

Distribution-Level Feature Distancing for Machine Unlearning: Towards a Better Trade-off Between Model Utility and Forgetting

Dasol Choi^{1,2}, Dongbin Na^{3*}

¹Yonsei University

²MODULABS

³POSTECH

dasolchoi@yonsei.ac.kr, dongbinna@postech.ac.kr

Abstract

With the explosive growth of deep learning applications and increasing privacy concerns, the right to be forgotten has become a critical requirement in various AI industries. For example, given a facial recognition system, some individuals may wish to remove their personal data that might have been used in the training phase. Unfortunately, deep neural networks sometimes unexpectedly leak personal identities, making this removal challenging. While recent machine unlearning algorithms aim to enable models to forget specific data, we identify an unintended utility drop—correlation collapse—in which the essential correlations between image features and true labels weaken during the forgetting process. To address this challenge, we propose Distribution-Level Feature Distancing (DLFD), a novel method that efficiently forgets instances while preserving task-relevant feature correlations. Our method synthesizes data samples by optimizing the feature distribution to be distinctly different from that of forget samples, achieving effective results within a single training epoch. Through extensive experiments on facial recognition datasets, we demonstrate that our approach significantly outperforms state-of-the-art machine unlearning methods in both forgetting performance and model utility preservation.

Introduction

Deep neural network models have achieved remarkable success in various computer vision applications (He et al. 2016; Tan and Le 2019; Han et al. 2022; Huang et al. 2017; Jiang et al. 2022). Especially, recent works show large-scale foundation models demonstrate superior classification performance across a range of tasks (Radford et al. 2021; Kolesnikov et al. 2020; Floridi and Chiriatti 2020; Han et al. 2022; Liu et al. 2023). However, alongside these advancements, concerns emerge regarding the unintentional leakage of sensitive information, such as personal identities from training data (Shokri et al. 2017a; Hu et al. 2022).

Machine unlearning has emerged as a promising solution to mitigate potential data leakage (Tarun et al. 2023; Golatkar, Achille, and Soatto 2020b; Gupta et al. 2021; Bourtole, Chandrasekaran et al. 2021; Foster, Schoepf, and Brintrup 2024), particularly in upholding the *right to be forgotten*,

which allows individuals to request the removal of their personal information from trained models. For example, in medical AI applications, a patient might request that their medical images, used during the training of a diagnostic model, be removed to protect their privacy. In such a scenario, machine unlearning enables the model to forget the patient’s data without compromising overall performance on other tasks. This growing need for privacy has driven interest in machine unlearning research within various AI-driven industries.

Despite advancements in machine unlearning algorithms, we identify a critical issue that has not been fully explored: the risk of **correlation collapse**. When simply applying existing error-maximizing methods (Tarun et al. 2023; Kurmanji, Triantafillou, and Triantafillou 2023; Chundawat et al. 2023b), unexpected outcomes can occur. For instance, these methods can inadvertently increase the magnitude of loss excessively, leading to additional data leakage by making certain data points appear *special*. Moreover, relying solely on these approaches may degrade the generalization performance of the model on the original task, introducing a trade-off between model utility and forgetting. We believe this degradation is due to correlation collapse, where the useful correlations between image features and their true labels are weakened. To prevent these unexpected performance drops, it is crucial to carefully adapt and improve upon the existing methods.

To address this challenge, we propose a novel framework, *Distribution-Level Feature Distancing (DLFD)* that enables unlearning of specific images while maintaining the accuracy of the original task. Our approach shifts the **feature distribution of the retain images** away from the distribution of the forget images, by leveraging the *Optimal Transport (OT)* problem (Peyré, Cuturi et al. 2019; Le et al. 2021; Cuturi 2013; Altschuler, Niles-Weed, and Rigollet 2017). Specifically, DLFD generates perturbed images by maximizing the distance between the optimized data distribution and the forget data distribution in the feature space using OT loss.

Our method demonstrates superior performance compared to state-of-the-art methods in a setting that closely reflects real-world scenarios. We also introduce and analyze the concept of correlation collapse, which has not been extensively addressed in previous works, and revisit the task-agnostic instance unlearning setting. Our contributions are as follows:

- We identify and address **correlation collapse**, a critical issue that can lead to a drop in model utility, and propose

*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

an effective solution to mitigate this risk.

- We propose a novel method, *Distribution-Level Feature Distancing (DLFD)*, that generates a proxy data distribution distinct from the distribution of data to be forgotten.
- Through extensive experiments, we demonstrate that our method outperforms previous SOTA methods in task-agnostic machine unlearning.

Related Work

The previous machine unlearning algorithms typically rely on two main concepts: (1) model manipulation, and (2) data manipulation. Firstly, various studies address the machine unlearning problem by directly manipulating the parameters of the model to erase specific information. For instance, the *Fisher Forgetting* (Golatkar, Achille, and Soatto 2020a) method scrubs the model by directly adding specific noises to the parameters using the inverse of the Fisher information matrix. Another approach, SCRUB (Kurmanji, Triantafillou, and Triantafillou 2023), improves forgetting performance by using a teacher model that is a clone of the original model. This method trains the unlearned model by minimizing the KL divergence between the output probability of the unlearned model ($\theta_{unlearned}$) and that of the teacher model ($\theta_{teacher}$). Similarly, the BadTeaching (Chundawat et al. 2023a) method employs three models: a competent teacher, an incompetent teacher, and a student (unlearned model $\theta_{unlearned}$). The student model is trained to mimic the competent teacher on the D_{retain} while following the incompetent teacher on the D_{forget} . These methods highlight the effectiveness of teacher-student models in enhancing unlearning performance.

