference (UCD)  $\Delta(\mathbf{x}, y)$  is defined as :*


$$\Delta(\mathbf{x}, y) = \theta^u(\mathbf{x})_y - \theta^o(\mathbf{x})_y,$$

where  $\theta^o$  represents the original trained model and  $\theta^u$  represents the unlearned model,  $\theta(\mathbf{x})_y$  denotes the probability that the model  $\theta$  maps the input  $\mathbf{x}$  to correct label  $y$ .

Figure 3 illustrates the distribution of UCD scores for member and non-member data, as well as the impact of unlearning a specific data group on both. We observe a clear separation between the two distributions: after forgetting a group of data, the model retains high confidence in the samples it has memorized, while its confidence in previously unseen data fluctuates more noticeably, which aligns with our analysis in Section 3.2. Moreover, the behavioral patterns reflected by the UCD scores in the black-box setting closely mirror the influence scores computed in the white-box setting. This consistency demonstrates that the UCD metric ef-



(a) Influence Score



(b) UCD

Figure 3: Density plots of influence score and UCD. It can be observed that the distributions of both show similarity.

---

#### Algorithm 1: Dual-View Inference Attack

---

- 1: **Input:** Data distribution  $\mathbb{D}$ , target dataset  $D_{\text{target}}$ .
  - 2: **Require:** Initialized model  $\theta$ , training algorithm  $A$ , unlearning algorithm  $U$ , number of shadow datasets  $K$ , decision threshold  $\Lambda$ .
  - 3:  $S_{\text{mem}} \leftarrow \emptyset$   $\triangleright$  Set to store UCD metrics for member
  - 4:  $S_{\text{non}} \leftarrow \emptyset$   $\triangleright$  Set to store UCD metrics for non-member
  - 5: **for**  $k = 1$  **to**  $K$  **do**
  - 6:   Sample shadow datasets:  $D_k^{\text{mem}}, D_k^{\text{non}}$  from  $\mathbb{D}$
  - 7:   Use  $D_k^{\text{mem}}, D_k^{\text{non}}, A$  and  $U$  to get the original model  $\theta_k^o$  and the unlearned model  $\theta_k^u$
  - 8:    $S_{\text{mem}} \leftarrow S_{\text{mem}} \cup \{\phi(\Delta(\mathbf{x}, y)) \mid (\mathbf{x}, y) \in D_k^{\text{mem}}\}$
  - 9:    $S_{\text{non}} \leftarrow S_{\text{non}} \cup \{\phi(\Delta(\mathbf{x}, y)) \mid (\mathbf{x}, y) \in D_k^{\text{non}}\}$
  - 10: **end for**
  - 11:  $\mu_{\text{mem}}, \sigma_{\text{mem}}^2 \leftarrow \text{mean}(S_{\text{mem}}), \text{var}(S_{\text{mem}})$
  - 12:  $\mu_{\text{non}}, \sigma_{\text{non}}^2 \leftarrow \text{mean}(S_{\text{non}}), \text{var}(S_{\text{non}})$
  - 13:  $S_t \leftarrow \{\phi(\Delta(\mathbf{x}, y)) \mid (\mathbf{x}, y) \in D_{\text{target}}\}$
  - 14:  $L(\alpha) \leftarrow \frac{p(\alpha | \mathcal{N}(\mu_{\text{mem}}, \sigma_{\text{mem}}^2))}{p(\alpha | \mathcal{N}(\mu_{\text{non}}, \sigma_{\text{non}}^2))}$  for each  $\alpha \in S_t$
  - 15: **Return:**  $\{1 \text{ if } L(\alpha) > \Lambda \text{ else } 0 \mid \alpha \in S_t\}$
- 

fectively captures the differential impact of unlearning on member and non-member data, as shown in Figures 2–3.

#### 4.2 Likelihood Ratio Inference

Figure 3 shows that there is a clear distinction between the UCD distributions of members and non-members, with member data being more likely to fall within the member distribution rather than the non-member distribution. Based

on this, we design the likelihood ratio module to perform non-forgotten inference by leveraging this noticeable distribution difference. Specifically, in a machine unlearning scenario, an attacker could use shadow models to obtain the distributions of member and non-member data, then calculate the likelihood ratios of target samples with respect to these distributions to determine their membership status  $b$ , classifying them as either 0 or 1.

