

Scissorhands: Scrub Data Influence via Connection Sensitivity in Networks

Jing Wu¹, Mehrtash Harandi¹

¹Department of Electrical and Computer Systems Engineering
Monash University, Australia
{jing.wu1, mehrtash.harandi}@monash.edu

Abstract. Machine unlearning has become a pivotal task to erase the influence of data from a trained model. It adheres to recent data regulation standards and enhances the privacy and security of machine learning applications. In this work, we present a new machine unlearning approach *Scissorhands*. Initially, *Scissorhands* identifies the most pertinent parameters in the given model relative to the forgetting data via connection sensitivity. By reinitializing the most influential top- k percent of these parameters, a trimmed model for erasing the influence of the forgetting data is obtained. Subsequently, *Scissorhands* fine-tunes the trimmed model with a gradient projection-based approach, seeking parameters that preserve information on the remaining data while discarding information related to the forgetting data. Our experimental results, conducted across image classification and image generation tasks, demonstrate that *Scissorhands*, showcases competitive performance when compared to existing methods. Source code is available at <https://github.com/JingWu321/Scissorhands>.

Keywords: Machine unlearning · Connection sensitivity · Diffusion model

Warning: This paper contains explicit sexual imagery that may be offensive.

1 Introduction

In this work, we aim to propose an effective machine unlearning method. Under data regulations like the European Union General Data Protection Regulation (GDPR) [53] and the California Consumer Privacy Act (CCPA) [22], all users are granted the *right to be forgotten*. In machine learning, these legal provisions empower data owners with the right not only to withdraw their data from trained models but also to ensure that their data’s influence on these models is erased.

The most direct approach to accomplish this objective is to retrain the model from scratch, excluding the data requested for deletion from the training process. Retraining from scratch is typically considered the gold standard in the field of machine unlearning [12, 52]. Yet, this poses a challenge as numerous in-production models require prolonged training periods and substantial computing resources. While retraining is feasible, it is often impractical.

Consequently, the proposal and development of efficient approximate unlearning methods [13, 23, 26, 50, 51] have become essential. Currently, most approximate unlearning techniques achieve forgetting by adding normally distributed noise to the parameters [19–21] or by estimating the influence of a particular data point on the model’s predictions [23, 37, 38, 48] based on the influence function method [10]. Jia *et al.* [27] demonstrate that model sparsity can help to unlearn and fuse model sparsity into the unlearning process.

Fan *et al.* [12] recently raised concerns regarding the instability in approximate unlearning methods, and highlighted that the current machine unlearning approaches, initially crafted for image classification tasks, fall short in effectively tackling the challenges of machine unlearning within the realm of image generation. The authors introduce a novel unlearning method SalUn that can effectively perform forgetting in both image classification and generation tasks. Their unlearning mechanism includes finding salient weights and then fine-tuning these weights using forgetting data assigned with random labels. However, studies like [59] show that the model can still learn the true data distribution even using data with random labels. While SalUn has shown state-of-the-art performance, it could potentially memorize information about forgetting data points, albeit associating them with random labels but still can be undesired in sensitive applications. In § 4, results also show that SalUn tends to memorize knowledge about the forgetting data.

This work. We present an unlearning approach, *Scissorhands*, designed to erase data influence in the classification models and eliminate a particular concept from a text-to-image model. To achieve this, inspired by [27], our key insight is to first erase the critical influence of forgetting data in models and then relearn. Unlike previous methods that employ the forgetting data with noisy labels as a part of relearning, we propose to use an efficient gradient projection method to relearn the critical features and patterns while ensuring the exclusion of influences associated with the forgetting data. Through a series of experiments and evaluations, including the classification on SVHN [39], CIFAR-10 and CIFAR-100 [30], CelebAMASK-HQ [31] datasets, as well as the open-source stable diffusion [44] text-to-image model, results demonstrate the viability and effectiveness of our technique in forgetting the influence of random samples, discrete classes, and sensitive content such as nudity.

2 Methodology

In this section, we propose *Scissorhands*, our unlearning framework that scrubs data from a model by eliminating their influence on model parameters. Throughout the paper, we denote scalars and vectors/matrices by lowercase and bold symbols, respectively (e.g., a , \mathbf{a} , and \mathbf{A}). We show a d -dimensional vector of ones by $\mathbf{1}_d$, and use \odot to denote the Hadamard product.

Overview. Consider a model f with parameters $\boldsymbol{\theta} \in \mathbb{R}^d$ trained on a dataset $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$. Suppose a subset of \mathcal{D} , denoted as \mathcal{D}_f , is requested for deletion.

The remaining data is defined as $\mathcal{D}_r := \mathcal{D} \setminus \mathcal{D}_f$. Our objective is to develop the unlearning algorithm to generate a scrubbed model f_u , effectively removing the influence of \mathcal{D}_f from model f , while maintaining its utility over \mathcal{D}_r .

In *Scissorhands*, we initially identify critical parameters w.r.t. the forgetting data via connection sensitivity analysis [32, 33]. We then reinitialize the top- k percent of these key parameters, resulting in a trimmed model with parameters θ_t . This step aims to diminish the model’s memory of the forgetting data \mathcal{D}_f . Nevertheless, this process runs the risk of erasing valuable information from the data set \mathcal{D}_r that we aim to retain.

To address this, the next phase of *Scissorhands* concentrates on restoring the performance of the trimmed model on remaining data \mathcal{D}_r , while ensuring that the influence of the forgetting data \mathcal{D}_f remains excluded. We accomplish this balance and obtain the final scrubbed model with parameters θ^* through a gradient projection-based approach.

2.1 Trimming

To effectively identify salient parameters of a network with respect to \mathcal{D}_f , one has to establish a good criterion to determine such salient connections. A commonly used criterion is the magnitude of the weights, with those exceeding a certain threshold deemed salient as suggested in [25]. Our requirement is more nuanced, as we need a measure that specifically discerns saliency based on \mathcal{D}_f . Therefore, we adopt a single-shot approach as proposed in [33], which we outline briefly to ensure clarity and completeness in our methodology. Consider

$$\mathbf{s}_j(\mathcal{D}) := \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \mathcal{D}} \left[\ell(\theta; \mathbf{x}, \mathbf{y}) - \ell((\mathbf{1}_d - \mathbf{e}_j) \odot \theta; \mathbf{x}, \mathbf{y}) \right], \quad (1)$$

which measures the influence of parameter $j \in \{1, \dots, d\}$ on a model in terms of the empirical risk for a given dataset \mathcal{D} . Here, $\ell : \mathbb{R}^d \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is the loss of the model and \mathbf{e}_j is the indicator vector of element j (i.e., a binary vector with its j -th component set to one). Note that computing \mathbf{s}_j for each $j \in \{1, \dots, d\}$ is prohibitively expensive, as it demands $d+1$ forward passes through the dataset. An approximation to Equation (1) [25] is in the form of

$$\mathbf{s}_j(\mathcal{D}) \approx \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \mathcal{D}} \left[\frac{\partial \ell(\theta; \mathbf{x}, \mathbf{y})}{\partial \theta_j} \theta_j \right]. \quad (2)$$

Detailed proofs can be found in Appendix A. Equation (2) defines the sensitivity of parameter j as the average (i.e., expectation) of a product, the gradient of the loss and the current value of the parameter. As such, \mathbf{s}_j encodes the loss sensitivity according to the parameter’s magnitude. Intuitively, a parameter with a large value that significantly affects the loss, as captured by the gradient, is considered more influential or salient. The form in Equation (2) offers several benefits, most importantly, the ability to calculate the sensitivity for all parameters with just a single sweep through the dataset.

