

Moderating the Generalization of Score-based Generative Model

Wan Jiang¹ He Wang² Xin Zhang³ Dan Guo¹ Zhaoxin Fan^{4,5} Yunfeng Diao^{1,6,7*} Richang Hong¹

¹Hefei University of Technology ²University College London ³San Diego State University

⁴Hangzhou International Innovation Institute, Beihang University

⁵Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing, AI, BU

⁶ Intelligent Interconnected Systems Laboratory of Anhui Province (HFUT)

⁷Institute of Artifical Intelligence, Hefei Comprehensive National Science Center

xjiangw000@gmail.com, he_wang@ucl.ac.uk, xzhang17@wpi.edu, guodan@hfut.edu.cn

zhaoxinf@buaa.edu.cn, diaoyunfeng@hfut.edu.cn, hongrc.hfut@gmail

Abstract

Score-based Generative Models (SGMs) have demonstrated remarkable generalization capabilities, e.g. generating unseen, but natural data. However, the greater the generalization power, the more likely the unintended generalization, and the more dangerous the abuse. Despite these concerns, research on unlearning SGMs has not been explored. To fill this gap, we first examine the current 'gold standard' in Machine Unlearning (MU), i.e., re-training the model after removing the undesirable training data, and find it does not work in SGMs. Further analysis of score functions reveals that the MU 'gold standard' does not alter the original score function, which explains its ineffectiveness. Building on this insight, we propose the first Moderated Score-based Generative Model (MSGM), which introduces a novel score adjustment strategy that redirects the score function away from undesirable data during the continuous-time stochastic differential equation process. Albeit designed for SGMs, MSGM is a general and flexible MU framework compatible with diverse diffusion architectures, training strategies and downstream tasks. Code is available at https://github.com/yunfengdiao/ Moderated-Score-based-Generative-Model.

1. Introduction

"The greater the power, the more dangerous the abuse."

- Edmund Burke Generative models has been a foundational topic in deep learning in the past decade, *e.g.* Generative Adversarial Networks [13], Variational Autoencoders [20] and normalizing flows [7]. Recently Score-based Generative models

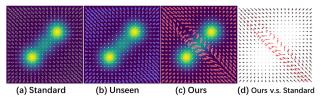


Figure 1. Comparsions of score functions of Standard VE SDE, Unseen by Re-training and MSGM (Ours) in the toy experiment. Unseen by Re-training retains almost the same score function with Standard VE SDE, while MSGM (red arrows in (d)) significantly alter the original score functions (black arrows in (d)).

(SGMs) [42], Denoising Diffusion Probabilistic Models (DDPMs) [17] and their variants become one of the dominating class of generative models. In generative models, one key research effort is to maximize their generalization ability, often meaning generating unseen, but natural data, e.g. images with realistic faces or stories not written by humans. However, such novel data can also be unintended and potentially cause privacy breaches, copyright infringements, misinformation spreading, etc. We refer to this phenomenon as unintended generalization. Taking SGMs and DDPMs as an example, they can reconstruct training data which should not to be accessible [3], generate faces that are similar to a specific person without permission [33, 35, 36], and mimic content styles of artists unintentionally [10, 38]. Overall, the bigger the generalization power is, the more likely the unintended generalization is, and therefore the more detrimental the intended/unintended harm becomes. However, it is not straightforward to design a new Machine Unlearning (MU) paradigm for SGMs and cannot be achieved by adapting existing MU methods for other models. First, MUs for DDPMs primarily aim to reduce the evidence lower bound (ELBO) on the distribution of the forgotten data, but SGMs need to estimate the score function with a continuous noise schedule. This difference in the optimization goals prevents

^{*}Corresponding author

existing DDPM MU methods from being directly applied to SGMs. More generally, most MU methods are designed for conditional generation [8, 10, 14, 15, 23, 25, 45, 48], especially on text-to-image generation with conditional DDPM, tightly coupling the MU with specific conditions, *e.g.* refining cross-attention layers [10–12, 25, 29]. They cannot be applied to unconditioned generation, which is widely used for *e.g.* data augmentation. Designing MU for unconditional generation can potentially provide a universal solution to a broader range of generative models.

To address the challenges, we propose the first Moderated Score-based Generative Model (MSGM) for controlling the unintended generalization. MSGM aims to overcome the limitations of current 'gold standard' in score-based generative unlearning, by introducing a straightforward yet effective strategy to alter the score function. Our key idea is to deform the original score function so that it circumvents the Not Suitable For Generation (NSFG) data in sampling, while ensuring that it still approximates the Suitable for Generation (SFG) data score to maintain the generation quality. To this end, we present two variants of MSGM for general MU. The first is to explore the orthogonal complement space of the original score function with respect to the NSFG data, so that sampling will steer away from the high density area of the NSFG data, which works well when the distributions of NSFG and SFG are mildly separable. The second is to explore the negatively correlated score subspace, to target scenarios where the distributions of NSFG and SFG are very similar. Additionally, if NSGF needs to be removed from a pre-trained model, we propose to negate part of the original score function locally.

Although designed for SGM, MSGM is a general and flexible MU framework compatible with various diffusion models and downstream tasks. To demonstrate its versatility, we conduct extensive experiments on five datasets using SGM, DDPM, and the latent diffusion Model. The results indicate that MSGM consistently achieves strong unlearning performance in both unconditional and conditional generation. Furthermore, it enables zero-shot transfer of unlearning pre-trained models to downstream tasks, including image inpainting and reconstruction.