However, a key challenge lies in constructing the member and non-member distributions to support membership inference tasks. Prior research (Carlini et al. 2022a) has shown that Gaussian fitting is a common and effective strategy, as it requires only four statistical parameters, namely the means and standard deviations of member and non-member data, to approximate the distributions. However, the domain of the UCD distribution is  $(-1, 1)$ , which is inconsistent with the support of the Gaussian distribution. Directly fitting a Gaussian distribution in the original space leads to boundary truncation issues, making such an approach inappropriate.

To address this, we consider applying a monotonic transformation to the UCD metric to expand its domain. Although (Carlini et al. 2022a) has proposed  $\phi_0(\Delta) = \frac{\Delta}{1-\Delta}$  as a common transformation function, this function is not well-defined for negative values, limiting its applicability to the UCD metric. Therefore, we propose the following transformation function:

$$\phi(\Delta) = \log\left(\frac{1+\Delta}{1-\Delta}\right). \quad (2)$$

which is strictly monotonic on  $x \in (-1, 1)$  and maps this interval to  $(-\infty, \infty)$ . This transformation not only aligns with the support of the Gaussian distribution but also maintains local linearity near  $x = 0$ , preserving the separability of the original UCD scores. Therefore, it provides a sound basis for subsequent Gaussian modeling, which underpins our attack design described in Algorithm 1. In a nutshell, the attacker first samples the shadow datasets from  $\mathbb{D}$  and trains shadow original and shadow unlearned models from these datasets (Lines 6-7). The attacker then collects UCD from the two shadow models for both member and non-member data (Lines 8-9) and use  $\phi$  to fit them into two Gaussian distributions (Lines 11-12). For any targeted data point  $(\mathbf{x}, y)$ , the attacker calculates  $\phi(\Delta(\mathbf{x}, y))$  (Line 13) and determines its likelihood ratio  $L(\alpha)$  relative to the member and non-member distributions (Line 14). We assume  $\Lambda = 1$  by default, which is a reasonable choice as it implies that a sample is classified as a member if its membership probability exceeds its non-membership probability.

**Advantages.** Compared to the inference approaches based on attack models (Shokri et al. 2017; Chen et al. 2021), our module does not require the selection and training of a dedicated attack model. Furthermore, unlike threshold-based methods (Lu et al. 2022), our module avoids the complexity of carefully tuning an appropriate threshold; a default threshold of 1 is sufficient for effective inference.

**Comparison with (Hayes et al. 2025).** While both our method and U-LIRA employ likelihood ratio-based inference to perform membership inference in the unlearning setting, they differ fundamentally in focus and granularity. U-

LIRA targets forgotten data using a confidence-based metric and performs inference at the single-sample level, requiring multiple shadow models to estimate membership distributions per target instance. In contrast, DVIA is the first to investigate the privacy risks for non-forgotten data. We introduce the UCD metric and focus on group-level separability between members and non-members by leveraging statistical patterns across multiple samples. This design enables efficient inference using only a few shadow models and reveals the systemic privacy leakage introduced by unlearning.

## 5 Experiment

### 5.1 Experimental Setup

We conduct our experiments in Python 3.8 using PyTorch (Paszke et al. 2019). Our evaluations are performed on two commonly used datasets, CIFAR10 (Krizhevsky, Hinton et al. 2009) and SVHN (Netzer et al. 2011). We divide the datasets into four parts to create the training set, validation set, shadow training set, and shadow validation set, with proportions of 50%, 10%, 20%, and 20%, respectively. For the models, we use ResNet18 (He et al. 2016) and the convolutional neural network ALLCNN (Springenberg et al. 2014), both of which are widely adopted in unlearning tasks (Golatkar, Achille, and Soatto 2020; Liu et al. 2023; Fan et al. 2023; Zhang et al. 2024). Following (Chen et al. 2021), we train 4 shadow original models and 16 shadow unlearned models for each shadow original model. We assume that the attacker invokes the unlearning APIs to request the removal of 5% of the training data. We use the best exact unlearning method, Retrain, as the default unlearning approach, and explore approximate unlearning methods in Section 5.4. Following (Shokri et al. 2017; Hayes et al. 2025), we use balanced accuracy as the primary evaluation metric, as it effectively measures the overall performance of the inference attack. We construct an evaluation dataset of 1,000 points, with half randomly sampled from the non-forgotten set and the other half from the validation set.

