UNO: Unlearning via Orthogonalization in Generative models

Pinak Mandal[*]
University of Sydney

Georg A. Gottwald[†]
University of Sydney

Abstract

As generative models become increasingly powerful and pervasive, the ability to unlearn specific data, whether due to privacy concerns, legal requirements, or the correction of harmful content, has become increasingly important. Unlike in conventional training, where data are accumulated and knowledge is reinforced, unlearning aims to selectively remove the influence of particular data points without costly retraining from scratch. To be effective and reliable, such algorithms need to achieve (i) forgetting of the undesired data, (ii) preservation of the quality of the generation, (iii) preservation of the influence of the desired training data on the model parameters, and (iv) small number of training steps. We propose fast unlearning algorithms based on loss gradient orthogonalization. We show that our algorithms are able to forget data while maintaining the fidelity of the original model. Using MNIST and CelebA data, we demonstrate that our algorithms achieve orders of magnitude faster unlearning times than their predecessors, such as gradient surgery.

1 Introduction

Machine learning models are often trained on datasets that contain personal or sensitive information, such as medical records, financial data, or social media activity [30, 39]. This reliance on personal data introduces substantial privacy risks, especially when models can unintentionally memorize or leak identifiable information; see [8] for an in-depth exploration of this issue in the context of large language models (LLMs). Legal frameworks such as the General Data Protection Regulation (GDPR) and related EU laws have been established to address these issues [1]. One of the central provisions is the right to be forgotten (RTBF), which grants individuals the ability to request the deletion of their personal data [24]. It is increasingly likely that this obligation will become a standard requirement for machine learning services. Retraining large models from scratch each time such a request is received is computationally infeasible since the training costs are substantial [5, 18]. Machine unlearning refers to removal of the influence of specific data points from a trained model without requiring full retraining. In the context of generative models this can be formalized as follows. Given a training dataset $\mathcal{D} = \mathcal{D}_r \sqcup \mathcal{D}_f$ partitioned into retain and forget datasets $\mathcal{D}_r$ and $\mathcal{D}_f$, respectively, and a model $\mathcal{M}_\theta$ trained on $\mathcal{D}$, the objective of unlearning is to update the model parameters θ in a way such that $P(\mathrm{sim}(x, \mathcal{D}_f) \geq \delta) \leq \varepsilon$ where P denotes probability, x is a sample generated by the updated model, sim is an appropriate similarity measure, and ε, δ are thresholds controlling the degree of forgetting. For an unlearning algorithm to be effective, it should (i) prevent the model from generating data resembling samples from $\mathcal{D}_f$, (ii) preserve the quality or fidelity of the generated samples, (iii) retain the influence of $\mathcal{D}_r$ on the model parameters, and (iv) require only a small number of training steps.

A simple approach to machine unlearning is to reverse the model update steps by performing gradient ascent on the loss computed over the forget dataset $\mathcal{D}_f$. However, this method is susceptible

[*]pinak.mandal@sydney.edu.au
[†]georg.gottwald@sydney.edu.au

Preprint. Under review.

to catastrophic forgetting, where the model loses knowledge far beyond just the targeted forget dataset [29, 28]. To mitigate this, several approaches combine gradient ascent on $\mathcal{D}_f$ with gradient descent on the retain dataset $\mathcal{D}_r$ [44]. The Gradient Difference (GDiff) method minimizes the difference of losses evaluated on the retain and forget datasets. Balancing the opposing updates in ascent-descent methods or weighing the loss terms properly in methods like GDiff is challenging since the forget and retain datasets might have significant size disparity, and the risk of catastrophic forgetting persists unless training hyperparameters such as the learning rate are finely tuned [6]. Recently, multi-task optimization (MTO) techniques [36, 45] have inspired several unlearning algorithms. One such algorithm is gradient surgery [2] where gradient ascent is performed in a direction that is orthogonal to the loss gradient computed over the retain dataset. While theoretically sound, this method remains sensitive to the choice of hyperparameters and can suffer from catastrophic forgetting without careful tuning, see Appendix 7 for an example.

In this work, we aim to advance the gradient surgery framework for unlearning in generative models. Although our proposed algorithms are general-purpose and presented accordingly, we demonstrate their effectiveness specifically using variational autoencoders (VAEs) [22, 34] for two widely used benchmark image datasets MNIST [11] and CelebA [27].

Contributions

Our main contributions are as follows:

1. We propose two new unlearning algorithms that regularize the main loss function with an additional term enforcing orthogonality between loss gradients computed over the retain and forget datasets.
2. We compare our algorithms against prior approaches, including gradient surgery and gradient ascent, evaluating both unlearning speed and the quality of generated samples. Our methods achieve orders of magnitude faster unlearning than gradient surgery, while retaining the influence of the desired training data unlike gradient ascent.
3. We explore unlearning in the presence of a classifier able to distinguish between the retain and forget data and assess its impact on accelerating unlearning algorithms.
4. We provide implementations of both the proposed and baseline algorithms, along with the experiment data, in this GitHub repository: https://github.com/pinakm9/forget.

2 Related work

Early foundational work by Koh and Liang [23] introduced influence functions as a principled approach for quantifying the impact of removing individual training points from machine learning models. Although influential, their technique is computationally demanding, limiting its scalability, particularly for large-scale neural networks [3, 15]. To address these computational challenges, recent studies have developed more efficient and scalable methodologies. For example, Schioppa et al. [35] and Guo et al. [15] proposed efficient approximations of influence functions that significantly reduce computational complexity. Further, innovative optimization-based frameworks such as SCRUB by Kurmanji et al. [25] approximate data removal for classification models such as ResNet [16] using a teacher-student distillation paradigm combined with checkpoint rewinding. Gradient-based methods have emerged as an effective paradigm for machine unlearning. Golatkar et al. [13] approximate the influence of individual data points on model parameters using the Fisher Information Matrix and use it to execute unlearning in deep networks. Building on this, Mixed-Privacy Forgetting [12] combines public and private data during training, enabling the selective removal of private data while preserving the utility of public data. Neel et al. [31] propose Descent-to-Delete, a gradient-based optimization technique that incrementally updates model parameters to approximate the behavior of a model trained without the forgotten data.

Unlearning in generative models introduces distinct challenges due to their capacity to implicitly memorize training data, complicating data removal without degrading generative quality. Addressing these, Bae et al. [37] introduced methods specifically designed to detect and mitigate unintended memorization in generative adversarial networks (GANs). Selective Amnesia [17], proposed by Heng and Soh, leverages continual learning frameworks to selectively remove specific concepts from deep generative models without compromising the overall data distribution learned by the model. In the

context of LLMs, recent works have tackled critical challenges such as selective forgetting of harmful or copyrighted content and aligning models to user preferences [20, 9, 33, 32]. These methods employ parameter-efficient fine-tuning, low-rank adaptations, and in-context learning strategies to remove specific learned knowledge while minimally impacting overall model performance.

Negative Preference Optimization (NPO) [46] offers an alignment-inspired approach to machine unlearning by assigning lower preference or likelihood to data from the forget set. Through preference-based training, the model learns to reduce its reliance on forget data, often using pairwise comparisons or preference signals. Normalized Gradient Difference (NGDiff) [6] approaches unlearning as a multi-task optimization problem, balancing the objectives of forgetting and retaining. By normalizing the gradient differences between these tasks and employing an adaptive learning rate scheduler, NGDiff provides stable training and effectively manages the trade-off between unlearning and model utility. Cao et al. [7] propose a projection residual based method to remove the influence of undesired data. In the same vein, gradient surgery [2] attempts to maximize the loss in a direction orthogonal to the loss gradient evaluated on the retain dataset. While promising for generative models, gradient surgery can suffer from inefficiency when there is significant overlap between loss gradients computed on the retain and forget data, or even cause catastrophic forgetting. We aim to improve upon this approach by explicitly enforcing orthogonality between these conflicting gradients. Our algorithms exhibit no catastrophic forgetting, achieve fast unlearning speeds, and are robust to hyperparameter selection.