On the other hand, some methods focus on data manipulation. For example, UNSIR (Tarun et al. 2023) generates noise that is added to the data to maximize the loss values for a specific target class that needs to be forgotten. Training on these error-maximized data points has shown good forgetting performance. Building on this, another method (Chundawat et al. 2023b) uses *samples to be retain* to improve unlearning scores, extending the work of UNSIR. Similarly, recent works (Cha et al. 2024) use perturbing noise to increase the loss value, focusing primarily on error-maximizing synthesized images to achieve a high forgetting score.

Despite their effectiveness in achieving high forgetting scores, we argue that this error-maximizing approach can easily lead to **correlation collapse** (Figure 1), where the useful correlations between features and labels degrade. Our work addresses these challenges by focusing on distribution-level changes rather than instance-level perturbations, which will be elaborated upon in the subsequent sections.

Motivation: Correlation Collapse

In the general computer vision domain, a feature vector $w \in \mathcal{W}$ corresponding to an image x might contain various semantic information (Na, Ji, and Kim 2022; Richardson et al. 2021). Some of these semantic features, which we denote as w_{task} , are highly correlated with the original task that the model $\theta_{original}$ is designed to solve. In addition, for personal identity unlearning tasks, another set of features, $w_{identity}$, represents information specific to personal information.

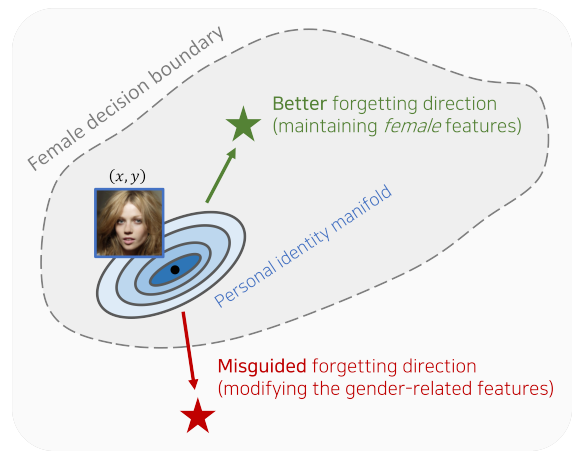


Figure 1: The concept of correlation collapse. If following the **misguided** forgetting direction, the correlation between the task-related useful features and labels can weaken.

In the latent space \mathcal{W} , we denote $\mathcal{W}_{identity} \subset \mathcal{W}$ and $\mathcal{W}_{task} \subset \mathcal{W}$ as the manifolds of identity and task features respectively, with $\mathcal{W}_{identity} \cap \mathcal{W}_{task} \neq \emptyset$. This feature space overlap manifests in individual feature vectors: for any image, its feature representations $w_{identity}$ and w_{task} share common elements. For example, in facial gender classification, attributes like hair length and facial structure exist in both identity and task-relevant features—they help identify an individual while also providing gender-related information.

This inherent overlap leads to what we term **feature entangling**, making it fundamentally challenging to separate identity information from task-relevant features. When error-maximizing methods attempt to remove identity information, they inevitably affect the shared features, resulting in **correlation collapse**: a phenomenon where the model’s ability to leverage task-relevant features deteriorates, leading to degraded classification performance.




Figure 2: Feature representations from age classification model. (a) demonstrates clear class distinctions, with age groups well-separated in feature space. (b), derived using Negative Gradient method, shows clustered features with less distinction, illustrating correlation collapse.

As illustrated in Figure 2, which visualizes feature representations in the latent space, the Original Model (Figure 2(a)) maintains clear class separations, preserving task-relevant features. In contrast, the Error Maximized model (Figure 2(b)) shows diminished class distinctions, highlighting correlation collapse caused by shared feature disruption.

To address this, we propose **Distribution-Level Feature**

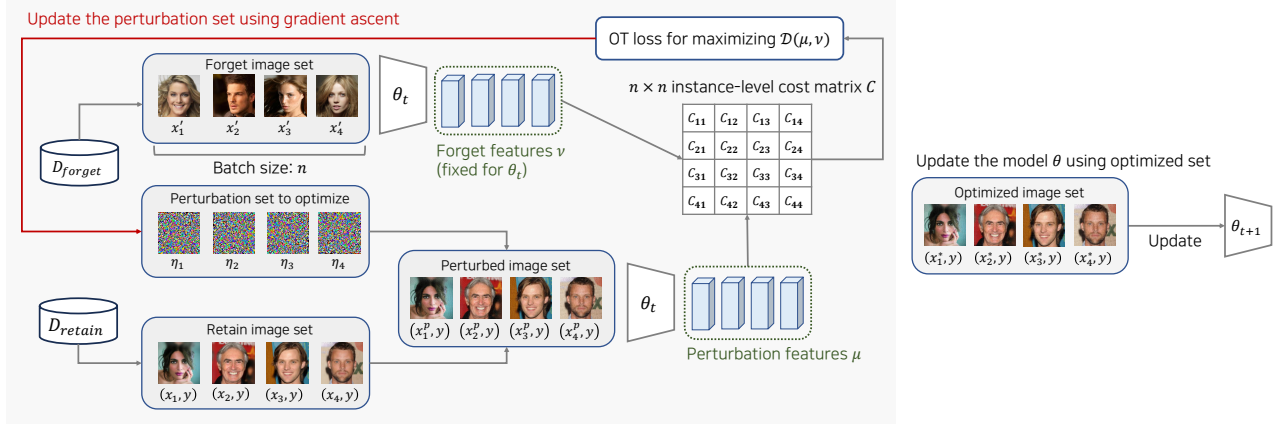


Figure 3: The core method of DLFD-feature distribution optimization through optimal transport. This component generates a synthesized dataset by maximizing the distance between retain and forget data distributions in the feature space. When combined with other components (detailed in Algorithm1), DLFD achieves a balance between model utility and forgetting performance.