2. Related Work

Diffusion Models and Security&Privacy Issues. Diffusion models progressively add noise to training data and learn to reverse this process.DDPMs [17, 28] use discrete denoising steps, while SGMs [39, 42] generalize to continuous time stochastic differential equations (SDEs). Song et al.[40] improve the training and sampling process and achieve high-fidelity image generation. Meng et al. [26] utilize high-order scores to accelerate the mixing speed of synthetic data and natural image sampling. In addition to technical improvements, SGMs have been shown to be effec-

tive across various applications, including natural language processing[31], computational physics [19], video prediction [9], audio codecs [46], medical imaging [5], etc. These improvements in diffusion models maximize their generalization ability, but it also raise the concerns about the unintended generalization. Rando et al. [33] observed that malicious users may bypass the safety filters in open-source diffusion models to create disturbing content, e.g. violence and gore. Beyond this, diffusion models are susceptible to create misleading videos or images of individuals without permission, potentially damaging their reputation or spreading misinformation [35, 36]. Additionally, these models can mimic various art styles, potentially infringing on portrait and intellectual property rights [10, 38].

Machine Unlearning in Generative Model. MU enables models to selectively forget the undesirable content for privacy, security or adaptability purposes. Early research has explored numerous MU methods on supervised learning tasks [37], such as image classification [44]. However, Fan et al. [8] demonstrated that existing MU for image classification cannot be applied in image generation. This gap highlights the urgent need for effective MU techniques in generative models. Very recently, new MU schemes have been proposed for different types of generative models, including unlearning in VAEs [2, 27], GANs [21, 43] and DDPMs [45], especially for text-to-image conditional generation [8, 10, 14, 15, 23, 48]. However, SGM for MU is still largely missing. More importantly, contrary to the common belief that Unseen by Re-training is the 'gold standard' in MU [44], we empirically found it is ineffective in SGM. This motivates us to develop a new SGM MU that can control the unintended generalization and surpass the previous 'gold standard' in MU.

3. Methodology

3.1. Preliminaries

Machine Unlearning in Generative Model. Let $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N \in \mathbb{R}^D$ be the training data, following the distribution $\mathbf{x}_i \sim p_d$. Let $\mathcal{D}_f = \{\mathbf{x}_i^u\}_{i=1}^M \subseteq D$ denote the Not Suitable For Generation (NSFG) data following the distribution $p_f(\mathbf{x})$. The remaining data, $\mathcal{D}_g = \mathcal{D} \backslash \mathcal{D}_f = \{\mathbf{x}_i^g\}_{i=1}^{N-M} \sim p_g(\mathbf{x})$, represents the Suitable For Generation (SFG) data. Our goal is to enable the generative model to avoid generating NSFG samples while maintaining the quality of image generation for SFG data. We refer to such a generative model as an unlearning generative model. We use the symbol p to denote either a probability distribution or its probability density or mass function depending on the context.

Score-Based Generative Modeling with SDEs. The two main components of a score-based SDE generative model [42] are the *forward process* and the *reverse pro-*

cess. The forward process $\{\mathbf{x}(t) \in \mathbb{R}^d\}_{t=0}^T$ transforming data from the distribution $p_{data}(\mathbf{x})$ to a simple noise distribution with a continuous-time stochastic differential equation (SDE):

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{w}, t \in [0, T], \tag{1}$$

where $\mathbf{f}: \mathbb{R}^d \to \mathbb{R}^d$ is called the drift, $g \in \mathbb{R}$ is called the diffusion, and \mathbf{w} represents the standard Brownian Motion. Let $p_t(\mathbf{x})$ denote the density of $\mathbf{x}(t)$. At time t=0, the initial distribution of $\mathbf{x}(0)$ follows $p_0:=p_{data}$, while at time $t=T,\mathbf{x}(T)$ adheres to p_T which is normally an easy-to-sample prior distribution such as Gaussian.

Given samples from the prior, the reverse process converts them into data samples via a reverse-time SDE:

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_{\mathbf{x}}\log p_t(\mathbf{x})]dt + g(t)d\bar{\mathbf{w}}, \quad (2)$$

where $\bar{\mathbf{w}}$ is a Brownian motion, and dt represents an infinitesimal negative time step. Running the reverse process requires estimating the score function of the forward process, which is typically done by training a neural network with a score-matching objective:

$$\min_{\theta} \mathbb{E}_{t} \lambda(t) \left\{ \mathbb{E}_{\mathbf{x}(0)} \mathbb{E}_{\mathbf{x}(t) \mid \mathbf{x}(0)} \left[\| s_{\theta}(\mathbf{x}(t), t) - \nabla_{\mathbf{x}(t)} \log p_{0t}(\mathbf{x}(t) \mid \mathbf{x}(0)) \|_{2}^{2} \right] \right\},$$
(3)

where $\mathbf{x}(0) \sim p_0(\mathbf{x})$ and $\mathbf{x}(t) \sim p_{0t}(\mathbf{x}(t) \mid \mathbf{x}(0))$, $t \sim \mathcal{U}(0,T)$ is a uniform distribution over [0,T], $p_{0t}(\mathbf{x}(t) \mid \mathbf{x}(0))$ denotes the transition probability from $\mathbf{x}(0)$ to $\mathbf{x}(t)$, and $\lambda(t) \in \mathbb{R}_{>0}$ denotes a positive weighting function. Other than Eq. (3), other score matching objectives, such as sliced score matching [41] and finite-difference score matching [30] are also applicable.

3.2. Motivation-Unintended Generalization

Two mainstream mechanisms for MU are commonly employed: (1) *Erasing by Fine-tuning*, which removes learned NSFG features from a pre-trained generator, and (2) *Unseen by Re-training*, which re-trains the generator on filtered data. The latter is widely regarded as the 'gold standard' in MU [8, 44], as its superior performance over fine-tuning methods [47].