To scrub the influence of the forgetting data \mathcal{D}_f in the model, we first obtain $\mathbf{s}_j(\mathcal{D}_f)$, saliency of parameters w.r.t. the forgetting data. We then re-initialize

the top- $k\%$ of the parameters based on their saliency rankings. This process is akin to performing a targeted ‘lobotomy’ on the model specifically concerning the data \mathcal{D}_f , thereby selectively erasing their influence. Unfortunately, this aggressive approach of reinitializing parameters can detrimentally impact the model’s performance on the remaining dataset \mathcal{D}_r , which we will correct in the next phase of *Scissorhands*.

Remark 1. In principle, Scissorhands can benefit from any algorithms that can identify important connections w.r.t. \mathcal{D}_f . The main reason behind choosing SNIP is its single-shot nature. As presented in the study [12], the gradient of loss w.r.t. the model parameters over the forgetting data can help identify important parameters w.r.t. \mathcal{D}_f and hence help with erasing.

Remark 2. In practice, we have observed that even a small subset of \mathcal{D}_f can be sufficient for the trimming. For example, to unlearn a large set of data on CIFAR-100, our algorithm outperforms baselines by utilizing approximately 3% of the forgetting data for trimming.

Remark 3. Initializing weights of the trimmed neurons can take various forms. Empirically, we observed initializing with a uniform distribution is particularly effective. Details on the influence of different initialization strategies and the choice of $k\%$ trimming on the model’s performance will be discussed in § 4.2.

2.2 Repairing

Following the parameter-trimming process, aimed at mitigating the influence of the forgetting data \mathcal{D}_f , a challenge presents itself: the potential erasure of crucial information associated with the remaining data \mathcal{D}_r . A straightforward solution to recover the model utility is to relearn the model using the remaining data. However, this approach risks biasing the model towards features of the remaining data and inadvertently reintegrating information about the forgetting samples. In the quest to ensure that models retain essential information from the remaining data \mathcal{D}_r while effectively forgetting \mathcal{D}_f , our strategy involves maximizing the loss over the forgetting data \mathcal{D}_f while concurrently minimizing the loss over the remaining data \mathcal{D}_r . Therefore, to achieve the balance between these goals, we propose the objective of an efficient practice for unlearning:

$$L(\boldsymbol{\theta}, \mathcal{D}_r, \mathcal{D}_f) := \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}_r) - \lambda \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}_f) . \quad (3)$$

We aim to optimize Equation (3) w.r.t. $\boldsymbol{\theta}$ through gradient descent, ensuring minimal reintroduction of information about \mathcal{D}_f . This is particularly challenging when there are similarities between samples in \mathcal{D}_f and \mathcal{D}_r . To achieve this desideratum, it is vital that updates to the model do not improve (i.e., reduce) $\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}_f)$. To this end, we define $\mathbf{g}_f = \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}_f)$ as the gradient direction for \mathcal{D}_f and $\mathbf{g}_o = \nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta}, \mathcal{D}_r, \mathcal{D}_f)$ as the gradient direction for optimizing Equation (3). The optimal descent direction $\mathbf{g} \in \mathbb{R}^d$ should thus satisfy two criteria: 1. It should exhibit maximum similarity to \mathbf{g}_o to ensure swift convergence. The

Algorithm 1 The procedure of *Scissorhands*.

```

1: procedure TRIMMING( $\theta, \mathcal{D}_f$ )
2:    $\mathbf{X}, \mathbf{y} \leftarrow \{\mathbf{x}_i, \mathbf{y}_i \sim \mathcal{D}_f | i \in [B]\}$ .
3:   // Compute connection sensitivity
4:   Get  $\mathbf{s}_j(\theta; \mathbf{X}, \mathbf{y}), \forall j \in [1, d]$ 
      using Eq. (2).
5:   // Re-initialization
6:   Re-initialize parameters according
      to the top- $k\%$  value of  $\mathbf{s}_j$ .
7:   return  $\theta_t$ 
8: end procedure

9: procedure REPAIRING( $\theta_t, \mathcal{D}_f, \mathcal{D}_r$ )
10:   $\theta^0 = \theta_t$ .
11:  for  $e = 0$  to  $E - 1$  do
12:    Get  $L(\theta, \mathcal{D}_r, \mathcal{D}_f)$  (cf. Eq. (3)).
13:    Get  $\mathbf{g}_o$  and  $\mathbf{g}_f$ .
14:    // Get the optimal direction
15:    if  $\langle \mathbf{g}_o, \mathbf{g}_f \rangle > 0$  then
16:      Compute  $v^*$  (cf. Eq. (5)).
17:       $\mathbf{g} = \mathbf{g}_o - v^* \mathbf{g}_f$ .
18:       $\mathbf{g}_o \leftarrow \mathbf{g}$ .
19:    end if
20:     $\theta^{e+1} \leftarrow \theta^e - \eta \mathbf{g}_o$ .
21:  end for
22:  return  $\theta^* = \theta^E$ .
23: end procedure

```

notion of similarity can be captured by \mathbf{g}_o using $\|\mathbf{g} - \mathbf{g}_o\|_2^2$. **2.** It should not align with \mathbf{g}_f to prevent improving $\mathcal{L}(\theta; \mathcal{D}_f)$. This can be mathematically formalized as $\langle \mathbf{g}_f, \mathbf{g} \rangle \leq 0$. Consequently, after computing $\mathbf{g}_o = \nabla_{\theta} L(\theta, \mathcal{D}_r, \mathcal{D}_f)$ and $\mathbf{g}_f = \nabla_{\theta} \mathcal{L}(\theta; \mathcal{D}_f)$ in each iteration, we formulate the following optimization challenge to identify the model update direction:

$$\begin{aligned} \arg \min_{\mathbf{g}} \quad & \frac{1}{2} \|\mathbf{g} - \mathbf{g}_o\|_2^2, \\ \text{s.t.} \quad & \langle \mathbf{g}_f, \mathbf{g} \rangle \leq 0. \end{aligned} \tag{4}$$

This problem may be addressed using the Frank-Wolfe algorithm. Given the high dimensionality of gradients in neural networks, directly addressing the constraint optimization in Equation (4) could become overwhelming. Inspired by [36], we propose to make use of the dual formulation of the problem Equation (4) as:

$$\begin{aligned} \sup_v \quad & \frac{1}{2} \mathbf{g}_f^{\top} \mathbf{g}_f v^2 - \mathbf{g}_o^{\top} \mathbf{g}_f v, \\ \text{s.t.} \quad & v \geq 0. \end{aligned} \tag{5}$$

Then, the optimal descent direction \mathbf{g} is given from the solution v^* in Equation (5) as $\mathbf{g} = \mathbf{g}_o - v^* \mathbf{g}_f$. Detailed proofs can be found in Appendix A. By doing so, we ensure a balance between unlearning and retaining utility.