Distancing (DLFD), which maintains task-related features during unlearning (Figure 3). Our method preserves the structure of \mathcal{W}_{task} while modifying identity-related features.

Proposed Methods

In this section, we introduce **Distribution-Level Feature Distancing (DLFD)**, our comprehensive framework for effective machine unlearning. DLFD consists of three key components designed to balance forgetting performance and model utility.

Feature Distribution Optimization

Traditional approaches to machine unlearning often focus on point-wise optimization, where individual data points are manipulated to maximize the loss for data that needs to be forgotten (Tarun et al. 2023; Chundawat et al. 2023b). However, such methods can lead to issues like label leakage and correlation collapse, where the underlying relationships between features are disrupted (Kurakin, Goodfellow, and Bengio 2017; Madry et al. 2018; Ilyas et al. 2019). To overcome these limitations, we propose a more holistic approach that considers the entire distribution of the data.

Our first component shifts the retain data distribution (μ) away from the forget data distribution (ν) by leveraging the optimal transport (OT) distance. Unlike simpler metrics such as KL or JS divergence, it captures the complex, high-dimensional relationships between data points (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017). The OT distance between the distributions μ and ν is defined as:

$$\mathcal{D}(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \mathbb{E}_{(w, w') \sim \gamma} [c(w, w')] \quad (1)$$

Here, γ is the set of all possible joint distributions that can transport μ to ν , and $c(w, w')$ represents the cost based on cosine similarity between feature vectors. To handle the complexity of solving this problem directly, we employ a differentiable Sinkhorn method (Cuturi 2013; Altschuler, Niles-Weed, and Rigollet 2017), which approximates the solution efficiently and reduces the computational complexity to $\mathcal{O}(n^2)$ for mini-batch computations.

To further refine the OT distance, we reformulate the problem to find an optimal transport plan T :

$$T^\lambda = \arg \min_{T \in \Pi(\mu, \nu)} \langle T, C \rangle - \frac{1}{\lambda} \sum_{i=1}^n \sum_{j=1}^n T_{ij} \log T_{ij} \quad (2)$$

In this equation, the cost matrix C captures pairwise distances between feature vectors from μ and ν . The regularization term λ keeps the transport plan T smooth, preventing overly concentrated mass transfers that could destabilize the model. Iteratively optimizing T^λ effectively separates the retain and forget data distributions, mitigating the risk of correlation collapse.

Classification Loss Preservation

To maintain model utility during the unlearning process, our second component incorporates a *classification loss* guiding the perturbation process to address correlation collapse. The classification loss ensures that the original class information of the retain data is preserved, even as the model attempts to forget data points. A critical aspect of this component is the use of a **linear weight** dynamically adjusting the importance of the classification loss throughout the training process.

The linear weight plays a crucial role in balancing the trade-off between maximizing the separation of distributions and preserving the model's utility. At the beginning of training, it is set lower, allowing the model to focus more on maximizing the distance between the retain and forget data distributions. As training progresses, the linear weight gradually increases, shifting the model's focus toward preserving the original class-specific features of the retain data. The perturbation applied to the retain data points x_i is computed as follows:

$$x_i^* \leftarrow x_i + \alpha \cdot \text{sign}(\nabla_{x_i} [l_{OT} - \lambda \cdot l_{CE}]) \quad (3)$$

Here, l_{OT} represents the OT loss between the retain and forget data distributions, while l_{CE} indicates the classification loss, weighted by the linear factor λ and computed as:

$$l_{CE} \leftarrow \text{CE}(y_i, \theta(x_i^*)) \quad (4)$$

Algorithm 1: Distribution-Level Feature Distancing (DLFD)

```
1: Input: Total batch iterations in one epoch  $K$ , Feature
   distancing steps  $M$ , learning rate  $\gamma$ , step size  $\alpha$ , batch
   size  $n$ , retain dataset  $\mathcal{D}_{retain}$ , forget dataset  $\mathcal{D}_{forget}$ ,
   model  $\theta_{original}$ 
2: Output: Unlearned model  $\theta^*$ 
3: Initialization:  $\theta^* \leftarrow \theta_{original}$ 
4: for  $k = 1$  to  $K$  do
5:   Sample retain and forget batches:
6:      $\{(x_i, y_i)\}_1^n \sim \mathcal{D}_{retain}$ 
7:      $\{(x'_i, y'_i)\}_1^n \sim \mathcal{D}_{forget}$ 
8:   Evaluate forgetting score  $F_{score}$ 
9:   if  $F_{score} \geq \text{threshold}$  then
10:    Initialize perturbed samples:  $\{x_i^*\}_1^n \leftarrow \{x_i\}_1^n$ 
11:    for  $m = 1$  to  $M$  do
12:      Compute OT loss for perturbation:
13:        Extract features  $F_{retain} \leftarrow F(\{x_i^*\}_1^n)$ 
14:         $F_{forget} \leftarrow F(\{x'_i\}_1^n)$ 
15:         $l_{OT} \leftarrow \text{Optimal Transport loss}(F_{retain}, F_{forget})$ 
16:      Compute classification loss:
17:         $l_{CE} \leftarrow \text{CE}(\{y_i\}_1^n, \theta^*(\{x_i^*\}_1^n))$ 
18:      Compute combined perturbation loss:
19:         $\lambda \leftarrow \text{linear\_weight}(k, K)$ 
20:         $l_{perturb} \leftarrow l_{OT} + \lambda \cdot (-l_{CE})$ 
21:      Update samples with perturbation loss:
22:         $\{x_i^*\}_1^n \leftarrow \{x_i^*\}_1^n + \alpha \cdot \text{sign}(\nabla_{\{x_i^*\}_1^n} l_{perturb})$ 
23:    end for
24:    Apply perturbation loss to the model:
25:     $l_{train} \leftarrow \text{CE}(\{y_i\}_1^n, \theta^*(\{x_i^*\}_1^n))$ 
26:  else
27:    Compute classification loss for model update:
28:     $l_{train} \leftarrow \text{CE}(\{y_i\}_1^n, \theta^*(\{x_i\}_1^n))$ 
29:  end if
30:  Update model parameters:
31:   $\theta^* \leftarrow \theta^* - \gamma \cdot \nabla_{\theta^*} l_{train}$ 
32: end for
```