Unfortunately, the 'gold standard' does not work for SGMs. Specifically, we demonstrate that even after removing NSFG data, SGMs can still generate samples resembling the unlearned features. To formalize this phenomenon, we design a synthetic experiment (Fig. 2) using a mixture of 2D Gaussians:

$$p_{\text{data}} = \underbrace{\frac{4}{5}\mathcal{N}((-2,-2),I)}_{\mathcal{D}_g} + \underbrace{\frac{2}{5}\mathcal{N}((0,0),I)}_{\mathcal{D}_f} + \underbrace{\frac{4}{5}\mathcal{N}((2,2),I)}_{\mathcal{D}_g},$$

where \mathcal{D}_f represents the NSFG data and \mathcal{D}_g is the rest data. The data distribution is shown in Fig. 2 (a)&(c) Left. We train a Variance Exploding Stochastic Differential Equation (VE SDE) model [42], referred to as the standard VE SDE,

Table 1. The Negative log-likelihood (NLL) values of different methods with respect to the data from $p_{\rm data}$.

Test	Standard	Unseen	MSGM
\mathcal{D}_g	10.91	10.63	10.64
$\mathcal{D}_f^{"}$	10.73	11.59	39.01

which after training generates a data distribution shown in Fig. 2 (a) Right. Clearly, the standard VE SDE learns to generate all data. After using Unseen by Re-training, VE SDE can forget some of the NSFG data, shown in Fig. 2 (b) Right, but not completely forget them. This is a typical example of unintended generalization.

Besides visual inspection, we quantify the generation probability of NSFG and SFG data in terms of Negative Loglikelihood (NLL) given different generators in Tab. 1. For Standard VE SDE, it is reasonable for \mathcal{D}_q and \mathcal{D}_f to have similar likelihoods, as both data are observed during training. However, for Unseen by Re-training, the likelihood of \mathcal{D}_f is almost the same as \mathcal{D}_q . This demonstrates that Unseen by Re-training cannot force the generator to forget the NSFG data well. The visualization of the score functions in Fig. 1 shows that Unseen by Re-training does not alter the original score function of Standard VE SDE, explaining why VE SDE under Unseen by Re-training can still generate NSFG data. We further investigate the existence of unintended generalization in widely-used real-world datasets. Specifically, we conduct experiments on the CELEBA [1] and MNIST [24] datasets. For the CELEBA dataset, we exclude training samples containing the 'bang' attribute. As illustrated in Fig. 3 (left), the model trained under the Unseen by Re-training paradigm retains the capability to generate realistic images with the 'bang' feature, despite its absence from the training data. Similarly, in the MNIST dataset, we remove all samples corresponding to the digits '3' and '7' from the training set. Remarkably, the model still generates digits that closely resemble '3' and '7'. These findings provide empirical evidence that unintended generalization persists in high-dimensional data, challenging the reliability of the current 'gold standard' for MU in SGMs. This raises significant concerns regarding the robustness and trustworthiness of existing MU methodologies.

3.3. Moderated Score-based Generative Model

To forget the NSFG data, the design object of Unseen by Re-training is to remove \mathcal{D}_f from training data and only approximate $p_g(\mathbf{x})$, *i.e.*, decreasing the distance between $p_g(\mathbf{x})$ and $p_\theta(\mathbf{x})$,

$$\theta^* = \arg\min_{\theta \in \Theta} D(p_{\theta}(\mathbf{x}), p_g(\mathbf{x})),$$
 (4)

where the distance between these two distributions $D(p_{\theta}(\mathbf{x}), p_{g}(\mathbf{x}))$ can be evaluated using some distance metrics. Using score matching, Unseen by Re-training in SGM

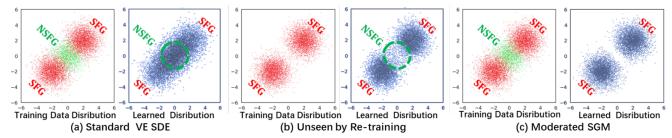


Figure 2. The samples from the mixture Gaussian distribution and the samples generated by the model trained by Standard VE SDE (a), Unseen by Re-training (b) and Unlearning Re-training (c). The left side of (a), (b) and (c) represents the training data, in which the green part is NSFG data, and the red part is SFG data. The right side of (a), (b) and (c) represents the data generated by diffusion models.



Figure 3. Unintended generalization in high-dimensional data via Unseen-by-Re-training: (Left) CELEBA: Persistent feature retention (yellow boxes: generated "bangs") despite exclusion from training data (green boxes). (Right) MNIST: Generation of unlearned classes '3' & '7' (yellow) omitted from training set (green).

can train a score network to estimate the score of the distribution $p_q(\mathbf{x})$,

$$L_g = \lambda(t) \left\{ \mathbb{E}_{\mathbf{x}(0)} \mathbb{E}_{\mathbf{x}(t)} [\| \mathbf{s}_{\theta}^{\mathbf{u}}(\mathbf{x}^g(t), t) - \nabla_{\mathbf{x}^g(t)} \log p_{0t}(\mathbf{x}^g(t) \mid \mathbf{x}^g(0)) \|_2^2] \right\}, \ \mathbf{x}^g \in \mathcal{D}_g.$$
 (5)