For a comprehensive overview of unlearning techniques for large language models, including method categorization and scale-specific challenges, see Blanco-Justicia et al. [4]. For a broad taxonomy of machine unlearning across centralized, distributed, and privacy-critical settings with a focus on open problems and verification, see Wang et al. [41].

3 Unlearning via orthogonalization

We now describe the unlearning algorithms used to produce the results presented in this paper. The pseudocode for all the algorithms presented in this section can be found in Appendix 6.

3.1 Gradient ascent

We begin by introducing the most primitive approach, namely, gradient ascent. Given a model $\mathcal{M}_\theta$ with trainable parameters θ, trained using a loss function $\mathcal{L}$ on dataset $\mathcal{D}$, unlearning can be induced by maximizing the loss on the forget data, which can be done with the update step:

$$\theta_{k+1} = \theta_k + \eta \mathbf{g_f}, \tag{A}$$

where θ_k represents the model parameters after the k-th training step, η is the learning rate, and $\mathbf{g_f}$ is the gradient of the loss evaluated over the forget data (we omit the index of θ in the definition below for brevity),

$$\mathbf{g_f} = \frac{1}{|\mathcal{D}_f|} \sum_{x \in \mathcal{D}_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x). \tag{1}$$

This approach, however, may delete knowledge acquired on the retain data $\mathcal{D}_r$ if $\mathbf{g_f}$ resembles $\mathbf{g_r}$, the gradient of loss evaluated over the retain data,

$$\mathbf{g_r} = \frac{1}{|\mathcal{D}_r|} \sum_{x \in \mathcal{D}_r} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x). \tag{2}$$

A naive way to prevent the model from forgetting retain data is to perform alternating ascent in $\mathbf{g_f}$ direction and descent in $\mathbf{g_r}$:

$$\theta_{k+1} = \begin{cases} \theta_k + \eta \mathbf{g_f}, & \text{if } k \text{ is even,} \\ \theta_k - \eta \mathbf{g_r}, & \text{if } k \text{ is odd.} \end{cases} \tag{A-D}$$

This simple modification, however, does not safeguard against catastrophic forgetting, as we will see in Section 4.

3.2 Gradient surgery

Similar challenges also arise in a related subfield of machine learning: multi-task optimization where a model must learn to perform new tasks without compromising performance on earlier tasks [10]. If the loss gradient corresponding to the new task points in a direction opposing the loss gradients corresponding to the old tasks, the model risks losing its previously learned skills with each new gradient descent step, paralleling catastrophic forgetting. In multi-task optimization, gradient surgery refers to techniques that modify task-specific gradients during training to reduce this interference between tasks. When gradients from different tasks conflict, i.e., point in opposing directions, methods like PCGrad project gradients to minimize this conflict, allowing the model to learn multiple tasks more effectively without one task hindering the progress of another [45].

Gradient surgery can be used to reduce the potential conflict between $\mathbf{g_f}$ and $\mathbf{g_r}$ to improve the vanilla gradient ascent [2]. We remove the orthogonal projection of $\mathbf{g_r}$ from $\mathbf{g_f}$ before taking the ascent step,

$$\begin{aligned} \bar{\mathbf{g}}_\mathbf{f} &= \mathbf{g_f} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\mathbf{g_r} \cdot \mathbf{g_r}} \mathbf{g_r}, \\ \theta_{k+1} &= \theta_k + \eta \bar{\mathbf{g}}_\mathbf{f}. \end{aligned} \tag{SA}$$

While this modified ascent reduces over-unlearning compared to vanilla ascent, it does not fully resolve the issue, and still requires careful tuning of η to avoid catastrophic forgetting. Therefore, we introduce another version of gradient surgery which we find to be more stable and use it throughout, for generating the results in Section 4. Rather than perform ascent along modified $\mathbf{g_f}$ direction, we perform descent along modified $\mathbf{g_r}$ direction resulting in the following update:

$$\begin{aligned} \bar{\mathbf{g}}_\mathbf{r} &= \mathbf{g_r} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\mathbf{g_f} \cdot \mathbf{g_f}} \mathbf{g_f}, \\ \theta_{k+1} &= \theta_k - \eta \bar{\mathbf{g}}_\mathbf{r}, \end{aligned} \tag{S}$$

which aims at minimizing the loss in directions orthogonal to $\mathbf{g_f}$. This form of gradient surgery does not suffer from catastrophic forgetting, is robust to the choice of η, and consequently can achieve faster unlearning speeds compared to (SA) with larger values of η. For a comparison of these two versions of gradient surgery: (SA) and (S), see Appendix 7.

3.3 UNO and UNO-S

In the ideal scenario, when $\mathbf{g_f}$ is orthogonal to $\mathbf{g_r}$, (SA) is equivalent to gradient ascent (A) without the risk of losing desired knowledge. In this case, (S) is equivalent to retraining the model on the retain data, without the risk of relearning about the forget data. Therefore, we propose a modified loss function that attempts to enforce this ideal scenario with the help of an orthogonality promoting regularization term,

$$\mathcal{L}_{\text{UNO}} = \frac{1}{|\mathcal{D}_r|} \sum_{x \in \mathcal{D}_r} \mathcal{L}(\mathcal{M}_\theta, x) + \beta_o \left(\frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_r}\| \|\mathbf{g_f}\|} \right)^2, \tag{3}$$

where β_o is a regularization parameter. The unlearning via orthogonalization algorithm (UNO), can be expressed as performing gradient descent on this modified loss,

$$\theta_{k+1} = \theta_k - \eta \nabla_{\theta_k} \mathcal{L}_{\text{UNO}}. \tag{UNO}$$

Note that we only use the retain data to construct the first term in (3) to mimic the ideal retraining scenario mentioned above.

We further propose a hybrid algorithm that applies the (UNO) update step and the (S) update step alternately which we refer to as UNO-S:

$$\theta_{k+1} = \begin{cases} \theta_k - \eta \nabla_{\theta_k} \mathcal{L}_{\text{UNO}}, & \text{if } k \text{ is even}, \\ \theta_k - \eta \bar{\mathbf{g}}_\mathbf{r}, & \text{if } k \text{ is odd}. \end{cases} \tag{UNO-S}$$

The UNO update step attempts to enforce orthogonality between $\mathbf{g_f}$ and $\mathbf{g_r}$, which helps the subsequent surgery step effectively resolve the conflict between them.

3.4 Unlearning in the presence of a classifier able to distinguish between $\mathcal{D}_r$ and $\mathcal{D}_f$

If we have access to a binary classifier that distinguishes forget data from retain data, we can leverage this extra information to accelerate unlearning algorithms. We can use this classifier to identify every sample generated by our model as either a retain or forget sample, and compute the probability p_r that a generated sample is a retain sample. This associates our generative model with a Bernoulli distribution with probability of success p_r. We would like this distribution to have probability of success close to 1 or $1-\alpha$ where α is a small positive threshold controlling the degree of forgetting. We can enforce this by simply adding the following term to our loss,

$$\beta_h d_{\mathrm{KL}} = \beta_h \left[p_r \log\left(\frac{p_r}{1-\alpha}\right) + (1-p_r)\log\left(\frac{1-p_r}{\alpha}\right)\right], \tag{4}$$

where β_h is a regularization parameter, and d_{KL} represents the KL divergence between the computed and desired Bernoulli distributions. Small positive values of α ensure stable computation of this KL divergence term. Recalling that p_r is a function of the model and its parameters, we can now use the modified loss function in place of the original loss in the previously described algorithms. We use the hat symbol ($\hat{}$) to denote unlearning algorithms that operate with the additional loss term (4). For example, gradient surgery (S), UNO, and UNO-S become $\hat{\mathrm{S}}$, UN$\hat{\mathrm{O}}$, and UN$\hat{\mathrm{O}}$-$\hat{\mathrm{S}}$, respectively, when (4) is utilized. Addition of the new term yields the following modified definitions of $\mathbf{g_f}$ and $\mathbf{g_r}$:

$$\mathbf{g_f} = \frac{1}{|\mathcal{D}_f|}\sum_{x\in\mathcal{D}_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h \nabla_\theta d_{\mathrm{KL}}, \tag{5}$$

$$\mathbf{g_r} = \frac{1}{|\mathcal{D}_r|}\sum_{x\in\mathcal{D}_r} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h \nabla_\theta d_{\mathrm{KL}}. \tag{6}$$