Remark 4. The primal problem (cf. Equation (4)) involves optimizing over the high-dimensional space of gradient vectors, while the dual problem (cf. Equation (5)) simplifies this problem to optimize over a single scalar variable $v \in \mathbb{R}$. This dimensionality reduction can significantly decrease the computational complexity, especially for large-scale neural networks.

2.3 Algorithm Description

Algorithm 1 describes the procedure of *Scissorhands* in detail. We first identify the important parameters w.r.t. the forgetting data in single-shot via connection sensitivity in networks, then re-initialize these parameters for erasing the influence of the forgetting data \mathcal{D}_f in the given model. To repair the performance of the trimmed model on the remaining data \mathcal{D}_r , while ensuring the exclusion of information associated with the forgetting data \mathcal{D}_f , we employ the gradient projection-based optimization algorithm. By doing so, we can scrub the influence of the forgetting data while retaining model utility on the remaining data. The loss depicted in Equation (3) can be viewed as a scalarization of a Multi-Objective Optimization (MOO) problem, namely minimizing $(\mathcal{L}(\theta; \mathcal{D}_r), -\mathcal{L}(\theta; \mathcal{D}_f))^{\top}$. A common issue in MOO is the gradient conflict, as optimizing one objective could hinder another one. Our proposal on gradient projection is designed to address and mitigate the gradient conflict. Empirically (see Table 3), we studied the effect of removing gradient projection.

3 Related Work

The development of efficient machine unlearning methods [2, 3, 11, 17, 18, 20, 23, 27, 28, 37, 40, 42, 45, 48, 51, 58, 61] has gained prominence. Applications of machine unlearning span various domains, including regression tasks [51], federated learning [24, 34, 35, 54, 56], graph neural networks [5, 7], and diffusion models [12, 15, 16, 26, 57, 60, 62], as well as scenarios where training data are not available [8, 50].

Retraining the model from scratch without forgetting data is typically considered the standard gold unlearning algorithm [12, 52]. However, this is often deemed impractical as most in-production models require extensive training periods and considerable computing resources. While fine-tuning the model for a new task might lead to forgetting previous knowledge (i.e., catastrophic forgetting) [36], it fails to adequately erase the data influence in the models. Approximate unlearning methods, as such, become attractive alternatives. We will briefly discuss these methods in image classification and image generation.

Unlearning in Classification Models. Most unlearning algorithms are based on the influence function [23, 38, 48] and the Fisher Information Matrix (FIM) [19–21]. Influence function-based methods estimate the influence of the particular training data points on the model’s predictions, and Fisher unlearning methods assume that the unlearned model and the retrained model are close to each other, Golatkar et al. [20] leverage FIM and hide the difference between the unlearned models and retrained models via adding noise. Golatkar et al. [20] also present an upper bound for the amount of retained information, offering a quantifiable measure of the effectiveness of the unlearning process. Such approximate unlearning methods need to compute the Hessian matrix, or FIM w.r.t. the data. This may render them impractical for certain scenarios, such as federated learning, where computational resources are distributed and limited. To mitigate these computational demands, Mehta et al. [37] introduce L-CODEC, a strategy aimed at

pre-selecting a subset of parameters, allowing for more efficient computation. Jia et al. [27] recently propose to fuse the model sparsity into the unlearning algorithm, helping improve approximate unlearning methods, and achieving effective and efficient unlearning effects.

Concept Erasure in Generative Models. The advent of generative models, particularly those converting text to images, has been a significant milestone in the field of artificial intelligence. A notable concern is the risk of it being tainted or manipulated [6,43], resulting in not-safe-for-work (NSFW) generations [46], as these models leverage training data from a wide array of open sources. To address these challenges, data censoring [1, 14, 41, 47] to exclude black-listed images is employed. Studies [15, 46] introduce methods to update the models away from the inappropriate concepts. Heng and Soh [26] recently propose to adopt Elastic Weight Consolidation (EWC) and Generative Replay (GR) to effectively unlearn without access to the training data for a wide range of generative models.

Unlearning across Domains. Recently, Fan et al. [12] highlight the difficulty in cross-domain applicability of the machine unlearning algorithms, i.e., existing machine unlearning methods designed for classification tasks cannot be effective when applied to image generation. To address this gap, they introduce a novel method, Saliency unlearning (SalUn), which shifts attention from the entire model to target parameters. SalUn achieves state-of-the-art performance and is effective in both image classification and image generation tasks.

In this work, we propose *Scissorhands* which could achieve improved performance compared to existing methods across various scenarios. Inspired by [27] that apply ℓ_1 regularization to simply fuse sparsity into the unlearning methods, we identify and re-initialize the key model parameters w.r.t. the forgetting data via connection sensitivity. Concurrently, Fan et al. [12] also explore unlearning via saliency scores, determining weight saliency via the gradient of the forgetting loss w.r.t. model parameters, and employing random labeling for model fine-tuning. While there are similarities in our overarching goals, the methodologies present unique perspectives and solutions to the problem of unlearning. Fan et al. [12] only focus on the salient weights, while our mechanism aims to first ‘destroy’ the key information and then relearn for the whole model. Additionally, using the forgetting data with random labels could potentially lead to memorizing information about the forgetting data points [59]. In contrast, *Scissorhands* excludes knowledge about the forgetting data via a gradient projection-based approach, resulting in the effective erasure of information about the forgetting data while preserving the knowledge about the retained data.

4 Experimental Evaluation

In this section, we illustrate how *Scissorhands* effectively eliminates the data influence in the models. For sample-wise forgetting, where the forgetting data shares the same distribution as the training data, we evaluate *Scissorhands* on SVHN [39], CIFAR-10 and CIFAR-100 [30]. We further extend our evaluation on CelebAMask-HQ [31] where targeting to remove the entirety of specific identities.