Although Unseen by Re-training has approximated $p_g(\mathbf{x})$ and the model generates data that follows $p_g(\mathbf{x})$ with high likelihood. However, it does not consider the likelihood of generating \mathcal{D}_f . If the distributions $p_f(\mathbf{x})$ and $p_g(\mathbf{x})$ are close or overlapping, Unseen by Re-training may not control the probability of generating \mathcal{D}_f (see Fig. 2). Therefore, we propose a new Moderated Generalization strategy to prevent the generator from generating undesired content by maximizing the distance between $p_f(\mathbf{x})$ and $p_\theta(\mathbf{x})$, while minimizing the distance between $p_g(\mathbf{x})$ and $p_\theta(\mathbf{x})$, i.e.,

$$\arg\min_{\theta\in\Theta} \left\{ D(p_{\theta}(\mathbf{x}), p_{g}(\mathbf{x})) - D(p_{\theta}(\mathbf{x}), p_{f}(\mathbf{x})) \right\}.$$
 (6)

Unlike Unseen by Re-training only approximating $p_g(\mathbf{x})$, Moderated Generalization aims to ensure that the generator assigns low likelihood to \mathcal{D}_f and high likelihood to \mathcal{D}_g . However, it is not straightforward to instantiate Moderated Generalization for SGMs. We need to re-consider the Moderated Generalization from the view of score functions. Score estimation is crucial in the generation process of SGMs, as it enables the model to capture data distributions accurately. Theoretically, as long as the score estimation is sufficiently accurate and the forward diffusion time is infinite (allowing the noise distribution to approach the prior distribution), diffusion models can approximate any continuous data distribution with polynomial complexity under weak conditions [4].

Consequently, Eq. (6) can be reframed as a score estimation problem, where different score functions are estimated for $p_g(\mathbf{x})$ and $p_f(\mathbf{x})$. The challenge then becomes how to train a time-dependent score-based model $s_{\theta}^{\mathbf{u}}(\mathbf{x}, \mathbf{t})$ to approximate $\nabla_{\mathbf{x}^g} \log p_t(\mathbf{x}^g)$ and deviate $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$. For approximating $p_g(\mathbf{x})$, we can directly use Eq. (5). For unlearning $p_f(\mathbf{x})$, if the estimated score at any moment deviates from the score of the NSFG data on the timeline from 0 to T, the samples generated during sampling will be far away from the data distribution $p_f(\mathbf{x})$. Under this goal, a straightforward idea is to reduce the correlation between $s_{\theta}^{\mathbf{u}}(\mathbf{x}, \mathbf{t})$ and $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$, i.e. minimizing the dot product of the two scores:

$$L_f = \lambda(t) \left\{ \mathbb{E}_{\mathbf{x}(0)} \mathbb{E}_{\mathbf{x}(t)} [\| s_{\theta}^{\mathbf{u}}(\mathbf{x}^f(t), t) + \nabla_{\mathbf{x}^f(t)} \log p_{0t}(\mathbf{x}^f(t) \mid \mathbf{x}^f(0)) \|_2^2 \right\}, \mathbf{x}^f \in \mathcal{D}_f.$$
(7)

Eq. (7) seeks for the orthogonal complement space of $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$, such that for $\forall \mathbf{x}^f \in \mathcal{D}_f$, $s_{\theta}^{\mathbf{u}}(\mathbf{x}^f, \mathbf{t})$. $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f) \to 0$. We refer to this unlearning optimization as Orthogonal-Moderated Score-based Generative Model, or *Orthogonal-MSGM*. However, in our preliminary experiments, we observed that when $p_q(\mathbf{x})$ and $p_f(\mathbf{x})$ are very close (e.g. when generating human faces where local features like bangs or beards are undesirable) or when $s_{\theta}^{\mathbf{u}}(\mathbf{x}^{\mathrm{f}}, \mathbf{t})$ has been learned well (e.g. erasing undesirable content from a converged pre-trained generator), strictly enforcing the orthogonality constraint in $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$ becomes challenging. To address this issue, we expand the search space by relaxing the orthogonality constraint to a negatively correlated score subspace, or an Obtuse half-space, defined by $s_{\theta}^{\mathbf{u}}(\mathbf{x}^f(t), t) \cdot \nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f) < 0, \ \forall \mathbf{x}^f \in \mathcal{D}_f$. This relaxation leads us to propose a new unlearning objective called *Obtuse-MSGM*:

$$L_f = \lambda(t) \left\{ \mathbb{E}_{\mathbf{x}(0)} \mathbb{E}_{\mathbf{x}(t)} [s_{\theta}^{\mathbf{u}}(\mathbf{x}^f(t), t) + \nabla_{\mathbf{x}^f(t)} \log p_{0t}(\mathbf{x}^f(t) \mid \mathbf{x}^f(0))] \right\}, \mathbf{x}^f \in \mathcal{D}_f.$$
(8)

The final loss of MSGM can be expressed as:

$$\min_{\theta} L_{MSGM} = \min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(0,T)} \left(\alpha L_g + (1 - \alpha) L_f \right), \tag{9}$$

where $\mathcal{U}(0,T)$ is a uniform distribution over [0,T], $p_{0t}(\mathbf{x}(t) \mid \mathbf{x}(0))$ denotes the transition probability from $\mathbf{x}(0)$

to $\mathbf{x}(t)$, $\lambda(t) \in \mathbb{R}_{>0}$ denotes a positive weighting function and α is a hyperparameter.

In contrast to Unseen by Re-training, MSGM modifies the original function of NSFG data. To demonstrate it, we conduct a quick experiment on the mixture Gaussian distribution to evaluate the effectiveness of MSGM. We plot the learned scores at a randomly selected generation process t=0.08 in Fig. 1. The results show that the scores for both Unseen by Re-training and Standard VESDE are quite similar, while our method alters the score distribution of NSFG data, causing the model to steer away from high probability density areas, thereby reducing the likelihood of generating NSFG data. As shown in Fig. 2, compared to Unseen by Re-training, samples generated by our method almost do not contain NSFG data. Meanwhile, the NLL values in Tab. 1 indicate a substantial decrease in the probability of generating NSFG data.