Using (5), (6) with (S) gives us $\hat{\mathrm{S}}$. Similarly, the update rule for UN$\hat{\mathrm{O}}$ can be written as,

$$\begin{aligned} \mathcal{L}_{\mathrm{UN\hat{O}}} &= \frac{1}{|\mathcal{D}_r|}\sum_{x\in\mathcal{D}_r} \mathcal{L}(\mathcal{M}_\theta, x) + \beta_o \left(\frac{\mathbf{g_r}\cdot\mathbf{g_f}}{\|\mathbf{g_r}\|\|\mathbf{g_f}\|}\right)^2 + \beta_h d_{\mathrm{KL}}, \\ \theta_{k+1} &= \theta_k - \eta \nabla_{\theta_k} \mathcal{L}_{\mathrm{UN\hat{O}}}. \end{aligned} \tag{UN$\hat{\mathrm{O}}$}$$

Alternating update steps of UN$\hat{\mathrm{O}}$ and $\hat{\mathrm{S}}$ gives us UN$\hat{\mathrm{O}}$-$\hat{\mathrm{S}}$. Since the KL divergence term promotes unlearning of the forget data by preventing generation of forget samples, we also test the following update rule which is equivalent to UN$\hat{\mathrm{O}}$ with $\beta_o = 0$,

$$\begin{aligned} \mathcal{L}_H &= \frac{1}{|\mathcal{D}_r|}\sum_{x\in\mathcal{D}_r} \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h d_{\mathrm{KL}}, \\ \theta_{k+1} &= \theta_k - \eta \nabla_{\theta_k} \mathcal{L}_H. \end{aligned} \tag{H}$$

We call the resulting unlearning algorithm histogram unlearning and denote it by H throughout the remainder of this work.

4 Results

We test the algorithms described in Section 3 and Appendix 6 on VAEs trained on MNIST [11] and CelebA [27]. Each algorithm was tested 10 times to generate statistics. For the training losses used to train the original VAEs, training data, experiment hyperparameters, and model sizes, refer to Appendix 9. The architecture of the models can be found in the code provided in Section 1. All experiments were done on an A100 GPU provided by Google Colab.

4.1 Performance metrics

In order to assess the speed of unlearning we use classifiers trained on the datasets and track the fraction of generated samples that are classified as forget samples after each model update or training step. Our classifiers achieve $\sim 98\%$ accuracy on unseen data. Therefore, we define the **time to**

unlearn as the execution time of the unlearning algorithm until the fraction of forget samples drops below 0.02. We evaluate the quality of the generated images by computing the **Fréchet Inception Distance (FID)** using $25,000$ samples from the model and an equal number of real images from the dataset. We also report the execution **time per training step**, however, we do not highlight these values in the tables, as a larger time per step does not necessarily indicate slower unlearning, and vice versa. For the algorithms in Section 3.4, whenever applicable, we report the **speed-up** in Table 2 which is simply the factor by which the time to unlearn decreases due to utilizing the extra information provided by the classifier.

4.2 MNIST

We use a 0.6M-parameter VAE with a 2-dimensional latent space, trained for 200 epochs on $60,000$ images, as our original model. We attempt to unlearn the digit "1" by running the algorithms for 530 training steps with a mini-batch size of 128, and a learning rate of 10^{-3}. Figure 1 shows samples generated before and after unlearning with UNO, using the same noise samples for ease of comparison. The 1's in the original generation (left) transform into $7, 8$ and 3 after unlearning (right). The non-1 digits remain nearly unchanged. Even though 1's can transform into many different digits, they have an affinity for turning into 8's, followed by 3's, as seen in Figure 2. If the goal is uniform generation across the retain classes, one may utilize a loss term similar to d_{KL} for enforcing uniformity, assuming the availability of a classifier for all classes.

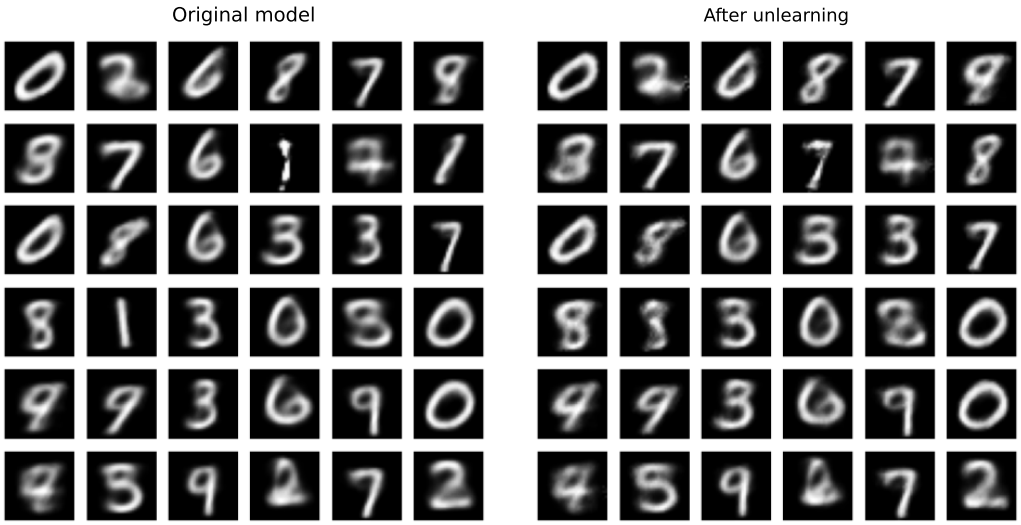

Figure 1: MNIST samples generated by the original model (left) and after unlearning digit "1" with UNO (right), using identical noise inputs for the decoder.

Figure 3 and Table 1 show that UNO-S achieves the fastest unlearning time, closely followed by UNO, with both having similar fidelity as the original model, indicated by the FID. Gradient ascent, while fast at unlearning, suffers from catastrophic forgetting, resulting in a large FID. Ascent-descent also experiences catastrophic forgetting and is significantly slower at unlearning than gradient ascent. Gradient surgery, while preserving image quality, is ~ 170 times slower than UNO and UNO-S at unlearning. Even though UNO takes ~ 1.6 times longer to execute a training step compared to gradient surgery, it still achieves orders of magnitude faster unlearning speed. Since one step of surgery is faster than one step of UNO, UNO-S overall is slightly faster than UNO, as the time per training step is averaged over the two algorithms. UNO-S, while taking twice as long as gradient ascent for a single training step, achieves the same unlearning speed.

4.3 CelebA

We use an 8.7M-parameter VAE with a 512-dimensional latent space, trained for 200 epochs on $202,599$ 64×64 images, downsampled from the original 178×178 resolution, as our original

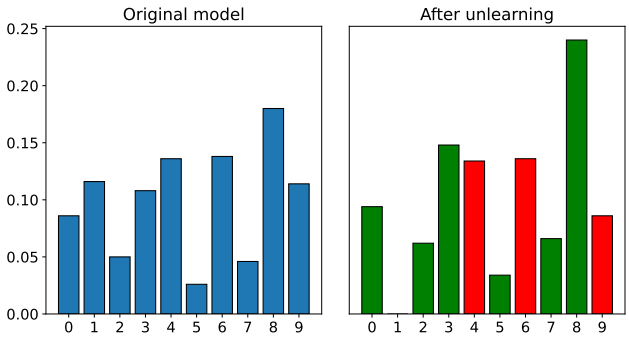

Figure 2: Distribution of generated digits before (left) and after unlearning (right), for a single run of UNO. Each histogram shows data for 500 generated samples. A bar in the right panel is colored green if the fraction of the corresponding digit increases after unlearning, and red if it decreases.