The linear weight λ is adjusted throughout the training process to balance the trade-off between maximizing the OT loss and preserving the classification accuracy. The perturbation is scaled by a step size α and applied in the direction that increases the OT loss and decreases the weighted classification loss. This ensures that the perturbed data not only becomes more distinct from the forget data but also maintains its original task-related features.

Dynamic Forgetting Strategy

Our third component introduces an adaptive approach designed to optimize the forgetting process by continuously monitoring the forgetting score during training. Specifically, a subset of the validation set is used to simplify the forgetting monitoring process. When the forgetting score, assessed through this subset, drops below a predefined threshold—indicating that the model has sufficiently forgotten the target data—the algorithm dynamically shifts its focus from using the optimal transport optimization to exclusively fine-tuning the model with classification loss.

This transition not only reduces the computational overhead by avoiding unnecessary further perturbations but also ensures that the model’s original task performance remains stable. By fine-tuning solely with classification loss at this stage, the strategy helps preserve the important task-related features, preventing potential degradation in model utility.

Experiments

Preliminaries

In machine unlearning research, an original model $\theta_{original}$ is trained on the dataset \mathcal{D}_{train} to solve a specific task. To evaluate the **model utility**, we measure the classification accuracy of the model on the test set \mathcal{D}_{test} . If the model achieves high accuracy on \mathcal{D}_{test} , it is considered to have a high utility *of the model* for the original task.

The goal of an ideal machine unlearning method is to remove the images that need to be forgotten (\mathcal{D}_{forget}) while maintaining the original classification performance. In this study, we adopt a common machine unlearning setting where the model has access to a subset of the training data, \mathcal{D}_{retain} , which the AI company may still possess. Formally, we assume that the training data \mathcal{D}_{train} is composed of \mathcal{D}_{retain} and \mathcal{D}_{forget} , following the general machine unlearning setting described by Choi and Na (2023). Our objective is to develop a machine unlearning algorithm that makes the unlearned model $\theta_{unlearned}$ as similar as possible to the retrained model $\theta_{retrained}$, which is considered the ground truth and is trained only on \mathcal{D}_{retain} .

We also introduce a dataset \mathcal{D}_{unseen} , which is never used during the training or testing phases of the model. This dataset serves as our test set \mathcal{D}_{test} and is exclusively used for evaluating the **forgetting score**. It is important to note that any subject targeted for unlearning should not simultaneously belong to the three datasets: \mathcal{D}_{forget} , \mathcal{D}_{retain} , and \mathcal{D}_{unseen} . This ensures that the subject to be forgotten is not present across multiple datasets in the machine unlearning setting.

Task Agnostic Instance-Unlearning

In this work, we adopt a task-agnostic machine unlearning setup, which ensures that unlearning specific target subjects does not affect the model’s original functionality. Traditional machine unlearning research has primarily focused on class-unlearning, where entire categories (classes) are removed from the model upon a data removal request (Tarun et al. 2023; Golatkar, Achille, and Soatto 2020b; Goel, Prabhu, and Kumaraguru 2022b). While this approach works in certain scenarios, it is not applicable in all cases. For instance, in a gender classification model, removing the *male* class would leave only the *female* class, rendering the model ineffective for its intended purpose of gender classification. Hence, class-unlearning is not always representative of real-world needs.

To address these limitations, we propose an instance-unlearning problem setting, which targets the removal of specific personal identities or data samples without changing the overall function of the model. This approach ensures that the model’s core functionality remains intact, making it more applicable to scenarios where the goal is to forget specific

	Evaluation Metrics	Original	Retrained	Fine-tuning	NegGrad	CF-k	EU-k	UNSIR	BadT	SCRUB	DLFD
Facial Age (8-classes)	Test Acc. \uparrow	0.6329	0.6050	0.6349	0.6283	0.6323	0.4767	0.5950	0.3663	0.6311	0.6166
	Top-2 Acc. \uparrow	0.8803	0.8430	0.8876	0.8736	0.8736	0.6901	0.8503	0.6828	0.8743	0.8806
	Forgetting Score \downarrow	0.1923	0.0767	0.1980	0.1880	0.1853	0.0438	0.0887	0.0455	0.1614	0.0385
	NoMUS \uparrow	0.6241	0.7258	0.61945	0.62615	0.6308	0.69455	0.7088	0.6376	0.6541	0.7698
Facial Emotion (7-classes)	Test Acc. \uparrow	0.7535	0.6897	0.7509	0.7506	0.7513	0.7511	0.5788	0.5176	0.7509	0.6613
	Forgetting Score \downarrow	0.1852	0.0195	0.1735	0.1862	0.1845	0.1585	0.0192	0.0250	0.1391	0.0372
	NoMUS \uparrow	0.6915	0.8253	0.7019	0.6891	0.6911	0.7171	0.7702	0.7338	0.73635	0.7934
Multi- Attributes (3-labels)	Average Test Acc. \uparrow	0.9212	0.8700	0.9218	0.4487	0.9192	0.9189	0.9233	0.8129	0.7057	0.9129
	Forgetting Score \downarrow	0.0501	0.0044	0.0443	0.0009	0.04663	0.0399	0.0511	0.0164	0.0184	0.0281
	NoMUS \uparrow	0.9105	0.9306	0.9166	0.7234	0.9129	0.9195	0.9105	0.8900	0.8344	0.9283
Facial Gender (binary-class)	Test Acc. \uparrow	0.9016	0.8493	0.9215	0.1733	0.9196	0.9216	0.9142	0.9046	0.9214	0.8997
	Forgetting Score \downarrow	0.0461	0.0149	0.0488	0.0895	0.0581	0.0576	0.0663	0.0453	0.0615	0.0306
	NoMUS \uparrow	0.9047	0.9097	0.9119	0.4971	0.9017	0.9031	0.8908	0.9070	0.8992	0.9192