Additionally, our experimentation encompasses the open-source text-to-image model, Stable Diffusion v1.4 [44], which is conditioned on CLIP text embeddings through the cross-attention mechanism. Further detailed experimental setups and additional results can be found in Appendix B.

Baselines. We primarily compare against the following standard baselines that are frequently employed in machine unlearning, alongside recently proposed SOTA methods: (i) **Retrain**: models obtained by only retraining from scratch on \mathcal{D}_r . This will provide the performance of the oracle for us. (ii) **Fine-tuning (FT)** [55]: models that are fine-tuned on \mathcal{D}_r , *i.e.*, taking advantage from catastrophic forgetting in neural networks to unlearn. (iii) **Gradient ascent (GA)** [52]: gradient ascent on \mathcal{D}_f . This will provide a simple way of unlearning by making the performance of the model worse on \mathcal{D}_f . (iv) **Influence unlearning (IU)** [29]: utilizes the influence function method to estimate the change in model parameters when transitioning from the unscrubbed model to the retrained model. (v) **Boundary shrink (BS)** [4] and (vi) **Boundary expanding (BE)** [4]: shift the decision boundary of the original model to imitate that of the retrained model. (vii) ℓ_1 -**sparse** [27]: fine-tuning models on \mathcal{D}_r with ℓ_1 -norm sparse regularization. (viii) **Saliency unlearning (SalUn)** [12]: adopt weight saliency and random labeling for unlearning.

Metrics. To assess the unlearning algorithms, we employ several metrics: (i) **Accuracy**: the accuracy of the model on \mathcal{D}_f (denoted as $\mathbf{Acc}_{\mathcal{D}_f}$), \mathcal{D}_r (denoted as $\mathbf{Acc}_{\mathcal{D}_r}$) and the test set (denoted as $\mathbf{Acc}_{\mathcal{D}_t}$). It provides insight into how well the model performs after undergoing the unlearning process. (ii) **Membership inference attack (MIA)**: a standard metric for verifying the unlearning effect. A classifier ϕ is trained using both training data (marked with a label of 1, indicating data points included in training) and test data (marked with a label of 0, indicating data points not seen during training), and subsequently assessed on \mathcal{D}_f . An effective unlearning method should make it challenging to infer whether a particular sample was part of the training set, mitigating potential privacy breaches. In this context, we define MIA as $P_\phi(\mathbf{y} = 0|\mathbf{x}_f)$, representing the probability assigned by the classifier ϕ that a given sample \mathbf{x}_f was not included in the training data. (iii) **Avg. Gap** [12]: the average of the performance gaps measured in accuracy-related metrics: $\text{Avg. Gap} = (|\mathbf{Acc}_{\mathcal{D}_t} - \mathbf{Acc}_{\mathcal{D}_t}^*| + |\mathbf{Acc}_{\mathcal{D}_f} - \mathbf{Acc}_{\mathcal{D}_f}^*| + |\mathbf{Acc}_{\mathcal{D}_r} - \mathbf{Acc}_{\mathcal{D}_r}^*| + |\text{MIA} - \text{MIA}^*|)/4$, where $\mathbf{Acc}_{\mathcal{D}_t}^*$, $\mathbf{Acc}_{\mathcal{D}_f}^*$, $\mathbf{Acc}_{\mathcal{D}_r}^*$ and MIA^* are metric values of the retrained model. A better performance of an unlearning method corresponds to a lower performance gap with retraining, measuring how close the unlearned model is to the retrained model. (iv) **Run-time efficiency (RTE)** [27]: evaluate the computational efficiency of the unlearning method. Specifically, we have $\text{RTE} = T/T_r$, where T and T_r present the time for the unlearning methods to get the scrubbed models and the time for retraining from scratch with \mathcal{D}_r , respectively. An efficient unlearning method should aim for minimal computational overhead compared to the baseline. (v) **Relearn time** [8, 20, 51]: the number of epochs to fine-tune the scrubbed models for regaining performance on \mathcal{D}_f (*i.e.*, reach the original models' accuracy on \mathcal{D}_f).

Table 1: Quantitative results for forgetting 10% data.

	Method	$\text{Acc}_{\mathcal{D}_f}(\downarrow)$	$\text{Acc}_{\mathcal{D}_t}(\uparrow)$	$\text{Acc}_{\mathcal{D}_r}(\uparrow)$	MIA(\uparrow)	Avg. Gap
CIFAR-100	Retrain	75.13 \pm 0.85	74.69 \pm 0.08	99.98 \pm 0.01	50.22 \pm 0.62	-
	FT [55]	97.98 \pm 1.36	75.28 \pm 0.12	99.95 \pm 0.02	9.64 \pm 3.6	16.01
	GA [52]	98.00 \pm 1.34	75.59 \pm 0.11	98.24 \pm 1.16	5.00 \pm 2.25	17.68
	IU [29]	95.67 \pm 4.82	72.13 \pm 4.58	96.14 \pm 4.51	9.43 \pm 5.98	16.93
	BE [4]	97.94 \pm 1.38	74.16 \pm 0.09	98.12 \pm 1.24	7.60 \pm 3.05	16.96
	BS [4]	97.65 \pm 1.48	73.20 \pm 0.18	97.93 \pm 1.30	8.24 \pm 3.23	17.01
	ℓ_1 -sparse [27]	96.35 \pm 0.67	70.06 \pm 0.46	96.35 \pm 0.67	21.33 \pm 1.95	14.59
	SalUn [12]	88.56 \pm 1.18	71.34 \pm 0.48	99.40 \pm 0.35	74.66 \pm 2.48	10.45
	Ours	68.76 \pm 1.81	73.17 \pm 0.24	99.24 \pm 0.30	42.42 \pm 2.06	4.11
CIFAR-10	Retrain	94.81 \pm 0.53	94.26 \pm 0.14	100.0 \pm 0.00	13.05 \pm 0.64	-
	FT [55]	99.15 \pm 0.46	93.83 \pm 0.45	99.84 \pm 0.11	3.01 \pm 0.93	3.74
	GA [52]	99.66 \pm 0.23	94.57 \pm 0.01	99.62 \pm 0.25	0.91 \pm 0.29	4.42
	IU [29]	98.08 \pm 2.10	91.91 \pm 2.73	98.01 \pm 2.26	4.01 \pm 3.44	4.16
	BE [4]	99.41 \pm 0.38	93.79 \pm 0.15	99.41 \pm 0.38	16.16 \pm 0.78	2.19
	BS [4]	99.60 \pm 0.25	94.24 \pm 0.07	99.56 \pm 0.54	4.46 \pm 0.33	3.46
	ℓ_1 -sparse [27]	94.17 \pm 0.49	90.64 \pm 0.52	96.64 \pm 0.54	11.87 \pm 0.61	2.20
	SalUn [12]	98.07 \pm 0.42	93.92 \pm 0.25	99.89 \pm 0.07	17.93 \pm 0.37	2.15
	Ours	95.40 \pm 1.48	92.92 \pm 0.48	98.93 \pm 0.57	9.56 \pm 2.13	1.62
SVHN	Retrain	91.81 \pm 1.11	91.17 \pm 1.77	97.73 \pm 1.12	15.74 \pm 1.28	-
	FT [55]	99.60 \pm 0.24	95.28 \pm 0.04	99.99 \pm 0.00	3.85 \pm 0.4	6.51
	GA [52]	98.41 \pm 0.23	92.87 \pm 0.06	98.52 \pm 0.29	5.96 \pm 0.38	4.72
	IU [29]	92.47 \pm 1.62	86.92 \pm 2.06	93.36 \pm 1.82	17.00 \pm 2.67	2.63
	BE [4]	98.48 \pm 0.33	92.62 \pm 0.05	98.44 \pm 0.29	6.77 \pm 0.36	4.45
	BS [4]	98.29 \pm 0.32	92.48 \pm 0.02	98.36 \pm 0.29	6.74 \pm 0.39	4.36
	ℓ_1 -sparse [27]	97.79 \pm 0.10	93.59 \pm 3.18	99.49 \pm 0.23	7.57 \pm 0.29	4.65
	SalUn [12]	96.29 \pm 4.80	93.59 \pm 3.18	97.41 \pm 4.96	24.47 \pm 4.08	3.99
	Ours	91.62 \pm 1.18	93.05 \pm 0.56	99.51 \pm 0.14	17.62 \pm 2.54	1.43