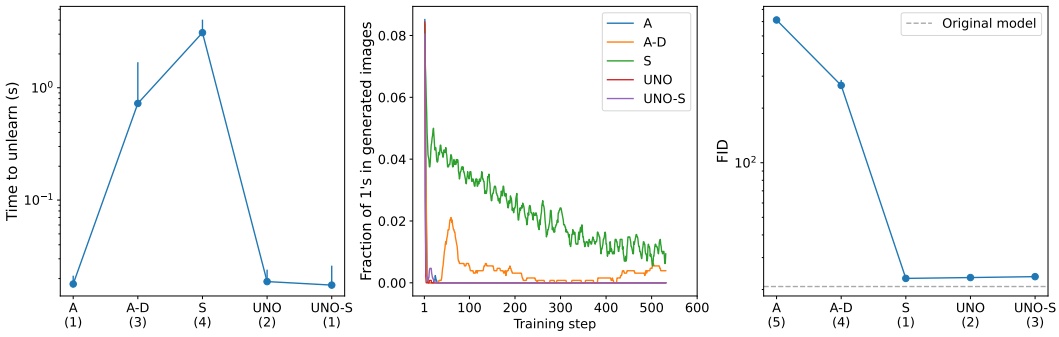

Figure 3: Time to unlearn (left) and FID (right) for various unlearning algorithms applied on MNIST. The middle panel shows the fraction of generated images classified as "1", averaged over 10 runs, as a function of training steps. Standard deviations over 10 independent runs are shown as one-sided error bars in the left and right panels. For time to unlearn and FID, the rank of each algorithm is also indicated below its name, with 1 being the best.

model. We attempt to unlearn "male" faces by running the algorithms for 659 training steps with a mini-batch size of 128, and a learning rate of 10^{-3}. Approximately 29% of the faces generated by the original model are male. Figure 4 shows samples generated before and after unlearning with UNO, using the same noise samples. We observe that male faces are successfully converted into female faces, and that feminine features are enhanced after unlearning, even when the originally generated face was already female. The original image remains nearly unchanged if it contains few or no male-specific features; see, for example, the last pair from the left in Figure 4. One notable effect of unlearning male-specific features is that the transformed images exhibit broader smiles. This is due to the sociological phenomenon wherein women tend to smile more in photographs [42]. For examples of these effects, see a larger collection of before/after unlearning pairs in Appendix 8.

Figure 5 and Table 1 show that UNO-S again achieves the fastest time to unlearn, followed by UNO. Even after spending ~ 200 times more execution time than the time to unlearn with UNO, gradient surgery is unable to achieve the desired $\leq 2\%$ male faces in the generated images. After the 659 allotted training steps gradient surgery is only able to reach $\sim 4\%$ male faces (see Figure 5, middle panel). All three algorithms result in similar values of FID, and the quality of the generated images is perceptually indistinguishable from the originally generated images, as seen in Figure 4.

4.4 Unlearning with a classifier

Table 2 shows that $\hat{\text{S}}$ achieves orders of magnitude speed-up over S for both MNIST and CelebA. UNO and UNO-S, already fast, gain an additional 10% and 20% speed-up respectively after incorporating the new information provided by the classifier. Histogram unlearning (H), although successful, is

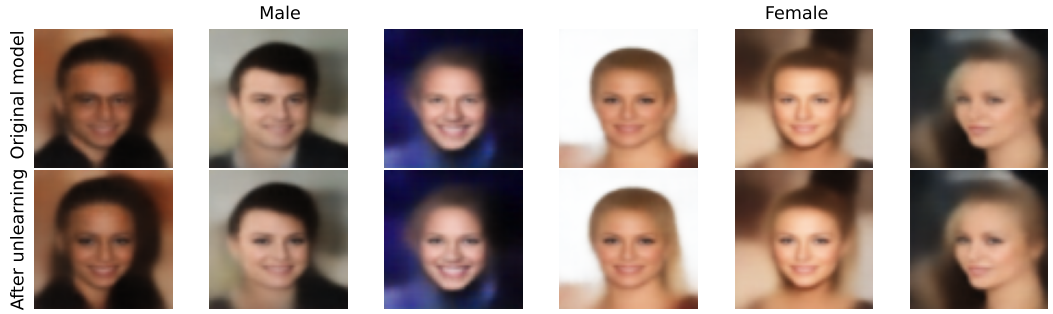

Figure 4: CelebA samples generated by the original model (top) and after unlearning "male" faces with UNO (bottom), using identical noise inputs for the decoder.

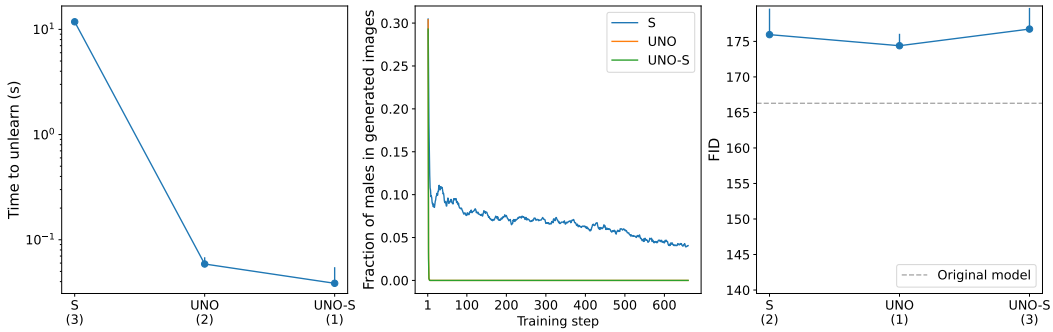

Figure 5: Time to unlearn (left) and FID (right) for various unlearning algorithms applied on CelebA. The middle panel shows the fraction of generated images classified as "male", averaged over 10 runs, as a function of training steps. Standard deviations over 10 independent runs are shown as one-sided error bars in the left and right panels. For time to unlearn and FID, the rank of each algorithm is also indicated below its name, with 1 being the best.

Table 1: Performance of various algorithms for class/feature unlearning with VAE on MNIST and CelebA. Each experiment is repeated 10 times, and the standard deviations are shown in parentheses. Bold indicates the best score. ✗ indicates that the generated samples after unlearning are unrecognizably different from the original model. ✓ indicates the generated samples after unlearning are perceptually indistinguishable from the original model in terms of visual fidelity. '*' indicates the algorithm was unable to achieve the desired fraction of forget samples in the generated images after the allotted number of training steps.

Dataset	**Algorithm**	Time to unlearn (s) ↓	FID ↓	Time per step (s)
MNIST (Class: 1) Original FID: 20.7	Gradient ascent (A)	**0.018** (0.003)	612.3 (4.9) ✗	0.005 (0.0001)
	Ascent descent (A-D)	0.725 (0.963)	266.9 (19.3) ✗	0.005 (0.0002)
	Gradient surgery (S)	3.094 (0.945)	**23.0** (0.2) ✓	0.007 (0.0007)
	UNO	0.019 (0.005)	23.3 (0.5) ✓	0.011 (0.0007)
	UNO-S	**0.018** (0.009)	23.6 (0.5) ✓	0.010 (0.0023)
CelebA (Feature: Male) Original FID: 166.3	Gradient surgery (S)	11.81* (0.688)	176.0 (3.7) ✓	0.018 (0.0003)
	UNO	0.059 (0.009)	**174.4** (1.7) ✓	0.031 (0.0006)
	UNO-S	**0.039** (0.016)	176.7 (3.0) ✓	0.024 (0.0067)

20–130 times slower than the other algorithms in Table 2. All algorithms in Table 2 preserve the fidelity of the original model, with UN$\hat{\text{O}}$ producing the lowest FID for both datasets.

Table 2: Performance of various algorithms for class/feature unlearning with VAE on MNIST and CelebA when a classifier able to distinguish between the retain and forget data is available. Each experiment is repeated 10 times, and the standard deviations are shown in parentheses. Bold indicates the best score. ✗ indicates that the generated samples after unlearning are unrecognizably different from the original model. ✓ indicates the generated samples after unlearning are perceptually indistinguishable from the original model in terms of visual fidelity. '*' indicates in the absence of the classifier, the algorithm did not reach the desired fraction of forget samples in the generated images after the allotted number of training steps.