Table 1: Overall performance of various machine unlearning methods on ResNet18 classification tasks. Our method achieves superior NoMUS scores across all tasks, with remarkable forgetting scores while maintaining competitive test accuracy. The best score is in **boldface** except for the ground-truth (*Retrained*). *Note*: Fine-tuning, NegGrad (Golatkar, Achille, and Soatto 2020b), CF-k, EU-k (Goel, Prabhu, and Kumaraguru 2022a), UNSIR (Tarun et al. 2023), BadTeaching (Chundawat et al. 2023a), SCRUB (Kurmanji, Triantafillou, and Triantafillou 2023)

data without compromising the model’s utility (Triantafillou et al. 2023; Choi and Na 2023; Choi et al. 2024).

While recent studies on instance-unlearning often focus on forcing misclassification of specific instances, this deviates from a truly task-agnostic approach (Liu et al. 2024; Shen et al. 2024; Cha et al. 2024). Our method differs by preserving the original task’s functionality while ensuring that specific instances are unlearned. For instance, consider a chest X-ray (CXR) disease classification model. This model uses chest X-ray images to predict the likelihood of diseases such as tuberculosis or pneumonia. Even if all images associated with a particular patient are removed, the model should still accurately diagnose these diseases for other patients. This task-agnostic approach ensures that the model’s core functionality is preserved, making it more robust and practical for real-world applications. Focusing on instance-unlearning within a task-agnostic framework, our method addresses a significant gap in current research, offering a solution that maintains the model’s task-related performance while effectively unlearning specific instances.

Evaluation Protocol

In this work, we evaluate the models using two metrics: (1) model utility and (2) forgetting score. The model utility is assessed by measuring the test accuracy on \mathcal{D}_{test} . A high accuracy on \mathcal{D}_{test} indicates that the model retains strong performance on its original task after the unlearning process.

For forgetting performance, we define a forgetting score based on the success rate of a Membership Inference Attack (MIA) (Shokri et al. 2017b). The MIA framework is formulated as follows:

$$\psi(x) = \begin{cases} 1 & \text{if } x \in \mathcal{D}_{forget} \\ 0 & \text{if } x \in \mathcal{D}_{unseen} \end{cases} \quad (5)$$

Given \mathcal{D}_{forget} and \mathcal{D}_{unseen} datasets, we train a binary classifier $\psi(\cdot)$ to distinguish between them.

The classifier $\psi(\cdot)$ is trained using binary cross-entropy loss on model predictions and loss values from $\theta_{original}$. The

forgetting score is then defined as:

$$\text{Forgetting Score} = |\text{MIA Acc.} - 0.5| \times 2 \quad (6)$$

where MIA Acc. is the binary classification accuracy of $\psi(\cdot)$.

A perfect forgetting score of 0.0 indicates that the model has completely forgotten the target data, as the MIA classifier achieves only random chance (0.5) accuracy in distinguishing between forget and unseen samples.

To capture both model utility and forgetting performance in a single metric, we calculate the **Normalized Machine Unlearning Score (NoMUS)** (Choi and Na 2023) as follows:

$$\text{NoMUS} = \frac{1}{2} (P(\hat{y} = y) + (1 - \text{Forgetting Score})) \quad (7)$$

where $P(\hat{y} = y)$ represents the model’s classification performance on \mathcal{D}_{test} . NoMUS ranges from 0 to 1, with higher values indicating better overall performance in both utility preservation and successful unlearning.

Datasets

For our experiments, we utilize three distinct facial datasets, each designed for specific classification tasks:

- **Age Estimation:** The MUFAC dataset (Choi and Na 2023) contains 13,068 facial images (128×128) in 8 age groups. The training set comprises 10,025 samples, with 8,525 retained and 1,500 designated for forgetting.
- **Emotion Recognition:** The RAF-DB dataset (Li, Deng, and Du 2017) contains 15,000 images across 7 emotional classes. The training set comprises 11,044 samples, with 7,730 retained and 3,314 designated for forgetting.
- **Multi-Attribute Classification:** The MUCAC dataset (Choi and Na 2023), derived from CelebA (Liu et al. 2018), consists of 30,000 facial images with three binary attributes: gender, age, and expression. The training set includes 25,933 samples, with 15,385 retained and 10,548 for forgetting.