4. Experiments

4.1. Experimental Setup

Datasets and Models. We evaluate MSGM on five datasets, including MNIST [1], CIFAR-10 [22], STL-10 [6], CelebA [24] and high-resoluation Imagenette [18] datasets. In addition to evaluating MSGM on score-based models such as the Variance Preserving (VP) SDE [42] and the VE SDE [42], we also employ DDPM [17] to verify the generalization of MSGM to different types of diffusion generative models, including the latent diffusion-based Stable Diffusion (SD) [34]. Based on the characteristics of the datasets, we performed class forgetting experiments on MNIST, CI-FAR10, STL-10, and Imagenette datasets, while applying attribute elimination generation on CelebA dataset.

Next, we outline the datasets preparation for the experiments. For MNIST, we trained the VE SDE models, selecting all instances of the digits '3' and '7' for \mathcal{D}_f . For CIFAR-10, we trained the VP SDE and DDPM models, selecting the data labeled as 'dog' and 'automobile' classes for \mathcal{D}_f . For STL-10, we trained the VP SDE models, selecting the data labeled as the 'airplane' class for \mathcal{D}_f . For CelebA, we trained the VP SDE models, selecting the feature 'Bangs' from the 40 available features provided for each image to form \mathcal{D}_f . For Imagenette, we fine-tuned SD v1.4, designating 'tench' class as \mathcal{D}_f .

Compared Methods. We establish the following baselines for comprehensive comparison in our experiments: **Standard**, the conventionally trained generative model without any unlearning intervention, serving as the performance reference; **Unseen by Re-training (Unseen)**, the gold standard in MU that retrains models from scratch on sanitized datasets; **EraseDiff** [45], a state-of-the-art diffusion model unlearning method leveraging gradient ascent; and **ESD** [10], a parameter-space editing approach for conditional predic-

tion adjustment. In our experiments, we evaluate two variants of our proposed method, MSGM: Orthogonal-MSGM (Ort) and Obtuse-MSGM (Obt). For unconditional generation evaluation, we benchmark both MSGM variants against Standard, EraseDiff, and Unseen baselines. In text-to-image generation tasks, we compare our method against EraseDiff and ESD, as they are specifically designed for conditional generation scenarios. For downstream tasks, we compare our method with Standard and Unseen, representing the extremes of naive training and full retraining, to highlight the trade-offs between performance and unlearning.

Evaluation Metric. We employ the Unlearning Ratio (UR) and Negative Log-Likelihood (NLL). UR measures the percentage of generated images containing NSFG content, where a lower UR indicates a stronger capability of the model to forget NSFG data. We use external classifiers or CLIP to distinguish whether NSFG categories or features have been removed from the synthetic image. For all experiments, we randomly sample 10,000 images from the model to calculate the unlearning ratio. Additionally, for SGMs, we can accurately calculate NLL to determine the likelihood of generating NSFG and SFG data. Higher values indicate a lower probability of generation. For visual quality evaluation, we follow the protocols of EraseDiff [45] to use FID [16], CLIP embedding distance [32], PSNR and SSIM.

4.2. Class-wise/Feature-wise Ungeneration

Quantitative Results. In Tab. 2, we compare the unlearning performance with baseline methods in unconditional generation. First, MSGM achieves the lowest unlearning rate compared to Unseen across all datasets, indicating that MSGM effectively unlearns the NSFG data. Second, for Unseen by Re-training, both SFG and NSFG data exhibit low NLL values, suggesting that despite the NSFG data never being observed during the training process, the generative model can still fit the distributions $p_f(\mathbf{x})$ well. EraseDiff only slightly decreases the likelihood of generating NSFG data. In contrast, MSGM significantly reduces the generation probability of \mathcal{D}_f via substantially increasing the NLL values of the NSFG data. Additionally, although both Orthogonal-MSGM and Obtuse-MSGM can successfully unlearn undesirable data/features, their performance varies across different scenarios. Orthogonal-MSGM is more effective for class unlearning, while Obtuse-MSGM is more effective for feature unlearning. We suspect that Orthogonal-MSGM seeks orthogonal complement space of $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$, so that $s_{\theta}^{\mathbf{u}}(\mathbf{x}, t)$ does not learn any semantic features(see Fig. 4), hence Orthogonal-MSGM is effective for most cases. However, when $p_q(\mathbf{x})$ and $p_f(\mathbf{x})$ are very close, the orthogonal complement space of $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$ is hard to be found, hence using Obtuse-MSGM to extend the search space $(s_{\theta}^{\mathbf{u}}(\mathbf{x}^f(t), t) \cdot \nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f) < 0, \forall \mathbf{x}^f \in \mathcal{D}_f)$ can improve the unlearning performance.

TD 11 0 0	` ''' '' 1.	c 1 ·	1 ' 11 C '		1'CC , 1	1 1	
Table 7 (Juantitative results	s of linlearnin	g undesirable features	or classes lising	t different linlearn	nno methods acro	se various datacets
14010 2.	Zuantituti ve resurt	or unicum	g undesirable realares	or crasses asing	, different diffedir	mig memous acro	ob various datasets.