4.1 Results on Classification Task

In this experiment on SVHN, CIFAR-10, and CIFAR-100, we try to forget randomly selected 10% of the training data; on CelebAMask-HQ, we attempt to forget randomly selected 10% identities among 307 identities. We further apply GradCAM [49] to visualize regions where models focus on **w/** and **w/o** machine unlearning algorithms. In brief, the results suggest that *Scissorhands* has successfully induced forgetting for the relevant samples and classes, with minor degradation in model performance over the remaining data and classes. We will discuss the results in depth below.

Sample-wise unlearning. Table 1 presents the results when forgetting randomly selected samples. *Scissorhands* achieves the lowest average performance gap on all the presented datasets and shows good generalization ability compared to

Table 2: Quantitative results for forgetting 10% identities on the CelebAMask-HQ.

Method	$\text{Acc}_{\mathcal{D}_f}(\downarrow)$	$\text{Acc}_{\mathcal{D}_t}(\uparrow)$	$\text{Acc}_{\mathcal{D}_r}(\uparrow)$	MIA(\uparrow)	Avg. Gap
Retrain	0.00±0.00	87.02±0.80	99.96±0.01	100.0±0.00	-
FT [55]	99.94±0.12	88.59±0.59	99.97±7.02	5.28±2.03	49.06
GA [52]	87.60±8.71	81.22±2.11	99.74±0.26	51.37±5.96	35.56
IU [29]	88.92±10.25	70.24±11.77	95.27±5.07	29.59±18.59	45.20
BE [4]	69.07±2.73	44.11±2.08	95.58±1.23	46.24±5.90	42.53
BS [4]	98.18±1.92	81.92±0.27	99.86±0.03	45.93±5.11	39.36
ℓ_1 -sparse [27]	98.81±0.72	89.37±0.70	99.97±0.00	76.78±5.66	31.10
SalUn [12]	0.00±0.00	78.36±1.34	96.90±1.11	100.0±0.00	2.93
Ours	1.52±2.73	80.18±6.60	97.20±3.81	99.83±0.35	2.82

**Fig. 1:** Visualizations of regions where models focus on generated by GradCAM [49].

Retrain. For example, on SVHN, *Scissorhands* achieves $\sim 91.62\%$ accuracy on the forgetting data \mathcal{D}_f , closely matching the 91.81% accuracy of a fully retrained model. Furthermore, the retrained model reaches around 91% accuracy on the test dataset, *Scissorhands* exhibits an accuracy of roughly 93%, showcasing the enhanced generalization of our scrubbed model.

When evaluating solely based on forgetting accuracy on CIFAR-10, the ℓ_1 -sparse method might seem to be the most effective baseline, achieving approximately 94.17%. Nonetheless, this perceived advantage is offset by a decline in both test accuracy and the accuracy of retained data. In contrast, *Scissorhands* shows a good trade-off between preserving model utility on the remaining data, generalization on unseen sets, and erasing the influence of the forgetting data.

Similarly, on CIFAR-100, SalUn may initially stand out with its impressive MIA accuracy of approximately 74.66%, relying solely on a single metric to gauge the performance of machine unlearning can be misleading, as it might not fully capture the method’s effectiveness [12]. While SalUn achieves a forgetting accuracy of about 88.56% and a test accuracy of 71.34%, *Scissorhands* demonstrates a slightly lower forgetting accuracy of around 68.76% but achieves a higher test accuracy of 73.17%. This indicates that *Scissorhands* achieves the best trade-off between eradicating data traces and generalizing to unseen data. Soon, we will show that on this dataset, the relearning time of *Scissorhands* is greater than 200 to regain the performance on \mathcal{D}_f , which indicates that *Scissorhands* has efficiently scrubbed information about \mathcal{D}_f .

Overall, *Scissorhands* illustrates its capability to adeptly balance the objectives of data removal, model utility preservation, and generalization to unseen data, ensuring that the influence of forgotten data is minimized without significantly compromising the overall performance and utility of the model.

Class-wise unlearning. Furthermore, Table 2 presents the outcomes of our analysis for the scenario where forgetting 10% of identities on the CelebAMask-HQ. Images are rescaled to 224×224 , and a model pre-trained with ImageNet1K [9] is employed. In this context, *Scissorhands* manifests the smallest average performance gap with retrained models. Notably, the Retrain method achieves perfect metrics in terms of MIA accuracy, which is 100%, and completely erases the identity information from the dataset, as indicated by a 0.00% accuracy on the forgotten data \mathcal{D}_f . However, retraining proves impractical and limited for unlearning purposes in general, functioning merely as an oracle in our experiments.


Among the baselines, fine-tuning (FT) and ℓ_1 -sparse exhibit higher accuracies on the remaining data \mathcal{D}_r and unseen data \mathcal{D}_t ; yet, their accuracies on forgetting remain high as well, highlighting a less effective balance between eradicating data traces and generalizing to unseen data. In contrast, *Scissorhands* and SalUn demonstrate superior capabilities in effectively forgetting identities with minimal impact on the models’ overall performance. SalUn achieves an MIA accuracy of 100% and a forgetting accuracy of 0.00%, effectively nullifying the identity information similar to the retrained models. *Scissorhands* slightly surpasses SalUn in terms of the average performance gap with retrained models, maintaining a high MIA accuracy of 99.83% and a significantly low accuracy on the forgotten data \mathcal{D}_f (i.e., 1.52%), while achieving a test accuracy of 80.18% and an accuracy of 97.20% on the retained data.