### 5.2 Comparative Experiment

Our comparative study includes the following state-of-the-art membership inference attacks: LIRA (Carlini et al. 2022a), SMIA (Shokri et al. 2017), and QMIA (Li and Zhang 2021). We evaluate their performance on both the original and unlearned models, denoted with the suffixes “-o” (original) and “-u” (unlearned), respectively. In addition, we consider advanced attacks that were originally designed to infer the membership of forgotten data, including MIAoMU (Lu et al. 2022), U-LIRA (Hayes et al. 2025), and U-MIA (Chen et al. 2021), and apply them to non-forgetting inference scenario. We use these methods with their default settings. The results are shown in Figure 4. It can be observed that LIRA, SMIA, and QMIA perform slightly better on the unlearned model than on the original model, and our method significantly outperforms all existing methods.

### 5.3 Ablation Experiments

**Different Attack Modules.** We replace the likelihood ratio inference (LRI) module with a decision tree (DT) (Rana

	CIFAR10		SVHN	
	ALLCNN	ResNet18	ALLCNN	ResNet18
LRI	0.9141	0.9386	0.8222	0.8210
DT	0.8665	0.8525	0.7321	0.7601
Threshold	0.8612	0.875	0.7422	0.7602

Table 1: Evaluation with different attack modules.

et al. 2014) and a threshold method (Li and Zhang 2021). As shown in Table 1, LRI achieves superior attack performance. This is because likelihood ratios typically capture distributional differences more effectively, leading to more accurate membership inference.

**Different Attack Metrics.** We evaluate the attack performance of the UCD metric, original model confidence  $\text{Conf}_o$ , and unlearned model confidence  $\text{Conf}_u$ , with the attack model using a likelihood ratio inference module. The results are shown in Table 2. It is evident that the inference attack based on the UCD metric performs the best, further highlighting the privacy risks posed by unlearning APIs.

#### 5.4 Approximate Unlearning

We evaluate our method with approximate unlearning from the perspectives of comp. overhead and attack effect. We select several advanced approximate unlearning methods and apply them to non-forgotten inference, including Gradient Ascent (GA) (Golatkar, Achille, and Soatto 2020; Graves, Nagisetty, and Ganesh 2021; Neel, Roth, and

	CIFAR10		SVHN	
	ALLCNN	ResNet18	ALLCNN	ResNet18
UCD	0.9141	0.9386	0.8222	0.8210
$\text{Conf}_o$	0.8221	0.8420	0.6874	0.6995
$\text{Conf}_u$	0.8330	0.8575	0.6960	0.7018

Table 2: Evaluation of attack metrics.

	CC	CR	SC	SR
(5%,5%)	0.9141	0.9386	0.8222	0.8210
(5%,10%)	0.9090	0.9372	0.8134	0.7956
(10%,5%)	0.9128	0.9344	0.8075	0.8165
(10%,10%)	0.9220	0.9425	0.8212	0.8182

Table 3: Impact of unlearning proportions in target and shadow models. Each  $(a\%, b\%)$  pair denotes the unlearning ratios of the target and shadow models, respectively.

Sharifi-Malvajardi 2021), Influence Unlearning (IU) (Liu et al. 2023; Fan et al. 2023) and Salun (Fan et al. 2023).

**Comp. Overhead.** Figure 7 (a) presents the comp. overhead of training unlearned shadow models using different unlearning methods. It is evident that approximate unlearning is significantly more efficient than exact unlearning.