Dataset	**Algorithm**	Time to unlearn (s) ↓	FID ↓	Time per step (s)	Speed-up ↑
MNIST (Class: 1) Original FID: 20.7	H	2.181 (0.853)	**23.4** (0.5) ✓	0.006 (0.0003)	-
	$\hat{\text{S}}$	**0.014** (0.003)	24.1 (0.5) ✓	0.008 (0.0008)	**221.0**
	UN$\hat{\text{O}}$	0.018 (0.006)	**23.4** (0.3) ✓	0.011 (0.0011)	1.1
	UN$\hat{\text{O}}$-$\hat{\text{S}}$	0.015 (0.006)	23.6 (0.5) ✓	0.009 (0.0022)	1.2
CelebA (Feature: Male) Original FID: 166.3	H	6.049 (2.155)	174.8 (1.9) ✓	0.016 (0.0002)	-
	$\hat{\text{S}}$	0.292 (0.247)	176.4 (3.5) ✓	0.023 (0.0003)	**40.4***
	UN$\hat{\text{O}}$	0.053 (0.009)	**174.4** (1.7) ✓	0.031 (0.0006)	1.1
	UN$\hat{\text{O}}$-$\hat{\text{S}}$	**0.046** (0.015)	175.4 (3.3) ✓	0.026 (0.0055)	1.2

5 Discussion

We advance the gradient surgery paradigm for machine unlearning by introducing two new algorithms UNO and UNO-S. We show that they are as fast as gradient ascent at unlearning without suffering from catastrophic forgetting, and substantially faster than gradient surgery. In the absence of a classifier, UNO-S outperforms all other algorithms, and can be up to 1.5 times faster than UNO at unlearning. The proposed algorithms preserve the quality of generation and the influence of the desired training data on the knowledge acquired by the model. We also demonstrate how incorporating the information provided by a classifier able to distinguish between desirable and undesired data, can accelerate unlearning algorithms. Table 3 in Appendix 9 documents the hyperparameter used in our experiments. The identical hyperparameter values across different scenarios in the table indicate that our algorithms are robust to the selection of hyperparameters.

Limitations

Although we test 9 different unlearning algorithms across 2 datasets using a variety of models, and report statistics over 10 independent runs, our study has the following limitations. The algorithms were not tested on model sizes currently considered large. Since these algorithms are general-purpose, it would also be interesting to apply them to other architectures, such as GANs [14] and transformers [40].

Future work

It is straightforward to conceptualize low-rank adapted [19, 43] variants of the unlearning algorithms presented here. Such modifications are essential for enabling efficient unlearning in large-scale generative models, and we leave their exploration to future research. The CelebA experiments show that, unlearning can easily produce male-to-female face filters. Applications of unlearning for designing a broader range of filters is an interesting topic for further exploration. Machine learning models used to simulate or predict physical systems, such as climate models, often generate unphysical states [26]. A similar issue arises in video generation models like Sora [21], which can produce physically implausible outputs. Use of unlearning to prevent generation of such unphysical outputs can be an extremely impactful research direction.

Acknowledgements

The authors acknowledge support from the Australian Research Council under Grant No. DP220100931.

References

[1] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). `https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng`, 2016. OJ L 119, 4.5.2016, pp. 1–88.

[2] Seohui Bae, Seoyoon Kim, Hyemin Jung, and Woohyung Lim. Gradient surgery for one-shot unlearning on generative model. *arXiv preprint arXiv:2307.04550*, 2023.

[3] Samyadeep Basu, Philip Pope, and Soheil Feizi. Influence functions in deep learning are fragile. *arXiv preprint arXiv:2006.14651*, 2020.

[4] Alberto Blanco-Justicia, Najeeb Jebreel, Benet Manzanares-Salor, David Sánchez, Josep Domingo-Ferrer, Guillem Collell, and Kuan Eeik Tan. Digital forgetting in large language models: A survey of unlearning methods. *Artificial Intelligence Review*, 58(3):90, 2025.

[5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.

[6] Zhiqi Bu, Xiaomeng Jin, Bhanukiran Vinzamuri, Anil Ramakrishna, Kai-Wei Chang, Volkan Cevher, and Mingyi Hong. Unlearning as multi-task optimization: A normalized gradient difference approach with an adaptive learning rate. *arXiv preprint arXiv:2410.22086*, 2024.

[7] Zihao Cao, Jianzong Wang, Shijing Si, Zhangcheng Huang, and Jing Xiao. Machine unlearning method based on projection residual. In *2022 IEEE 9th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–8. IEEE, 2022.

[8] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX security symposium (USENIX Security 21)*, pages 2633–2650, 2021.

[9] Jiaao Chen and Diyi Yang. Unlearn what you want to forget: Efficient unlearning for llms. *arXiv preprint arXiv:2310.20150*, 2023.

[10] Michael Crawshaw. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*, 2020.

[11] Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.

[12] Aditya Golatkar, Alessandro Achille, Avinash Ravichandran, Marzia Polito, and Stefano Soatto. Mixed-privacy forgetting in deep networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 792–801, 2021.

[13] Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9304–9312, 2020.

[14] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[15] Han Guo, Nazneen Fatema Rajani, Peter Hase, Mohit Bansal, and Caiming Xiong. Fastif: Scalable influence functions for efficient model interpretation and debugging. *arXiv preprint arXiv:2012.15781*, 2020.

[16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 770–778, 2016.

[17] Yuan Heng and Yew-Soon Soh. Selective amnesia: Learning to forget in generative models. *arXiv preprint arXiv:2301.13580*, 2023.

[18] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Eliza Noland, Kate Millican, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

[19] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. arxiv 2021. *arXiv preprint arXiv:2106.09685*, 2021.

[20] Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models. *arXiv preprint arXiv:2210.01504*, 2022.

[21] Bingyi Kang, Yang Yue, Rui Lu, Zhijie Lin, Yang Zhao, Kaixin Wang, Gao Huang, and Jiashi Feng. How far is video generation from world model: A physical law perspective. *arXiv preprint arXiv:2411.02385*, 2024.

[22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[23] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR, 2017.

[24] Christopher Kuner, Lee A. Bygrave, and Christopher Docksey, editors. *The EU General Data Protection Regulation (GDPR): A Commentary*. Oxford University Press, Oxford, UK, 2020.

[25] Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded machine unlearning. *Advances in neural information processing systems*, 36:1957–1987, 2023.

[26] Ching-Yao Lai, Pedram Hassanzadeh, Aditi Sheshadri, Maike Sonnewald, Raffaele Ferrari, and Venkatramani Balaji. Machine learning for climate physics and simulations. *Annual Review of Condensed Matter Physics*, 16, 2024.

[27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3730–3738, 2015.

[28] Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *arXiv preprint arXiv:2308.08747*, 2023.

[29] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, 24:109–165, 1989.

[30] Fatemehsadat Mireshghallah, Huseyin A Inan, Marcello Hasegawa, Victor Rühle, Taylor Berg-Kirkpatrick, and Robert Sim. Privacy regularization: Joint privacy-utility optimization in language models. *arXiv preprint arXiv:2103.07567*, 2021.

[31] Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, pages 931–962. PMLR, 2021.

[32] Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.

[33] Youyang Qu, Ming Ding, Nan Sun, Kanchana Thilakarathna, Tianqing Zhu, and Dusit Niyato. The frontier of data erasure: A survey on machine unlearning for large language models. *Computer*, 58(1):45–57, 2025.

[34] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1278–1286, 2014.

[35] Andrea Schioppa, Polina Zablotskaia, David Vilar, and Artem Sokolov. Scaling up influence functions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8179–8186, 2022.

[36] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[37] Hui Sun, Tianqing Zhu, Wenhan Chang, and Wanlei Zhou. Generative adversarial networks unlearning. *IEEE Transactions on Dependable and Secure Computing*, 2025.

[38] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.

[39] Nguyen Truong, Kai Sun, Siyao Wang, Florian Guitton, and YiKe Guo. Privacy preservation in federated learning: An insightful survey from the gdpr perspective. *Computers & Security*, 110:102402, 2021.

[40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[41] Weiqi Wang, Zhiyi Tian, Chenhan Zhang, and Shui Yu. Machine unlearning: A comprehensive survey. *arXiv preprint arXiv:2405.07406*, 2024.