We further employ the GradCAM [49] to illustrate the focus area of models $\mathbf{w}/$ and \mathbf{w}/\mathbf{o} machine unlearning algorithms. As shown in Figure 1, when evaluated on the remaining data \mathcal{D}_r , both SalUn and *Scissorhands* even pay more attention to the facial features for identity classification than the retrained model. Conversely, when evaluated on the forgetting data \mathcal{D}_f , the scrubbed model by *Scissorhands* shifts its attention towards regions that are least associated with the facial features. We hypothesize that in the scenario where the task involves forgetting identities, both the remaining data and the forgetting data are face images, the scrubbed models would need a more nuanced approach to distinguish between individuals. Specifically, because the task at hand requires the discernment of subtle differences across facial features to effectively differentiate and forget specific identities while retaining others, the models adapt by developing a refined sensitivity towards those facial characteristics that are most indicative of individual identity. This adaptation enables the models to maintain high accuracy on the remaining data by focusing more intently on the facial features that are crucial for identity recognition, thereby enhancing their ability to generalize and accurately classify identities.

These results underscore our proposed machine unlearning algorithm *Scissorhands*’s superior capabilities in balancing identity forgetting and model utility

Table 3: Influence of projection and initialization strategies (re-initializing parameters with uniform distribution \mathcal{U} , Gaussian distribution \mathcal{N} , and constant value of $\mathbf{0}/\mathbf{1}$).

Projection	Initialization	Acc $_{\mathcal{D}_f}(\downarrow)$	Acc $_{\mathcal{D}_t}(\uparrow)$	Acc $_{\mathcal{D}_r}(\uparrow)$	MIA(\uparrow)	Avg. Gap
✓	\mathcal{U}	95.40 \pm 1.48	92.92 \pm 0.48	98.93 \pm 0.57	9.56 \pm 2.13	1.62
✗	\mathcal{U}	98.14 \pm 0.60	93.14 \pm 0.16	99.73 \pm 0.16	6.08 \pm 0.78	2.92
✓	\mathcal{N}	94.68 \pm 2.50	92.13 \pm 2.20	97.94 \pm 2.71	10.18 \pm 3.81	1.80
✓	$\mathbf{0}$	96.10 \pm 2.05	92.92 \pm 0.99	99.09 \pm 1.19	8.76 \pm 3.30	1.96
✗	$\mathbf{1}$	84.28 \pm 0.49	83.38 \pm 0.58	87.18 \pm 0.90	20.67 \pm 0.41	10.46

**Fig. 2:** Influence of the percent value k in the trimming process, the balance term λ of Equation (3), and the ratio of forgetting data used in the trimming process.

performance, *Scissorhands* not only minimizes privacy risks but also maintains the integrity and applicability of the model to unseen data.

4.2 Ablation Study

Component analysis. As detailed in Table 3, the inclusion of different procedures in *Scissorhands* has a noticeable impact. Notably, the employment of the projection mechanism plays a crucial role in eliminating the influence of forgetting data \mathcal{D}_f , affecting the overall utility of the model. The rationale behind the projection process is to preclude the reacquisition of critical information about the forgetting data \mathcal{D}_f during the repair phase. Moreover, Table 3 reveals that initializing parameters with a uniform distribution \mathcal{U} or opting for zero initialization facilitates effective unlearning, in contrast to re-initializing parameters with a constant value of one, which impedes finding a viable solution to Equation (5).

Hyper-parameter impact. Figure 2 presents the effects of varying hyperparameters, such as the percentage value $k\%$ for the trimming process, the balance term λ of Equation (3) in the repairing phase, and the ratio of forgetting data used in the trimming process. An increase k results in higher accuracies on data, as fewer parameters associated with the forgetting data \mathcal{D}_f are re-initialized. Empirically, we observe choosing a value of $k \geq 0.9$ will result in stable performances. That said, this is a hyperparameter of the algorithm and can be chosen by measuring the performances on \mathcal{D}_f and \mathcal{D}_r during training. The hyper-parameter λ controls the balance between retaining and forgetting data, the accuracies on test data and remaining data show a peak at $\lambda = 0.05$

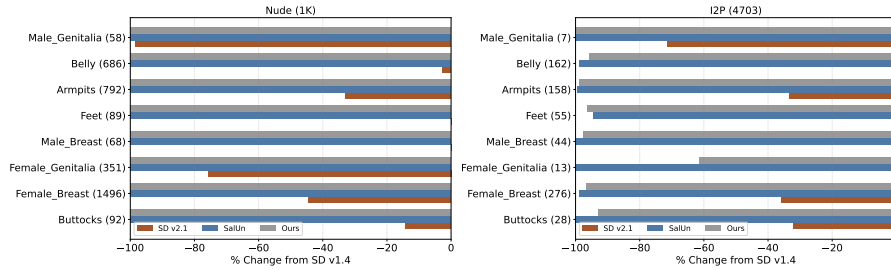


Fig. 3: Quantity of nudity content detected using the NudeNet classifier from 1K sampled images and I2P data. We observed a high false positive rate for exposed female genitalia/breast using the NudeNet classifier on generated I2P images. The flagged images can be found in Appendix B.




Fig. 4: Sample images with the prompt from $c_f = \{\text{'nudity'}, \text{'naked'}, \text{'erotic'}, \text{'sexual'}\}$ generated by SDs w/ and w/o machine unlearning algorithms. Best viewed in color.

when forgetting 10% data on CIFAR-10. Generally, as λ escalates, it places greater emphasis on the process of data omission, leading to a decrease in forgetting accuracy. As the ratio p increases, all accuracy metrics exhibit a decline, suggesting that employing more forgetting data for the purpose of forgetting in order to isolate critical knowledge about the forgetting data also inadvertently results in the loss of pertinent information regarding the remaining data. This phenomenon occurs particularly when the data designated for forgetting shares the same distribution as the data that is retained. In essence, while aiming to enhance the specificity of the forgetting process by pinpointing essential details about the data to be forgotten, there’s a risk of simultaneously diminishing knowledge pertinent to the retained data.

4.3 Case Study: Stable Diffusion

We also employed our algorithm to mitigate the generation of inappropriate content in Stable Diffusion (SD) models [44]. Specifically, our primary objective is to effectively eliminate the concept of nudity from the model’s generative capabilities. To this end, followed [26] and employed SD v1.4 for sampling with 50 time steps. We evaluate on 1K generated images with prompts $c_f = \{\text{'nudity'}, \text{'naked'}, \text{'erotic'}, \text{'sexual'}\}$ and 4703 generated images with I2P [46] using the open-source NudeNet classifier.