**Attack Evaluation.** Figures 4 and 5 indicate that membership inference attacks are generally less effective under approximate unlearning compared to retraining unlearning. This is because approximate unlearning often has a greater impact on non-forgotten data (Graves, Nagisetty, and Ganesh 2021). Although DVIA, based on approximate unlearning, is less effective than exact unlearning methods in attack performance, it still outperforms comparative membership inference methods, highlighting the elevated privacy risks in the dual-view setting. In future work, we aim to further explore the trade-off between attack effectiveness and comp. efficiency in non-forgotten inference.

#### 5.5 Further Discussions

**Impact of Unlearning Ratio Knowledge.** As shown in Table 3, the attack remains effective even when the attacker does not know the exact unlearning ratio. This is because machine unlearning significantly affects non-members while having minimal impact on members, and this trend holds robustly across different unlearning ratio settings.

**Defense.** We implement differential privacy (Dwork 2006; Abadi et al. 2016) by using Opacus (Yousefpour et al. 2021) and setting the noise\_multiplier to 1.0. The experimental results are shown in Table 4. Although differential privacy can defend against our attacks, it may make the model unusable.

**More Metrics.** We present additional comparisons using AUC. As shown in Figure 6, our method outperforms others. Additional metrics can be found in the appendix.

**Impact of Different Unlearning Ratios.** Figure 7 (b) illustrates the performance of DVIA under varying unlearning ratios. Notably, even with as little as 2.5% of data unlearned, DVIA achieves strong performance.

**Evaluation on CIFAR100 (Krizhevsky, Hinton et al.**

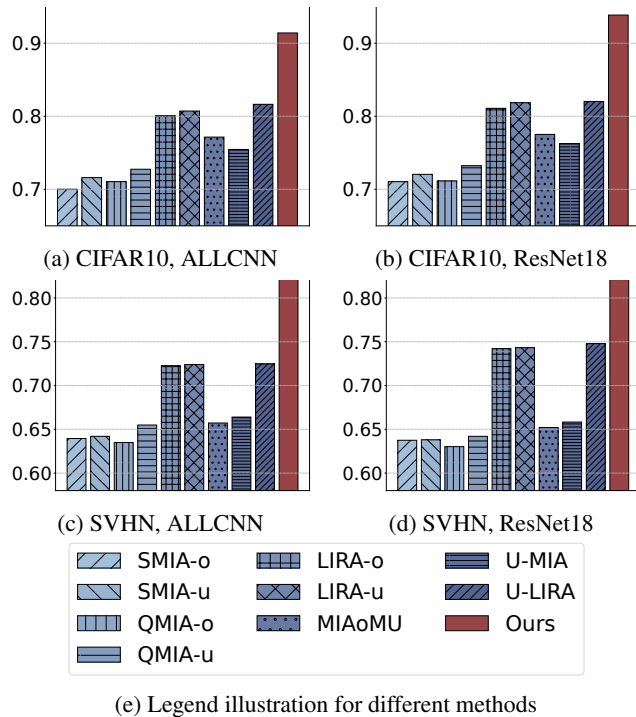


Figure 4: Evaluation across different datasets and models, where the blue bars represent the comparative methods and the red bars represent our approach.