ID	ResNet18				DenseNet121				EfficientNetB0			
	Age	Emotion	Multi-Attr.	Gender	Age	Emotion	Multi-Attr.	Gender	Age	Emotion	Multi-Attr.	Gender
Original	0.6242	0.6915	0.9105	0.9047	0.6813	0.7259	0.9099	0.9088	0.6419	0.6960	0.9389	0.8916
Retrained	0.7258	0.8253	0.9306	0.9097	0.7617	0.8504	0.9063	0.8901	0.7580	0.8375	0.9347	0.8964
Fine-tuning	0.6195	0.7019	0.9166	0.9119	0.6671	0.7260	0.9174	0.9147	0.6412	0.7026	0.9324	0.8974
NegGrad	0.6946	0.6891	0.7234	0.4971	0.7375	0.7289	0.9220	0.5353	0.6943	0.6938	0.9344	0.8901
CF-k	0.6262	0.6911	0.9129	0.9017	0.6677	0.7286	0.9200	0.9103	0.6478	0.6931	0.9294	0.8993
EU-k	0.6309	0.7171	0.9195	0.9031	0.6900	0.7433	0.9183	0.9101	0.6507	0.6258	0.9328	0.8908
UNSIR	0.7088	0.7702	0.9105	0.8908	0.6321	0.7524	0.9206	0.9017	0.7315	0.8009	0.9304	0.8876
Bad Teaching	0.6377	0.7338	0.8900	0.9070	0.6211	0.6736	0.9170	0.9009	0.7202	0.6711	0.9352	0.8757
SCRUB	0.6554	0.7363	0.8344	0.8992	0.6989	0.6727	0.8916	0.9036	0.6574	0.67245	0.8960	0.8890
DLFD	0.7698	0.7935	0.9283	0.9192	0.7733	0.7884	0.9433	0.9202	0.7617	0.7801	0.9529	0.9208

Table 2: The overall results of the major machine unlearning methods. The results are calculated using NoMUS. Our method shows superior performance compared to SOTA methods. The best scores are in **boldface** except the ground-truth (*Retrained*).

Experimental Setup

For experiments, we utilize various deep-neural network architectures including ResNet (He et al. 2016), DenseNet (Huang et al. 2017), and EfficientNet (Tan and Le 2019), widely adopted in computer vision. To ensure fair comparison, all machine unlearning methods start from the same $\theta_{original}$ for each task. Specifically, methods fine-tune $\theta_{unlearned}$, initialized as $\theta_{original}$, except for the Retrained model. Serving as ground truth, the Retrained model is trained from scratch on \mathcal{D}_{retain} , excluding data to be forgotten, to fully represent the desired unlearning outcome.

Given the computational complexity of our method, which involves calculating OT loss and performing MIA evaluations, we limit the training to a single epoch. Other machine unlearning methods are also trained for 1-2 epochs to ensure a fair comparison. Additionally, we find that learning rates between 0.001 and 0.005 are effective across all models and methods, consistent with previous work (Tarun et al. 2023).




Figure 4: Feature representations from the age classification model trained with DLFD. The model preserves class separation similar to the original model (Figure 2(a)), retaining task-relevant features while mitigating correlation collapse.

Performance of DLFD Method

We evaluate our method across four classification tasks: facial age prediction, emotion recognition, multi-attribute classification, and gender classification. The multi-attribute model includes three binary labels: gender (female/male), age (old/young), and expression (smiling/unsmiling), with the average classification accuracy reported as the model utility. Gender classification, originally part of the multi-attribute, is also evaluated as an independent binary classification task.

We compare our method with various previously proposed methods. As shown in Table 1, our method demonstrates superior performance in the comprehensive metric, NoMUS. Although *Fine-tuning*, *CF-k*, and *EU-k* can generally achieve high test accuracy, their forgetting scores are generally low, indicating insufficient unlearning performance. On the other hand, the *Retrained (ground-truth)* model shows excellent forgetting performance but suffers a significant drop in test accuracy, which adversely impacts model utility.

Across all experiments, our method consistently delivers competitive or superior performance in both metrics. We demonstrate that DLFD effectively unlearns the forget data while maintaining model utility. As shown in Figure 5c, the loss distributions of \mathcal{D}_{unseen} and \mathcal{D}_{forget} in our method closely resemble those of the *Retrained* model, considered the ground truth (Figure 5b). This similarity indicates that the unlearning algorithm works effectively.

In particular, the DLFD method shows more substantial improvements in complex, multi-class tasks such as age estimation and emotion recognition, where the feature entanglement is more significant. Conversely, the improvements in multi-attribute and gender classification tasks are relatively smaller, likely due to the binary nature of these classifications, where the complexity of feature entanglement is inherently lower. These results highlight the effectiveness of our approach in scenarios where maintaining feature integrity amid complex and overlapping feature spaces is more challenging.

Moreover, Figure 4 displays the feature representations extracted by the DLFD model. The figure demonstrates how DLFD maintains clear class distinctions, preventing correlation collapse and preserving essential task-related features.

Ablation Study

We perform ablation studies to evaluate each component of DLFD. Using only feature distribution optimization initially achieves success in machine unlearning, as shown in Table 3. While effective in separating retain and forget data distributions, this component alone may reduce model utility without the support of other components.

The addition of classification loss preservation as the second component significantly improves performance, especially in Age and Emotion tasks (NoMUS increased by 7.3% and 4.7%), highlighting its role in maintaining model utility.