Figures 3 and 4 present the results for SD v2.1 (trained on a dataset filtered for nudity), SalUn, and *Scissorhands*. In Figure 3, y-axis denotes the number of

exposed body parts generated by the SD v1.4 model, it presents the percentage change in exposed body parts w.r.t. SD v1.4. When comparing on I2P data, our algorithms outperform SalUn on ‘Feet’, while coming short on ‘Female Breast and Genitalia’. While this could suggest the superiority of SalUn on this data, the gap is not that big in reality, as out of the 9 images flagged for containing exposed female breasts, 7 were inaccurately identified, and none of the flagged images depicted exposed female genitalia. The use of the NudeNet classifier on generated I2P images exhibits a significant rate of false positives, as highlighted in [26] as well. When we use prompts c_f explicitly associated with nudity, no exposed sensitive content is detected for both *Scissorhands* and SalUn methods. We can conclude that, *Scissorhands* significantly reduces the amount of nudity content compared to SD v1.4, SD v2.1, showing comparable results to SalUn.

5 Limitations and Broader Impact

Our proposed unlearning algorithm, *Scissorhands*, can achieve comparable performance to state-of-the-art unlearning methods. We addressed potential issues of conflicting gradients arising from direct conflicts between features or patterns in the forgetting and remaining data, ensuring training stability. However, the limitation is the incomplete concept erasure shown in Figure 8 in Appendix B, and the potential for biased models due to the selective deletion of information. This is a crucial consideration, especially when unlearning algorithms are applied with malicious intent. Furthermore, in experiments on stable diffusion, while our method can successfully erase specific concepts, finding an optimal balance between forgetting and maintaining high image quality remains challenging. There is also a risk that unlearning algorithms might be used to alter or remove concepts in unethical or malicious ways. It is our hope that *Scissorhands* offers a fresh perspective on practical and effective machine unlearning, emphasizing the need for ethical guidelines to prevent misuse in its application.

6 Conclusion

In this work, we proposed *Scissorhands*, an effective and practical machine unlearning algorithm. Our algorithm identified the important parameters w.r.t. the forgetting data in single-shot via connection sensitivity, then re-initialized these parameters for scrubbing the influence of the forgetting data. An effective gradient projection-based technique is applied to enhance the model utility while excluding the information w.r.t. the forgetting data. Our evaluations on various datasets showed that, compared with other methods, our approach offers superior performance across image classification and image generation tasks. Future works could explore how *Scissorhands* performs on regression and NLP tasks, and further investigation into the scenario where data are not available, as well as unlearning with sequential/time-series data where could introduce unique challenges, e.g., the need to consider temporal dependencies, which may require new frameworks specifically tailored for sequential data.

Acknowledgements

Mehrtash Harandi is supported by funding from the Australian Research Council (ARC) Discovery Program DP230101176.

References

1. Birhane, A., Prabhu, V.U.: Large image datasets: A pyrrhic win for computer vision? In: 2021 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1536–1546. IEEE (2021)
2. Bourtole, L., Chandrasekaran, V., Choquette-Choo, C.A., Jia, H., Travers, A., Zhang, B., Lie, D., Papernot, N.: Machine unlearning. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 141–159 (2021)
3. Cao, Y., Yang, J.: Towards making systems forget with machine unlearning. In: 2015 IEEE Symposium on Security and Privacy (SP). pp. 463–480 (2015)
4. Chen, M., Gao, W., Liu, G., Peng, K., Wang, C.: Boundary unlearning: Rapid forgetting of deep networks via shifting the decision boundary. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7766–7775 (2023)
5. Chen, M., Zhang, Z., Wang, T., Backes, M., Humbert, M., Zhang, Y.: Graph unlearning. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. pp. 499–513 (2022)
6. Chen, W., Song, D., Li, B.: Trojdiff: Trojan attacks on diffusion models with diverse targets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4035–4044 (2023)
7. Cheng, J., Dasoulas, G., He, H., Agarwal, C., Zitnik, M.: Gnndelete: A general strategy for unlearning in graph neural networks. In: International Conference on Learning Representations (ICLR) (2023)
8. Chundawat, V.S., Tarun, A.K., Mandal, M., Kankanhalli, M.: Zero-shot machine unlearning. *IEEE Transactions on Information Forensics and Security* (2023)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
10. Dennis, C.R., Weisberg, S.: Residuals and influence in regression (1982)
11. Fan, C., Liu, J., Hero, A., Liu, S.: Challenging forgets: Unveiling the worst-case forget sets in machine unlearning. *arXiv preprint arXiv:2403.07362* (2024)
12. Fan, C., Liu, J., Zhang, Y., Wei, D., Wong, E., Liu, S.: Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. In: International Conference on Learning Representations (ICLR) (2024)
13. Foster, J., Schoepf, S., Brintrup, A.: Fast machine unlearning without retraining through selective synaptic dampening. *arXiv preprint arXiv:2308.07707* (2023)
14. Gandhi, S., Kokkula, S., Chaudhuri, A., Magnani, A., Stanley, T., Ahmadi, B., Kandaswamy, V., Ovenc, O., Mannor, S.: Scalable detection of offensive and non-compliant content/logo in product images. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2247–2256 (2020)
15. Gandikota, R., Materzynska, J., Fiotto-Kaufman, J., Bau, D.: Erasing concepts from diffusion models. In: 2023 IEEE International Conference on Computer Vision (ICCV) (2023)