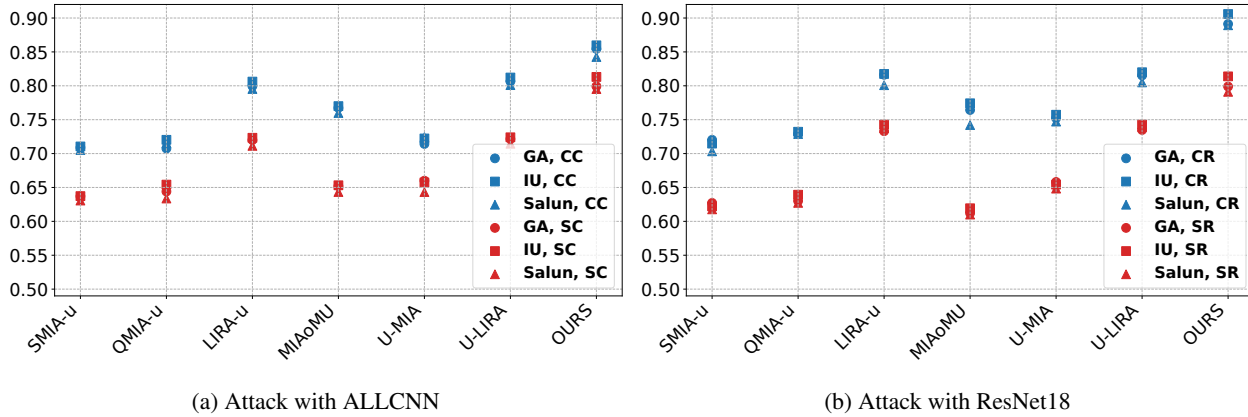


Figure 5: The evaluation of attack performance under different unlearning methods.

	Drop (%)	Retrain	GA	IU	Salun
CC	32.64	0.5248	0.5135	0.5293	0.5250
CR	33.14	0.5293	0.5351	0.4925	0.5162
SC	42.81	0.5128	0.4919	0.5048	0.5043
SR	39.73	0.5151	0.5103	0.5103	0.5120

Table 4: Evaluation under differential privacy. Drop indicates the drop in model accuracy on the validation set.




Figure 6: The AUC evaluation across different methods.

**2009**). Following (Carlini et al. 2022a), we evaluate DVIA on CIFAR100, where it achieves 0.8975 accuracy on All-CNN and 0.9444 on ResNet18, demonstrating strong performance.




Figure 7: More evaluation results.

## 6 Conclusion, Limitation, and Future

In this paper, we challenge the notion that machine unlearning ensures privacy and, for the first time, reveal the privacy risks it poses to non-forgotten data through an information-theoretic lens. We introduce the novel dual-view setting, where adversaries can query both the original and the unlearned models, and demonstrate that such dual access significantly amplifies privacy leakage. To formalize this advantage, we propose the notion of privacy information gain, which highlights the increased privacy exposure resulting from comparing the outputs of the two models. To exploit this vulnerability in practice, we propose DVIA, which uses black-box queries and a lightweight likelihood ratio module to infer membership of retained data. Our experiments across diverse datasets and model architectures validate the effectiveness of DVIA and underscore the privacy vulnerabilities introduced by unlearning in dual-view scenarios.

In terms of limitations, we focus solely on image classification tasks, which are the primary application area for machine unlearning. In future work, we plan to broaden our analysis to other domains such as natural language processing and graph learning. Furthermore, we intend to investigate the trade-offs between inference accuracy, computational efficiency, and privacy guarantees in approximate unlearning frameworks.

## Acknowledgments

Minghui’s work is supported in part by the National Natural Science Foundation of China (Grant No.62572206).