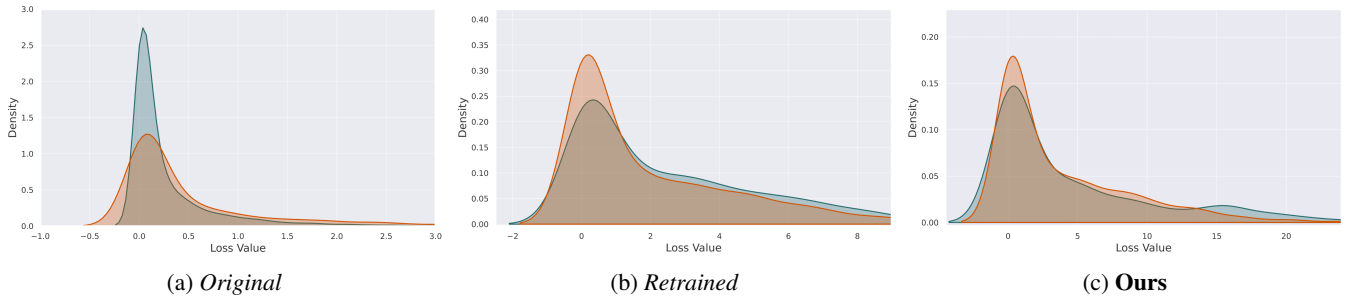


Figure 5: The loss distributions for two baselines and ours. The orange space represents the loss distribution for unseen data, while the green represents the loss distribution for forget data. (a) illustrates loss distributions for *Original* model. (b) shows loss distributions for *Retrained* model. Finally (c) represents loss distributions for $\theta_{unlearned}$ fine-tuned on DLFD-optimized images.

Finally, integrating dynamic forgetting further enhances performance, with additional improvements in Age (2.%) and Emotion (5.3%) tasks. This component effectively prevents correlation collapse by balancing forgetting and utility preservation. The complete framework, combining all three components, achieves superior NoMUS scores across all tasks, showing the effectiveness of their synergistic interaction.

Feature Dist. Optim.	Cls Loss	Dynamic Forgetting	Age	Emotion	Multi-Attr.	Gender
			NoMUS \uparrow	NoMUS \uparrow	NoMUS \uparrow	NoMUS \uparrow
✓			0.7021	0.7199	0.9423	0.8898
✓	✓		0.7536	0.7536	0.9420	0.9170
✓	✓	✓	0.7698	0.7934	0.9433	0.9192

Table 3: Ablation study results for each component in DLFD, showing cumulative performance improvements.

Discussion

Information Leakage in Error-Maximization

A trained model generally shows lower loss values for training data compared to unseen data, which can lead to data leakage. Methods like UNSIR (Tarun et al. 2023) and SCRUB (Kurmanji, Triantafillou, and Triantafillou 2023) that maximize loss for data intended to be forgotten may inadvertently increase the loss for forget data beyond that of unseen data, making the model vulnerable to membership inference attacks. Our findings reveal that even with unlearning, naive error-maximization can still result in information leakage. Specifically, when the number of forget samples is small (< 100), the loss values for forget data can abnormally increase, exceeding those of unseen data (Figure 6). This issue highlights a risk that has been overlooked in prior studies.

Trade-off between Model Utility and Forgetting

Our method reveals a trade-off between test accuracy and forgetting score. As we increase the loss for the data intended to be forgotten (x_{forget}), the forgetting score improves, but this comes at the cost of test accuracy. This is likely due to *correlation collapse*, where essential label-related features of the retain data are altered. Moreover, the effectiveness of unlearning strategies can vary depending on the dataset’s characteristics, scale, and task complexity. This variability suggests the challenges of setting up robust unlearning experiments and the need for adaptive unlearning methods adjusted to different scenarios. Our findings emphasize the importance of

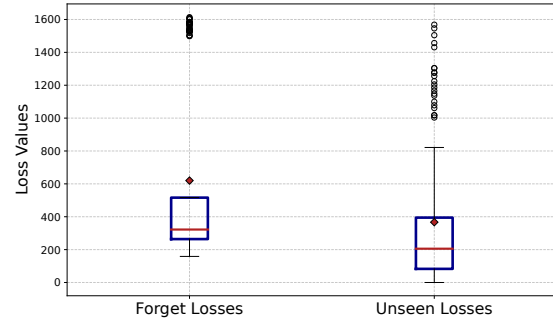


Figure 6: The naive instance-level error-maximizing can induce excessively high loss value for *forget* data even larger than the *unseen* data, which might not be desirable.

carefully balancing model utility and forgetting performance to achieve optimal unlearning results.

Practical Considerations and Future Work

One potential limitation of our method could arise when retain and forget datasets have overlapping features in the feature space. While our current implementation demonstrates strong performance in settings with minimal overlap, handling heavily overlapping feature distributions remains a challenging scenario that warrants further investigation.

Moreover, although MIA is widely used as a metric to assess forgetting performance, it may not fully capture unlearning effectiveness across all scenarios. In scenarios where the model is exceptionally well-trained, the distinction between forget and unseen data may become minimal, leading to MIA scores that do not adequately reflect the true forgetting performance. This suggests the need for the unlearning community to develop more robust evaluation metrics.

Conclusion

We address key challenges in machine unlearning, including information leakage in error-maximizing methods, task-specific settings, and the critical trade-off between model utility and effective forgetting. Our proposed DLFD method effectively mitigates these issues by reducing the risk of correlation collapse while maintaining high model utility. Experimental results consistently demonstrate that DLFD outperforms existing methods across multiple benchmarks, underscoring its robustness and effectiveness.

Acknowledgements

This research was supported by Brian Impact, a non-profit organization dedicated to advancing science and technology.