Dataset	Model	Feature/Class		Unlear	ning Ra	atio (%) (↓)	_ Test	Negative Log-Likelihood ($\mathcal{D}_g \left(\downarrow\right)$ and $\mathcal{D}_f \left(\uparrow\right)$)					
Dunger	1110001		Standard	Ort	Obt	Unseen	EraseDiff	- 1050	Standard	Ort	Obt	Unseen	EraseDiff	
MNIST	VESDE	3 7	11.0 15.8	0.4 0.8	1.5 3.6	1.8 2.3	1.4 1.1	\mathcal{D}_g	2.82	3.92	3.70	3.07	3.30	
		3 and 7	26.8	1.2	5.1	4.1	2.5	\mathcal{D}_f	2.78	13.23	12.08	3.01	3.74	
CIFAR-10	VPSDE	automobile dog	11.2 13.4	1.9 10.0	0.9 11.5	3.4 10.8	9.7 8.2	\mathcal{D}_g	3.12	3.22	3.28	3.09	3.09	
		automobile and dog	24.6	11.9	12.4	14.2	17.9	\mathcal{D}_f	3.20	5.94	4.37	3.21	4.10	
STL-10	VPSDE	airplane	12.1	2.4	3.6	3.8	2.6	$\left \begin{array}{c} \mathcal{D}_g \\ \mathcal{D}_f \end{array} \right $	2.90 2.19	2.90 8.94	2.92 9.25	2.90 2.32		
CelebA	VPSDE	bangs	19.6	3.5	0.7	6.7	1.2	-	-	-	-	-	-	



Figure 4. Image generation using different unlearning methods for VE SDE on MNIST and VP SDE on CelebA. The top, middle, and bottom rows show images generated by MU strategy Unseen, Ort and Obt respectively. NSFG images sampled from the forgetting dataset \mathcal{D}_f are enclosed in the green box. Images generated by the different unlearning methods are enclosed in the yellow box.

Table 3. Fine-tune quantitative results for unleaning feature or class on different datasets. The Unlearning Ratio represents the degree of forgetting, measured by predicting the proportion of \mathcal{D}_f data in the generated 10,000 images using CLIP.

Dataset	Model	Feature/Class	Unlearning Ratio (%) (↓)								
Dutuset	Model	1 cuture/ Class	Stand	Ort	Obt	Unseen	EraseDiff				
		automobile	11.2	2.7	0.6	3.4	9.4				
	VPSDE	dog	13.4	8.7	8.9	10.8	5.2				
CIFAR-10		automobile and dog	24.6	11.4	9.5	14.2	14.6				
CII AIC-10		automobile	13.1	3.3	1.6	2.7	3.0				
	DDPM	dog	13.9	5.4	3.6	4.5	4.4				
		automobile and dog	27.0	8.7	5.2	7.2	7.4				
CelebA	VPSDE	bangs	19.6	2.6	0.1	6.7	1.9				

Table 4. Fine-tuned negative log-likelihood (NLL) values of \mathcal{D}_g and \mathcal{D}_f for CIFAR-10 data on VP SDE.

Test	Standard	Unseen	Unlearning	Unseen	EraseDiff
\mathcal{D}_q	2.89	3.06	4.36	2.92	3.06
\mathcal{D}_f^{σ}	2.91	10.36	14.96	2.95	4.38

Qualitative Results. We report the qualitative visualization comparison in Fig. 4 and observe that Unseen by Re-training may not fully erase the 'bangs' features. For example, facial images generated by Unseen by Re-training may still exhibit few 'bangs' features, even though the 'bangs' features are not as long as those in \mathcal{D}_f . In contrast, MSGM completely erases the 'bangs' features. An interesting phenomenon is that Orthogonal-MSGM and Obtuse-MSGM forget 'bangs' in different ways. For the unwanted fea-

ture, Orthogonal-MSGM replaced the 'bangs' with noisy images, while Obtuse-MSGM generate features opposite to the 'bangs' in the score distribution, such as 'no bangs' or 'hat'. This occurs because Orthogonal-MSGM seeks the orthogonal complement space of $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$, resulting in $s^{\mathbf{u}}_{\mathbf{u}}(\mathbf{x},t)$ learning nothing. In contrast, $s^{\mathbf{u}}_{\mathbf{u}}(\mathbf{x},t)$ in Obtuse-MSGM learns the inverse of $\nabla_{\mathbf{x}^f} \log p_t(\mathbf{x}^f)$, which may generate the 'inverse' feature of 'bangs'. The visual results in other datasets also have the similar phenomenon, as shown in the right side of the Fig. 4. Additionally, for SFG content generation, MSGM shows competitive generative performance compared to the source images, and performs well with high-resolution images.

4.3. Unlearning DDPM and Fine-tune

MSGM is a general and flexible framework that is compatible with DDPM models and fine-tuning training. The technical details of MSGM application to DDPM can be found in supplementary material. To demonstrate this, we conduct both class and feature unlearning on pre-trained VP SDE and DDPM models. Tab. 3 and Tab. 4 present quantitative results for fine-tuning experiments on different datasets using the MSGM method. We conduct 80,000 and 30,000 iterations of fine-tuning on SGM and DDPM architecture respectively, across all datasets. Notably, MSGM not only achieves the best unlearning performance on SGM but also outperforms other baseline on DDPM, even exceeding EraseDiff, an MU method specifically designed for DDPM-based models.

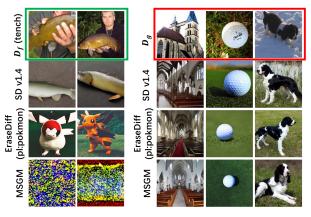


Figure 5. Visualization of diverse unlearning methods applied to fine-tune SD v1.4 on the Imagenette dataset. The left green box displays NSFG images sampled from forgetting datasets. 'pl' indicates the pseudo-label used during training.



Figure 6. Comparison of EraseDiff using semantically similar pseudo-label 'cyprinoid' (for 'tench') vs. our pseudo-label-free MSGM approach.