[42] Taylor R Wondergem and Mihaela Friedlmeier. Gender and ethnic differences in smiling: A yearbook photographs analysis from kindergarten through 12th grade. *Sex Roles*, 67:403–411, 2012.

[43] Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen, Xiaopeng Zhang, and Qi Tian. Qa-lora: Quantization-aware low-rank adaptation of large language models. *arXiv preprint arXiv:2309.14717*, 2023.

[44] Jin Yao, Eli Chien, Minxin Du, Xinyao Niu, Tianhao Wang, Zezhou Cheng, and Xiang Yue. Machine unlearning of pre-trained large language models. *arXiv preprint arXiv:2402.15159*, 2024.

[45] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in neural information processing systems*, 33:5824–5836, 2020.

[46] Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. *arXiv preprint arXiv:2404.05868*, 2024.

Appendix

6 Pseudocode for unlearning algorithms

This section presents the pseudocode for the unlearning algorithms used in this work.

6.1 Gradient ascent

Algorithms 1 and 2 describe the gradient ascent (A), and alternating gradient ascent-descent (A-D), respectively.

Algorithm 1 Gradient ascent (A)

Input: Loss function $\mathcal{L}$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, learning rate η, number of training steps K, batch size B.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
    Acquire mini-batch D_f of size B from $\mathcal{D}_f$.
    $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
    $\theta \leftarrow \theta + \eta \mathbf{g_f}$
end for
return θ

Algorithm 2 Alternating gradient ascent and descent (A-D)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, learning rate η, number of training steps K, batch size B.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
    Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
    if k is odd **then**
        $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
        $\theta \leftarrow \theta + \eta \mathbf{g_f}$
    else
        $\mathbf{g_r} \leftarrow \frac{1}{B} \sum_{x \in D_r} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
        $\theta \leftarrow \theta - \eta \mathbf{g_r}$
    end if
end for
return θ

6.2 Gradient surgery

Algorithms 3 and 4 describe gradient surgery with ascent in the forget direction (SA) and descent in the retain direction (S), respectively; in particular, the former appears in [2].

6.3 UNO and UNO-S

Algorithms 5 and 6 describe unlearning via orthogonalization (UNO), and alternating orthogonalization and surgery (UNO-S), respectively.

6.4 Unlearning in the presence of a classifier able to distinguish between $\mathcal{D}_r$ and $\mathcal{D}_f$

Algorithms 7, 8, 9, and 10 describe $\hat{\text{S}}$, UN$\hat{\text{O}}$, UN$\hat{\text{O}}$-$\hat{\text{S}}$, and histogram unlearning, respectively. While in practice it is sufficient for a binary classifier to output a logit or probability, for simplicity of presentation we assume the classifier outputs 1 for retain samples and 0 otherwise.

Algorithm 3 Gradient surgery with ascent in forget direction (SA)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, learning rate η, number of training steps K, batch size B.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
    Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
    $\mathbf{g_r} \leftarrow \frac{1}{B} \sum_{x \in D_r} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
    $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
    $\mathbf{g_f} \leftarrow \mathbf{g_f} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_r}\|^2} \mathbf{g_r}$
    $\theta \leftarrow \theta + \eta \mathbf{g_f}$
end for
return θ

Algorithm 4 Gradient surgery with descent in retain direction (S)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, learning rate η, number of training steps K, batch size B.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
    Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
    $\mathbf{g_r} \leftarrow \frac{1}{B} \sum_{x \in D_r} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
    $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
    $\mathbf{g_r} \leftarrow \mathbf{g_r} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_f}\|^2} \mathbf{g_f}$
    $\theta \leftarrow \theta - \eta \mathbf{g_r}$
end for
return θ

Algorithm 5 Unlearning via orthogonalization (UNO)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, weight for orthogonalization loss term β_o, learning rate η, number of training steps K, batch size B.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
    Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
    $L_r \leftarrow \frac{1}{B} \sum_{x \in D_r} \mathcal{L}(\mathcal{M}_\theta, x)$
    $\mathbf{g_r} \leftarrow \nabla_\theta L_r$
    $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
    $L \leftarrow L_r + \beta_o \left(\frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_r}\| \|\mathbf{g_f}\|} \right)^2$
    $\theta \leftarrow \theta - \eta \nabla_\theta L$
end for
return θ

Algorithm 6 Alternating orthogonalization and surgery (UNO-S)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, weight for orthogonalization loss term β_o, learning rate η, number of training steps K, batch size B.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
  Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
  $L_r \leftarrow \frac{1}{B} \sum_{x \in D_r} \mathcal{L}(\mathcal{M}_\theta, x)$
  $\mathbf{g_r} \leftarrow \nabla_\theta L_r$
  $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x)$
  if k is odd **then**
    $L \leftarrow L_r + \beta_o \left(\frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_r}\| \|\mathbf{g_f}\|} \right)^2$
    $\theta \leftarrow \theta - \eta \nabla_\theta L$
  else
    $\mathbf{g_r} \leftarrow \mathbf{g_r} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_f}\|^2} \mathbf{g_f}$
    $\theta \leftarrow \theta - \eta \mathbf{g_r}$
  end if
end for
return θ

Algorithm 7 Gradient surgery with histogram unlearning ($\hat{\text{S}}$)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, learning rate η, number of training steps K, batch size B, number of samples to generate N_g, classifier model $\mathcal{C}_\phi$, weight for KL divergence loss term β_h, a small positive threshold for stabilizing KL divergence computation α.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
  Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
  Generate N_g samples $\{y_i\}_{i=1}^{N_g}$ using $\mathcal{M}_\theta$.
  $p_r \leftarrow \frac{1}{N_g} \sum_{i=1}^{N_g} \mathcal{C}_\phi(y_i)$
  $d_{\text{KL}} \leftarrow p_r \log \left(\frac{p_r}{1-\alpha} \right) + (1 - p_r) \log \left(\frac{1-p_r}{\alpha} \right)$
  $\mathbf{g_r} \leftarrow \frac{1}{B} \sum_{x \in D_r} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h \nabla_\theta d_{\text{KL}}$
  $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h \nabla_\theta d_{\text{KL}}$
  $\mathbf{g_r} \leftarrow \mathbf{g_r} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_f}\|^2} \mathbf{g_f}$
  $\theta \leftarrow \theta - \eta \mathbf{g_r}$
end for
return θ

7 Comparison of two variants of gradient surgery

We now compare two variants of gradient surgery: 1) gradient surgery with descent in retain direction (S), described in Algorithm 4, used throughout this paper and 2) gradient surgery with ascent in forget direction (SA), described in Algorithm 3 which appears in [2]. Figure 6 shows that SA is prone to catastrophic forgetting and requires a carefully tuned, small learning rate to mitigate this effect. But even with a small learning rate the generated samples might look significantly different from the original model; for samples generated by the original model, see Figure 1. On the other hand, Figure 7 shows that S does not suffer from catastrophic forgetting, even for a large learning rate applied for many training steps, and produces samples that are much closer to the original model.