References

- Altschuler, J.; Niles-Weed, J.; and Rigollet, P. 2017. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. *Advances in neural information processing systems*, 30.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, 214–223.
- Bourtoule, L.; Chandrasekaran, V.; et al. 2021. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, 141–159. IEEE.
- Cha, S.; Cho, S.; Hwang, D.; Lee, H.; Moon, T.; and Lee, M. 2024. Learning to unlearn: Instance-wise unlearning for pre-trained classifiers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11186–11194.
- Choi, D.; Choi, S.; Lee, E.; Seo, J.; and Na, D. 2024. Towards Efficient Machine Unlearning with Data Augmentation: Guided Loss-Increasing (GLI) to Prevent the Catastrophic Model Utility Drop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 93–102.
- Choi, D.; and Na, D. 2023. Towards machine unlearning benchmarks: Forgetting the personal identities in facial recognition systems. *arXiv preprint arXiv:2311.02240*.
- Chundawat, V. S.; Tarun, A. K.; Mandal, M.; and Kankanhalli, M. 2023a. Can bad teaching induce forgetting? Unlearning in deep networks using an incompetent teacher. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7210–7217.
- Chundawat, V. S.; Tarun, A. K.; Mandal, M.; and Kankanhalli, M. 2023b. Zero-shot machine unlearning. *IEEE Transactions on Information Forensics and Security*, 18: 2345–2354.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Floridi, L.; and Chiriatti, M. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30: 681–694.
- Foster, J.; Schoepf, S.; and Brintrup, A. 2024. Fast machine unlearning without retraining through selective synaptic dampening. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12043–12051.
- Goel, S.; Prabhu, A.; and Kumaraguru, P. 2022a. Evaluating inexact unlearning requires revisiting forgetting. *CoRR abs/2201.06640*.
- Goel, S.; Prabhu, A.; and Kumaraguru, P. 2022b. Evaluating inexact unlearning requires revisiting forgetting. *arXiv preprint arXiv:2201.06640*.
- Golatkar, A.; Achille, A.; and Soatto, S. 2020a. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9304–9312.
- Golatkar, A.; Achille, A.; and Soatto, S. 2020b. Forgetting outside the box: Scrubbing deep networks of information accessible from input-output observations. In *ECCV 2020: 16th European Conference, Glasgow, UK, 2020, Proceedings*, 383–398. Springer.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C. 2017. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30.
- Gupta, V.; Jung, C.; Neel, S.; Roth, A.; Sharifi-Malvajerdi, S.; and Waites, C. 2021. Adaptive machine unlearning. *Advances in Neural Information Processing Systems*, 34: 16319–16330.
- Han, K.; Wang, Y.; Chen, H.; Chen, X.; Guo, J.; Liu, Z.; Tang, Y.; Xiao, A.; Xu, C.; Xu, Y.; et al. 2022. A survey on vision transformer. *IEEE transactions on pattern analysis and machine intelligence*, 45(1): 87–110.
- 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.
- Hu, H.; Salcic, Z.; Sun, L.; Dobbie, G.; Yu, P. S.; and Zhang, X. 2022. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)*, 54(11s): 1–37.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.
- Ilyas, A.; Santurkar, S.; Tsipras, D.; Engstrom, L.; Tran, B.; and Madry, A. 2019. Adversarial examples are not bugs, they are features. *Advances in neural information processing systems*, 32.
- Jiang, P.; Ergu, D.; Liu, F.; Cai, Y.; and Ma, B. 2022. A Review of Yolo algorithm developments. *Procedia Computer Science*, 199: 1066–1073.
- Kolesnikov, A.; Beyer, L.; Zhai, X.; Puigcerver, J.; Yung, J.; Gelly, S.; and Houlsby, N. 2020. Big transfer (bit): General visual representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, 491–507. Springer.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2017. Adversarial Machine Learning at Scale. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*.
- Kurmanji, M.; Triantafillou, P.; and Triantafillou, E. 2023. Towards Unbounded Machine Unlearning. *arXiv:2302.09880*.
- Le, K.; Nguyen, H.; Nguyen, Q. M.; Pham, T.; Bui, H.; and Ho, N. 2021. On robust optimal transport: Computational complexity and barycenter computation. *Advances in Neural Information Processing Systems*, 34: 21947–21959.
- Li, S.; Deng, W.; and Du, J. 2017. Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2852–2861.
- Liu, J.; Ram, P.; Yao, Y.; Liu, G.; Liu, Y.; SHARMA, P.; Liu, S.; et al. 2024. Model sparsity can simplify machine unlearning. *Advances in Neural Information Processing Systems*, 36.

Liu, X.; Zheng, Y.; Du, Z.; Ding, M.; Qian, Y.; Yang, Z.; and Tang, J. 2023. GPT understands, too. *AI Open*.

Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2018. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August, 15(2018)*: 11.

Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.

Na, D.; Ji, S.; and Kim, J. 2022. Unrestricted Black-Box Adversarial Attack Using GAN with Limited Queries. In *European Conference on Computer Vision*, 467–482. Springer.

Peyré, G.; Cuturi, M.; et al. 2019. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6): 355–607.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

Richardson, E.; Alaluf, Y.; Patashnik, O.; Nitzan, Y.; Azar, Y.; Shapiro, S.; and Cohen-Or, D. 2021. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2287–2296.

Shen, S.; Zhang, C.; Zhao, Y.; Bialkowski, A.; Chen, W.; and Xu, M. 2024. Label-agnostic forgetting: A supervision-free unlearning in deep models. *arXiv preprint arXiv:2404.00506*.

Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017a. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, 3–18. IEEE.

Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017b. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, 3–18. IEEE.

Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, 6105–6114. PMLR.

Tarun, A. K.; Chundawat, V. S.; Mandal, M.; and Kankanhalli, M. 2023. Fast yet effective machine unlearning. *IEEE Transactions on Neural Networks and Learning Systems*.

Triantafillou, E.; Pedregosa, F.; Hayes, J.; Kairouz, P.; and Guyon, I. e. a. 2023. NeurIPS 2023 - Machine Unlearning.