Table 5. Unlearning text-to-image models on high-resolution Imagenette dataset. CR measures the classification rate of a pre-trained classifier on generated images conditioned on the forgetting classes.

	SD v1.4	ESD	EraseDiff	MSGM
FID of $\mathcal{D}_g(\downarrow)$	4.89	3.09	3.09	3.08
CR(↓)	0.74	0.00	0.00	0.00

4.4. Unlearning Text-to-image Generation on Highresolution Datesets

Although MSGM was initially proposed for unconditional generation, it is a plug-and-play unlearning strategy and also effective for conditional generation, e.g. text-to-image generation. To demonstrate this, we fine-tune Stable Diffusion(SD) v1.4 [34] on its cross-attention layer for class-wise forgetting. We compare with other text-to-image models including EraseDiff [45] and ESD [10] on high-resolution datatset Imagenette [18] and choose the 'tench' as the forgetting class. As shown in Tab. 5, MSGM achieves an FID of 3.08, outperforming the original SD v1.4 and even slightly improving upon the performance of ESD and EraseDiff. Furthermore, while ESD, EraseDiff, and MSGM all achieve a CR value of 0, MSGM differs from other methods by directly generating noisy images for the forgotten class, eliminating the need for pseudo-labels from other classes. In contrast, existing methods often rely on pseudo-labels to replace the forgetton class, which can lead to ineffective forgetting when semantically similar categories are used. This limitation fre-

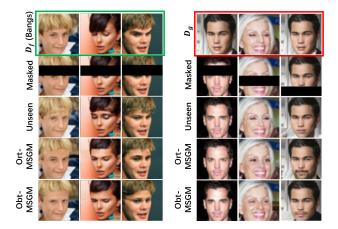


Figure 7. The comparison of inpainting results on the CelebA dataset. The mask size is 64×16 . The restored results on \mathcal{D}_f are displayed on the left. The restored results on \mathcal{D}_g are displayed on the right.



Figure 8. The comparison of reconstruction results on the CIFAR10 dataset. The top, middle and bottom columns are the original images, reconstruction images by Unseen, and reconstruction images by Ort respectively.

quently results in the failure to fully remove the target class. As shown in Fig. 6 and the CR value of 0.47, EraseDiff fails to remove the 'tench' semantic when 'cyprinoid' is used as a pseudo-label. In contrast, MSGM avoids this issue, showing greater robustness and precision in text-to-image unlearning.

4.5. Application to Downstream Tasks

Unleanring Inpainting. MSGM enables zero-shot transfer of the SGM MU to downstream tasks. We first test MSGM on inpainting task. For the class inpainting on CIFAR-10 and STL-10, we mask the upper half of the image and attempt to restore the whole image. For feature inpainting, we mask the region of the feature to be restored (covering 1/4 of the whole image, shown in Fig. 7). The quantitative restoring results on \mathcal{D}_f and \mathcal{D}_g are reported in Tab. 6. We regard the classification as correct if the predicted class of the restored image matches that of the corresponding original image. Obtuse-MSGM still contains a high classification accuracy for restored images on \mathcal{D}_g while significantly decrease the accuracy on restored images on \mathcal{D}_f . This indicates that restored image by Obtuse-MSGM still retains similar semantics on \mathcal{D}_g , while altering the source semantics on \mathcal{D}_g .

Table 6. The inpainting comparison results, where 'stand' refers to the 'Standard' method.

Dataset	Dataset		ACC (%) ($\mathcal{D}_g(\uparrow)$ and $\mathcal{D}_f(\downarrow)$)						FID of $\mathcal{D}_g(\downarrow)$ d Ort Obt Unseen Stand 1 15.96 13.46 13.64 6.80				CLIP of $\mathcal{D}_g(\downarrow)$			PSNR of $\mathcal{D}_g(\uparrow)$			SSIM of $\mathcal{D}_g(\uparrow)$			(†)
											!											
CIFAR-10	\mathcal{D}_g \mathcal{D}_f	95.4 95.5	72.5 75.0	75.5 57.2	74.7 49.6	75.8 59.7	13.11	15.96	13.46	13.64	6.80	6.80	6.77	6.72	31.09	31.09	31.01	31.03	0.56	0.55	0.54	0.54
STL-10	\mathcal{D}_g \mathcal{D}_f	96.3 96.3	83.4 84.1	83.6 59.5	83.1 50.3	84.5 54.9	28.48	29.95	28.55	28.56	8.50	8.51	8.50	8.50	31.18	31.17	31.17	31.18	0.59	0.58	0.57	0.59
CelebA	\mathcal{D}_g \mathcal{D}_f	98.3 98.3				98.0 2.0	29.42	30.31	29.43	30.42	8.96	8.96	8.97	8.94	34.54	34.52	34.50	34.54	0.83	0.82	0.81	0.82

Table 7. The comparison results of reconstruction, where 'stand' denotes the 'Standard' method.

Dataset	ACC (%) ($\mathcal{D}_g(\uparrow)$ and $\mathcal{D}_f(\downarrow)$)				FID of $\mathcal{D}_g(\downarrow)$			CLIP of $\mathcal{D}_g(\downarrow)$			PSNR of $\mathcal{D}_g(\uparrow)$			SSIM of $\mathcal{D}_g(\uparrow)$						
Butuset	Clean	Stand	Ort	Obt	Unseen Stan	d Ort	Obt	Unseen	Stand	Ort	Obt	Unseen	Stand	Ort	Obt	Unseen	Stand	Ort	Obt	Unseen
CIFAR-10 $\frac{\mathcal{D}_g}{\mathcal{D}_f}$	95.4 95.5	88.1 74.4	87.7 48.4	87.0 69.6	87.9 70.3 5.52	5.71	5.57	5.94	6.91	6.90	6.89	6.90	31.91	32.15	32.19	31.82	0.92	0.92	0.93	0.91

Next, for image quality metrics (FID, CLIP, PSNR and SSIM) on \mathcal{D}_g , MSGM matches the standard model, showing high visual quality. Furthermore, we compare the visual results on Fig. 7. When the masked image is from \mathcal{D}_f , Unseen by Re-training still has the probability to restore the 'bangs' feature in the masked region, while MSGM effectively erase the bangs on \mathcal{D}_f . When the masked image is from \mathcal{D}_g , MSGM can successfully restore realistic masked features, such as forehead, nose, mouth etc.