## References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security (CCS’16)*, 308–318.
- Bertran, M.; Tang, S.; Kearns, M.; Morgenstern, J.; Roth, A.; and Wu, Z. S. 2024. Reconstruction Attacks on Machine Unlearning: Simple Models are Vulnerable. *arXiv preprint arXiv:2405.20272*.
- Bourtole, L.; Chandrasekaran, V.; Choquette-Choo, C. A.; Jia, H.; Travers, A.; Zhang, B.; Lie, D.; and Papernot, N. 2021. Machine unlearning. In *Proceedings of the 2021 IEEE Symposium on Security and Privacy (SP’21)*, 141–159. IEEE.
- Carlini, N.; Chien, S.; Nasr, M.; Song, S.; Terzis, A.; and Tramer, F. 2022a. Membership inference attacks from first principles. In *Proceedings of the 2022 IEEE Symposium on Security and Privacy (SP’22)*, 1897–1914. IEEE.
- Carlini, N.; Jagielski, M.; Zhang, C.; Papernot, N.; Terzis, A.; and Tramer, F. 2022b. The privacy onion effect: Memorization is relative. *Proceedings of the 2022 Neural Information Processing Systems (NeurIPS’22)*, 35: 13263–13276.
- Chen, M.; Zhang, Z.; Wang, T.; Backes, M.; Humbert, M.; and Zhang, Y. 2021. When machine unlearning jeopardizes privacy. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security (CCS’21)*, 896–911.
- Chen, Y.; Shen, C.; Shen, Y.; Wang, C.; and Zhang, Y. 2022. Amplifying membership exposure via data poisoning. *Proceedings of the 2022 Neural Information Processing Systems (NeurIPS’22)*, 35: 29830–29844.
- Choquette-Choo, C. A.; Tramer, F.; Carlini, N.; and Papernot, N. 2021. Label-only membership inference attacks. In *Proceedings of the 2021 International conference on machine learning (ICML’21)*, 1964–1974. PMLR.
- Doughan, Z.; and Itani, S. 2024. Machine Unlearning, A Comparative Analysis. In *Proceedings of the 2024 International Conference on Engineering Applications of Neural Networks (EANN’2025)*, 558–568. Springer.
- Dwork, C. 2006. Differential privacy. In *Proceedings of the 2006 International colloquium on automata, languages, and programming (ICALP’06)*, 1–12. Springer.
- Fan, C.; Liu, J.; Zhang, Y.; Wei, D.; Wong, E.; and Liu, S. 2023. Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. *arXiv preprint arXiv:2310.12508*.
- Golatkar, A.; Achille, A.; and Soatto, S. 2020. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR’20)*, 9304–9312.
- Graves, L.; Nagisetty, V.; and Ganesh, V. 2021. Amnesiac machine learning. In *Proceedings of the 2021 AAAI Conference on Artificial Intelligence (AAAI’21)*, volume 35, 11516–11524.
- Hayes, J.; Shumailov, I.; Triantafillou, E.; Khalifa, A.; and Papernot, N. 2025. Inexact unlearning needs more careful evaluations to avoid a false sense of privacy. In *Proceedings of the 2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML’25)*, 497–519. IEEE.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the 2016 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR’16)*, 770–778.
- Hu, H.; Wang, S.; Dong, T.; and Xue, M. 2024a. Learn What You Want to Unlearn: Unlearning Inversion Attacks against Machine Unlearning. In *Proceedings of the 2024 IEEE Symposium on Security and Privacy (SP’24)*, 262–262. IEEE Computer Society.
- Hu, Z.; Zhang, Y.; Xiao, M.; Wang, W.; Feng, F.; and He, X. 2024b. Exact and efficient unlearning for large language model-based recommendation. *arXiv preprint arXiv:2404.10327*.
- Hu, Z.; Zhang, Y.; Xiao, M.; Wang, W.; Feng, F.; and He, X. 2024c. Exact and Efficient Unlearning for Large Language Model-based Recommendation. *arXiv preprint arXiv:2404.10327*.
- Koh, P. W.; and Liang, P. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 2017 International conference on machine learning (ICML’17)*, 1885–1894. PMLR.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, Z.; and Zhang, Y. 2021. Membership leakage in label-only exposures. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS’21)*, 880–895.
- Lindstrom, C. 2024. Exact Unlearning with Convex and Non-Convex Functions. *Preprints preprints202409*, 2061: v1.
- Liu, J.; Ram, P.; Yao, Y.; Liu, G.; Liu, Y.; SHARMA, P.; Liu, S.; et al. 2023. Model sparsity can simplify machine unlearning. In *Proceedings of the 2023 Neural Information Processing Systems (NeurIPS’23)*, volume 36.
- Liu, S.; Yao, Y.; Jia, J.; Casper, S.; Baracaldo, N.; Hase, P.; Xu, X.; Yao, Y.; Li, H.; Varshney, K. R.; et al. 2024a. Rethinking machine unlearning for large language models. *arXiv preprint arXiv:2402.08787*.
- Liu, Z.; Dou, G.; Tan, Z.; Tian, Y.; and Jiang, M. 2024b. Towards safer large language models through machine unlearning. *arXiv preprint arXiv:2402.10058*.
- Lu, Z.; Liang, H.; Zhao, M.; Lv, Q.; Liang, T.; and Wang, Y. 2022. Label-only membership inference attacks on machine unlearning without dependence of posteriors. *International Journal of Intelligent Systems*, 37(11): 9424–9441.
- Neel, S.; Roth, A.; and Sharifi-Malvajerdi, S. 2021. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, 931–962. PMLR.

- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, 4. Granada.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library.
- Rana, K. K.; et al. 2014. A survey on decision tree algorithm for classification. *International journal of Engineering development and research*, 2(1): 1–5.
- Sai, S.; Mittal, U.; Chamola, V.; Huang, K.; Spinelli, I.; Scardapane, S.; Tan, Z.; and Hussain, A. 2024. Machine unlearning: An overview of techniques, applications, and future directions. *Cognitive Computation*, 16(2): 482–506.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3): 379–423.
- Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017. Membership inference attacks against machine learning models. In *Proceedings of the 2017 IEEE symposium on security and privacy (SP'17)*, 3–18. IEEE.
- Singh, S. P.; and Alistarh, D. 2020. Woodfisher: Efficient second-order approximation for neural network compression. *Proceedings of the 2020 Neural Information Processing Systems (NeurIPS'20)*, 33: 18098–18109.
- Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2014. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.
- Tang, Y.; Cai, Z.; Liu, Q.; Zhou, T.; and Ni, Q. 2023. Ensuring User Privacy and Model Security via Machine Unlearning: A Review. *Computers, Materials & Continua*, 77(2).
- Thudi, A.; Deza, G.; Chandrasekaran, V.; and Papernot, N. 2022. Unrolling sgd: Understanding factors influencing machine unlearning. In *Proceedings of the IEEE 7th European Symposium on Security and Privacy (EuroS&P'22)*, 303–319. IEEE.
- Wang, Q.; Han, B.; Yang, P.; Zhu, J.; Liu, T.; and Sugiyama, M. 2024. Unlearning with Control: Assessing Real-world Utility for Large Language Model Unlearning. *arXiv preprint arXiv:2406.09179*.
- Xu, H.; Zhu, T.; Zhang, L.; and Zhou, W. 2024a. Really unlearned? verifying machine unlearning via influential sample pairs. *arXiv preprint arXiv:2406.10953*.
- Xu, J.; Wu, Z.; Wang, C.; and Jia, X. 2024b. Machine unlearning: Solutions and challenges. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
- Xue, L.; Hu, S.; Lu, W.; Shen, Y.; Li, D.; Guo, P.; Zhou, Z.; Li, M.; Zhang, Y.; and Zhang, L. Y. 2025. Towards Reliable Forgetting: A Survey on Machine Unlearning Verification, Challenges, and Future Directions. *arXiv preprint arXiv:2506.15115*.
- Yan, H.; Li, X.; Guo, Z.; Li, H.; Li, F.; and Lin, X. 2022. ARCANE: An Efficient Architecture for Exact Machine Unlearning. In *Proceedings of the 2022 International Joint Conferences on Artificial Intelligence (IJCAI'22)*, volume 6, 19.
- Yang, R.; He, M.; He, Z.; Qiu, Y.; and Huang, X. 2024. MUSO: Achieving Exact Machine Unlearning in Over-Parameterized Regimes. *arXiv preprint arXiv:2410.08557*.
- Yousefpour, A.; Shilov, I.; Sablayrolles, A.; Testuggine, D.; Prasad, K.; Malek, M.; Nguyen, J.; Ghosh, S.; Bharadwaj, A.; Zhao, J.; Cormode, G.; and Mironov, I. 2021. Opacus: User-Friendly Differential Privacy Library in PyTorch. *arXiv preprint arXiv:2109.12298*.
- Zhang, B.; Dong, Y.; Wang, T.; and Li, J. 2024. Towards Certified Unlearning for Deep Neural Networks. *arXiv preprint arXiv:2408.00920*.
- Zhang, H.; Nakamura, T.; Isohara, T.; and Sakurai, K. 2023. A review on machine unlearning. *SN Computer Science*, 4(4): 337.