Algorithm 8 UNO with histogram unlearning (UNÔ)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, weight for orthogonalization loss term β_o, learning rate η, number of training steps K, batch size B, number of samples to generate N_g, classifier model $\mathcal{C}_\phi$, weight for KL divergence loss term β_h, a small positive threshold for stabilizing KL divergence computation α.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
  Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
  Generate N_g samples $\{y_i\}_{i=1}^{N_g}$ using $\mathcal{M}_\theta$.
  $p_r \leftarrow \frac{1}{N_g} \sum_{i=1}^{N_g} \mathcal{C}_\phi(y_i)$
  $d_{\mathrm{KL}} \leftarrow p_r \log\left(\frac{p_r}{1-\alpha}\right) + (1-p_r)\log\left(\frac{1-p_r}{\alpha}\right)$
  $L_r \leftarrow \frac{1}{B} \sum_{x \in D_r} \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h d_{\mathrm{KL}}$
  $\mathbf{g_r} \leftarrow \nabla_\theta L_r$
  $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h \nabla_\theta d_{\mathrm{KL}}$
  $L \leftarrow L_r + \beta_o \left(\frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_r}\| \|\mathbf{g_f}\|}\right)^2$
  $\theta \leftarrow \theta - \eta \nabla_\theta L$
end for
return θ

Algorithm 9 Alternating orthogonalization and surgery with histogram unlearning (UNÔ-Ŝ)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, weight for orthogonalization loss term β_o, learning rate η, number of training steps K, batch size B, number of samples to generate N_g, classifier model $\mathcal{C}_\phi$, weight for KL divergence loss term β_h, a small positive threshold for stabilizing KL divergence computation α.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
  Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
  Generate N_g samples $\{y_i\}_{i=1}^{N_g}$ using $\mathcal{M}_\theta$.
  $p_r \leftarrow \frac{1}{N_g} \sum_{i=1}^{N_g} \mathcal{C}_\phi(y_i)$
  $d_{\mathrm{KL}} \leftarrow p_r \log\left(\frac{p_r}{1-\alpha}\right) + (1-p_r)\log\left(\frac{1-p_r}{\alpha}\right)$
  $L_r \leftarrow \frac{1}{B} \sum_{x \in D_r} \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h d_{\mathrm{KL}}$
  $\mathbf{g_r} \leftarrow \nabla_\theta L_r$
  $\mathbf{g_f} \leftarrow \frac{1}{B} \sum_{x \in D_f} \nabla_\theta \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h \nabla_\theta d_{\mathrm{KL}}$
  if k is odd **then**
    $L \leftarrow L_r + \beta_o \left(\frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_r}\| \|\mathbf{g_f}\|}\right)^2$
    $\theta \leftarrow \theta - \eta \nabla_\theta L$
  else
    $\mathbf{g_r} \leftarrow \mathbf{g_r} - \frac{\mathbf{g_r} \cdot \mathbf{g_f}}{\|\mathbf{g_f}\|^2} \mathbf{g_f}$
    $\theta \leftarrow \theta - \eta \mathbf{g_r}$
  end if
end for
return θ

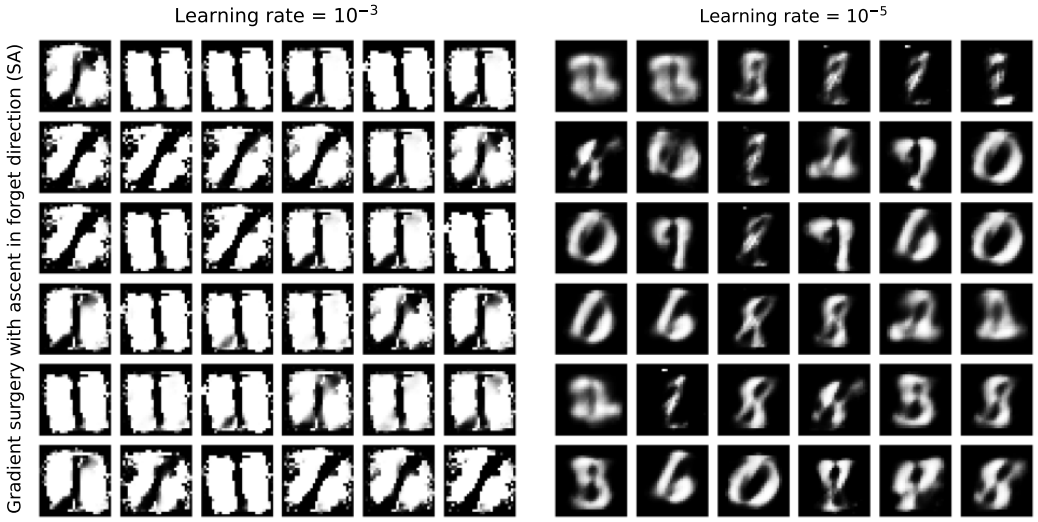

Figure 6: Generated samples after unlearning digit 1 via gradient surgery with ascent in forget direction (SA), described in Algorithm 3, for two different learning rates: 10^{-3} (left), 10^{-5} (right). SA was run for $K = 53$ training steps on the left and $K = 530$ training steps on the right.

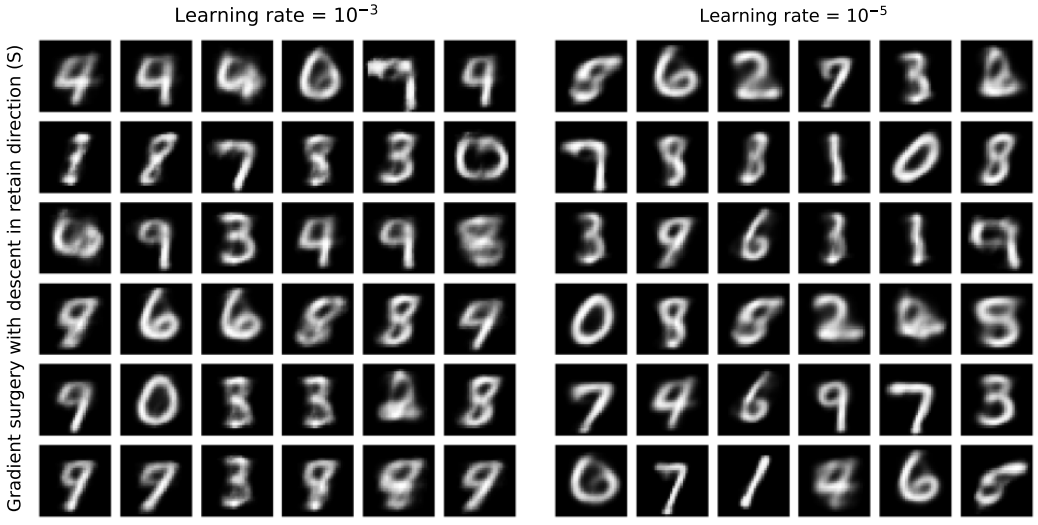

Figure 7: Generated samples after unlearning digit 1 via gradient surgery with descent in retain direction (S), described in Algorithm 4, for two different learning rates: 10^{-3} (left), 10^{-5} (right). S was run for $K = 530$ training steps for both learning rates.

Algorithm 10 Histogram unlearning (H)

Input: Loss function $\mathcal{L}$, retain dataset $\mathcal{D}_r$, forget dataset $\mathcal{D}_f$, trained model requiring unlearning $\mathcal{M}_\theta$, learning rate η, number of training steps K, batch size B, number of samples to generate N_g, classifier model $\mathcal{C}_\phi$, weight for KL divergence loss term β_h, a small positive threshold for stabilizing KL divergence computation α.
Output: Updated model parameters θ.
for $k = 1$ to K **do**
    Acquire retain and forget mini-batches D_r, D_f of size B from $\mathcal{D}_r, \mathcal{D}_f$ respectively.
    Generate N_g samples $\{y_i\}_{i=1}^{N_g}$ using $\mathcal{M}_\theta$.
    $p_r \leftarrow \frac{1}{N_g} \sum_{i=1}^{N_g} \mathcal{C}_\phi(y_i)$
    $d_{\mathrm{KL}} \leftarrow p_r \log\left(\frac{p_r}{1-\alpha}\right) + (1-p_r)\log\left(\frac{1-p_r}{\alpha}\right)$
    $L \leftarrow \frac{1}{B} \sum_{x \in D_r} \mathcal{L}(\mathcal{M}_\theta, x) + \beta_h d_{\mathrm{KL}}$
    $\theta \leftarrow \theta - \eta \nabla_\theta L$
end for
return θ

8 More generated samples for CelebA before and after unlearning

Figure 8 shows 18 pairs of images generated before and after unlearning with UNO for CelebA. In many of these pairs the after image shows a subtly larger smile than the before image, see for example the fourteenth pair.

9 Experiment setup

In this section we provide additional details of the experimental setup used to produce the reported results.

9.1 VAE loss functions and training data

We now document the loss functions used to train the original model. For each input image x, the encoder outputs $\mu(x) \in \mathbb{R}^{d_z}$ and $\sigma(x) \in \mathbb{R}^{d_z}$, which parameterize the approximate posterior distribution. Here d_z is the latent dimension. The corresponding reconstruction of x by the decoder is denoted by $\bar{x}$, with $\bar{x}_i$ referring to its i-th pixel.