16. Gandikota, R., Orgad, H., Belinkov, Y., Materzyńska, J., Bau, D.: Unified concept editing in diffusion models. *arXiv preprint arXiv:2308.14761* (2023)
17. Ginart, A., Guan, M., Valiant, G., Zou, J.Y.: Making ai forget you: Data deletion in machine learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 32 (2019)
18. Goel, S., Prabhu, A., Sanyal, A., Lim, S.N., Torr, P., Kumaraguru, P.: Towards adversarial evaluations for inexact machine unlearning. *arXiv preprint arXiv:2201.06640* (2022)
19. Golatkar, A., Achille, A., Ravichandran, A., Polito, M., Soatto, S.: Mixed-privacy forgetting in deep networks. In: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 792–801 (2021)
20. Golatkar, A., Achille, A., Soatto, S.: Eternal sunshine of the spotless net: Selective forgetting in deep networks. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9301–9309 (2020)
21. Golatkar, A., Achille, A., Soatto, S.: Forgetting outside the box: Scrubbing deep networks of information accessible from input-output observations. In: *European Conference on Computer Vision (ECCV)*. pp. 383–398. Springer (2020)
22. Goldman, E.: An introduction to the california consumer privacy act (ccpa). Santa Clara Univ. Legal Studies Research Paper (2020)
23. Guo, C., Goldstein, T., Hannun, A., Van Der Maaten, L.: Certified data removal from machine learning models. In: *Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 119, pp. 3832–3842. PMLR (2020)
24. Halimi, A., Kadhe, S., Rawat, A., Baracaldo, N.: Federated unlearning: How to efficiently erase a client in fl? In: *International conference on machine learning*. PMLR (2022)
25. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. vol. 28 (2015)
26. Heng, A., Soh, H.: Selective amnesia: A continual learning approach to forgetting in deep generative models. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2023)
27. Jia, J., Liu, J., Ram, P., Yao, Y., Liu, G., Liu, Y., Sharma, P., Liu, S.: Model sparsification can simplify machine unlearning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2023)
28. Karasuyama, M., Takeuchi, I.: Multiple incremental decremental learning of support vector machines. *IEEE Transactions on Neural Networks* **21**(7), 1048–1059 (2010)
29. Koh, P.W., Liang, P.: Understanding black-box predictions via influence functions. In: *International conference on machine learning*. pp. 1885–1894. PMLR (2017)
30. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
31. Lee, C.H., Liu, Z., Wu, L., Luo, P.: Maskgan: Towards diverse and interactive facial image manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5549–5558 (2020)
32. Lee, N., Ajanthan, T., Gould, S., Torr, P.H.: A signal propagation perspective for pruning neural networks at initialization. In: *International Conference on Learning Representations (ICLR)* (2020)
33. Lee, N., Ajanthan, T., Torr, P.H.: Snip: Single-shot network pruning based on connection sensitivity. In: *International Conference on Learning Representations (ICLR)* (2019)

34. Liu, G., Ma, X., Yang, Y., Wang, C., Liu, J.: Federaser: Enabling efficient client-level data removal from federated learning models. In: 2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS). pp. 1–10 (2021)
35. Liu, Y., Xu, L., Yuan, X., Wang, C., Li, B.: The right to be forgotten in federated learning: An efficient realization with rapid retraining. In: IEEE INFOCOM 2022 - IEEE Conference on Computer Communications. pp. 1749–1758 (2022)
36. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. *Advances in neural information processing systems* **30** (2017)
37. Mehta, R., Pal, S., Singh, V., Ravi, S.N.: Deep unlearning via randomized conditionally independent Hessians. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10412–10421 (2022)
38. Neel, S., Roth, A., Sharifi-Malvajerdi, S.: Descent-to-delete: Gradient-based methods for machine unlearning. In: Proceedings of the 32nd International Conference on Algorithmic Learning Theory. pp. 931–962. Proceedings of Machine Learning Research, PMLR (2021)
39. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
40. Nguyen, T.T., Huynh, T.T., Nguyen, P.L., Liew, A.W.C., Yin, H., Nguyen, Q.V.H.: A survey of machine unlearning. arXiv preprint arXiv:2209.02299 (2022)
41. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
42. Peste, A., Alistarh, D., Lampert, C.H.: Ssse: Efficiently erasing samples from trained machine learning models. arXiv preprint arXiv:2107.03860 (2021)
43. Rando, J., Paleka, D., Lindner, D., Heim, L., Tramèr, F.: Red-teaming the stable diffusion safety filter. arXiv preprint arXiv:2210.04610 (2022)
44. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
45. Romero, E., Barrio, I., Belanche, L.: Incremental and decremental learning for linear support vector machines. In: International Conference on Artificial Neural Networks. pp. 209–218. Springer (2007)
46. Schramowski, P., Brack, M., Deiseroth, B., Kersting, K.: Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22522–22531 (2023)
47. Schramowski, P., Tauchmann, C., Kersting, K.: Can machines help us answering question 16 in datasheets, and in turn reflecting on inappropriate content? In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency. pp. 1350–1361 (2022)
48. Sekhari, A., Acharya, J., Kamath, G., Suresh, A.T.: Remember what you want to forget: Algorithms for machine unlearning. *Advances in Neural Information Processing Systems (NeurIPS)* **34**, 18075–18086 (2021)
49. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 618–626 (2017)
50. Tarun, A.K., Chundawat, V.S., Mandal, M., Kankanhalli, M.: Fast yet effective machine unlearning. *IEEE Transactions on Neural Networks and Learning Systems* (2023)

51. Tarun, A.K., Chundawat, V.S., Mandal, M., Kankanhalli, M.: Deep regression unlearning. In: International Conference on Machine Learning. pp. 33921–33939. PMLR (2023)
52. Thudi, A., Deza, G., Chandrasekaran, V., Papernot, N.: Unrolling sgd: Understanding factors influencing machine unlearning. In: 2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P). pp. 303–319. IEEE (2022)
53. Voigt, P., Von dem Bussche, A.: The eu general data protection regulation (gdpr). A Practical Guide, 1st Ed., Cham: Springer International Publishing **10**(3152676), 10–5555 (2017)
54. Wang, J., Guo, S., Xie, X., Qi, H.: Federated unlearning via class-discriminative pruning. In: Proceedings of the ACM Web Conference 2022. pp. 622–632 (2022)
55. Warnecke, A., Pirch, L., Wressnegger, C., Rieck, K.: Machine unlearning of features and labels. arXiv preprint arXiv:2108.11577 (2021)
56. Wu, C., Zhu, S., Mitra, P.: Federated unlearning with knowledge distillation. arXiv preprint arXiv:2201.09441 (2022)
57. Wu, J., Le, T., Hayat, M., Harandi, M.: Erasediff: Erasing data influence in diffusion models. arXiv preprint arXiv:2401.05779 (2024)
58. Wu, Y., Dobriban, E., Davidson, S.: DeltaGrad: Rapid retraining of machine learning models. In: Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 119, pp. 10355–10366. PMLR (13–18 Jul 2020)
59. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning (still) requires rethinking generalization. In: International Conference on Learning Representations (ICLR) (2017)
60. Zhang, E., Wang, K., Xu, X., Wang, Z., Shi, H.: Forget-me-not: Learning to forget in text-to-image diffusion models. arXiv preprint arXiv:2303.17591 (2023)
61. Zhang, Y., Zhang, Y., Yao, Y., Jia, J., Liu, J., Liu, X., Liu, S.: Unlearncanvas: A stylized image dataset to benchmark machine unlearning for diffusion models. arXiv preprint arXiv:2402.11846 (2024)
62. Zhang, Y., Chen, X., Jia, J., Zhang, Y., Fan, C., Liu, J., Hong, M., Ding, K., Liu, S.: Defensive unlearning with adversarial training for robust concept erasure in diffusion models. arXiv preprint arXiv:2405.15234 (2024)