The VAE used as the original model for MNIST was trained using the loss function

$$\mathcal{L}_{\mathrm{MNIST}} = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \Big[- \sum_{i=1}^{784} (x_i \log(\bar{x}_i) + (1-x_i)\log(1-\bar{x}_i)) + \frac{1}{2} \sum_{i=1}^{d_z} \left(\mu_i^2(x) + \sigma_i^2(x) - \log \sigma_i^2(x) - 1\right) \Big]. \tag{7}$$

The $60,000$ training images were normalized such that pixel values lie in $[0, 1]$, following standard practice.

The VAE used as the original model for CelebA was trained using the loss function

$$\mathcal{L}_{\mathrm{CelebA}} = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \left[\|x - \bar{x}\|^2 + \frac{1}{2} \sum_{i=1}^{d_z} \left(\mu_i^2(x) + \sigma_i^2(x) - \log \sigma_i^2(x) - 1\right) \right]. \tag{8}$$

We worked with $202,599$ cropped and aligned images in CelebA which originally have resolution 178×178 pixels. We downsampled these images to 64×64 resolution for training.

9.2 Hyperparameters

Table 3 lists the hyperparameters used in the unlearning experiments presented here. Here η is the learning rate, K is the number of training steps executed, β_o is the weight for the orthogonalization

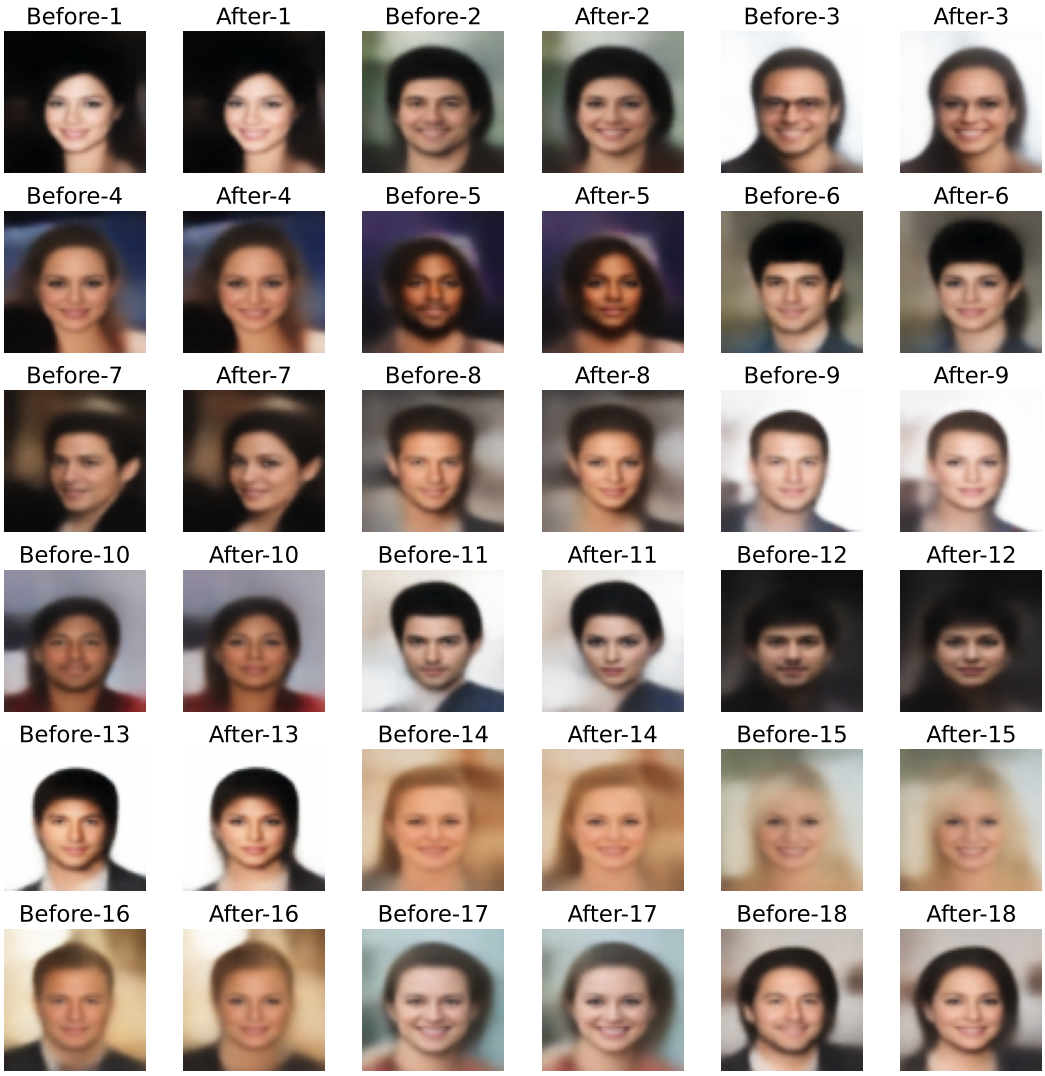

Figure 8: Results for unlearning on CelebA with UNO, illustrated using 18 pairs of generated images. The images labeled "Before" were generated using the original model. Each image labeled "After" was generated after unlearning using the same noise sample as the corresponding "Before" image.

loss term in (3), β_h is the weight for the KL divergence loss term in (4), α is a small positive threshold for stable computation of KL divergence in (4), B is the batch size, and N_{FID} is the number of samples used for calculating FID. We use $N_g = B$ for all the algorithms in Appendix 6.4, which determines the number of samples to be generated using the generative model. Each method was tested 10 times for each dataset. For MNIST, FID was computed using features extracted from the classifier model, whereas for CelebA, features were computed using the InceptionV3 model [38]. All experiments were done on an A100 GPU provided by Google Colab.

Table 3: Experiment hyperparameters

Dataset	**Algorithm**	η	K	$B\beta_o$	$B\beta_h$	α	B	N_{FID}
MNIST (Class: 1)	Gradient ascent (A)	10^{-3}	530	-	-	-	128	25,000
	Ascent descent (A-D)	10^{-3}	530	-	-	-	128	25,000
	Gradient surgery (S)	10^{-3}	530	-	-	-	128	25,000
	UNO	10^{-3}	530	10^3	-	-	128	25,000
	UNO-S	10^{-3}	530	10^3	-	-	128	25,000
	H	10^{-3}	530	-	10^3	10^{-8}	128	25,000
	$\hat{\text{S}}$	10^{-3}	530	-	10^3	10^{-8}	128	25,000
	UN$\hat{\text{O}}$	10^{-3}	530	10^3	10^3	10^{-8}	128	25,000
	UN$\hat{\text{O}}$-$\hat{\text{S}}$	10^{-3}	530	10^3	10^3	10^{-8}	128	25,000
CelebA (Feature: Male)	Gradient surgery (S)	10^{-3}	659	-	-	-	128	25,000
	UNO	10^{-3}	659	10^3	-	-	128	25,000
	UNO-S	10^{-3}	659	10^3	-	-	128	25,000
	H	10^{-3}	659	-	10^3	10^{-8}	128	25,000
	$\hat{\text{S}}$	10^{-3}	659	-	10^3	10^{-8}	128	25,000
	UN$\hat{\text{O}}$	10^{-3}	659	10^3	10^3	10^{-8}	128	25,000
	UN$\hat{\text{O}}$-$\hat{\text{S}}$	10^{-3}	659	10^3	10^3	10^{-8}	128	25,000

9.3 Model sizes

The VAE models for MNIST and CelebA have 632,788 and 8,742,659 parameters with latent dimension $d_z = 2$ and $d_z = 512$, respectively. The classifier models for MNIST and CelebA have 159,410 and 2,190,913 parameters, respectively. For the exact model implementations, please refer to the code linked in Section 1. The VAEs were trained for 200 epochs on $60,000$ and $202,599$ images in MNIST and CelebA, respectively. The classifiers were trained for 10 epochs on the same data.