Unlearning Reconstruction. Generative models can learn the latent representations of data and reconstruct images. Through the reconstruction, we use these latent representations as guidance to verify whether our method effectively achieves unlearning. To maintain the similarity between reconstruction results and original images on \mathcal{D}_q , we set t=0.02 for the continuous-time SDE schedule. We reconstruct images using VP SDE model trained by standard training, Unseen by Re-training and MSGM, and report the comparison results in Tab. 7. We utilize the classification accuracy to assess whether the reconstructed images still be classed by the original class. Obtuse-MSGM significantly decrease the accuracy for reconstructed \mathcal{D}_f data while maintaining the original semantic information for reconstructed \mathcal{D}_q . Additionally, we calculate the CLIP distance, PSNR and SSIM between original images and reconstructed images on \mathcal{D}_q . The images reconstructed by MSGM and the standard model have nearly the same numerical results, indicating that the images reconstructed by MSGM have high visual quality. Next, we visualize the reconstruction results in Fig. 8. Unlike Unseen by Re-training, where the reconstruction images on \mathcal{D}_f have the same semantics with the original images, Orthogonal-MSGM reconstructs \mathcal{D}_f as noisy images, indicating that Orthogonal-MSGM has completely unlearned the \mathcal{D}_f distribution.

4.6. Ablation Study

Optimization Choices. We explore two strategies for optimizing Eq. (9): (1) Simultaneously Updating: both L_g and L_f are simultaneously updated in each iterative sampling,

Table 8. Ablation studies on MNIST with Orthogonal-MSGM using different parameters α . $\alpha=1$ means Unseen.

α	Class	$UR(\%)(\downarrow)$	NLI	L Test	$ \alpha$	Class	$UR(\%)(\downarrow)$	NLI	_ Test
0.7	3 7 3 and 7	0.7 1.2 1.9	\mathcal{D}_g \mathcal{D}_f	4.17 13.32	0.99	3 7 3 and 7	0.4 0.8 1.2	\mathcal{D}_g \mathcal{D}_f	3.92 14.75
0.9	3 7 3 and 7	0.5 0.9 1.4	\mathcal{D}_g \mathcal{D}_f	4.05 13.88	1.0	3 7 3 and 7	1.8 2.3 4.1	\mathcal{D}_g \mathcal{D}_f	3.07 3.01

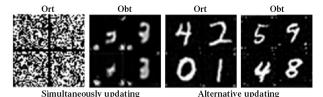


Figure 9. The visual results of Simultaneously Updating vs. Alternative Updating.

and (2) Alternative Updating: L_g is updated in each iterative sampling, while L_f is updated at intervals of every four iterations. We plot the loss curve in supplementary material and visualize the visual results in Fig. 9. We find that L_g converges much more easily than L_f . As a result, alternative updating facilitates better convergence of L_g and improves visual quality. Thus, we adopt this strategy.

 α **Setting.** MSGM uses two losses in Eq. (9): αL_g for generation quality and $(1-\alpha)L_f$ for unlearning. An ablation in Tab. 8 shows $\alpha=0.99$ best balances both, and is thus used by default.

5. Conclusion and Future work

In this work, we introduce MSGM, the first unlearning framework for score-based generative models (SGMs). MSGM outperforms the current "gold standard" and DDPM baselines, erasing unwanted content while preserving the quality of retained data. The plug-and-play design adapts to diverse diffusion backbones, supports zero-shot transfer to downstream tasks, and stays robust under improper guidance. Future extensions will target additional modalities such as skeletal time-series and video.

Acknowledgments. This work was supported by the National Key Research and Development Program of China (No. 2023YFC2506804), the Jointly Funded Project by the National Natural Science Foundation (No. U23B2031), the Key project of National Natural Science Foundation of China (No. 61932009), the National Natural Science Foundation of China (No. 62302139, 62441617 and 62302003), FRFCU China (No. PA2025IISL0113, JZ2025HGTB0227, KG16336301), the Postdoctoral Fellowship Program and China Postdoctoral Science Foundation (No. 2024M764093, BX20250485), the Beijing Natural Science Foundation (No. 4254100), the Anhui Provincial Natural Science Research Program for Higher Education Institutions (No. 2022AH050082) and by Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing.

References

- Areej Alsaafin and Ashraf Elnagar. A minimal subset of features using feature selection for handwritten digit recognition.
 Journal of Intelligent Learning Systems and Applications, 9 (04):55–68, 2017. 3, 5
- [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. 2
- [3] Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramer, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In 32nd USENIX Security Symposium (USENIX Security 23), pages 5253–5270, 2023. 1
- [4] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *International Conference on Learning Representations*, 2023. 4
- [5] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical image analysis*, 80: 102479, 2022. 2
- [6] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In Proceedings of the fourteenth international conference on artificial intelligence and statistics, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. 5
- [7] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In 5th International Conference on Learning Representations, ICLR 2017, 2017.
- [8] Chongyu Fan, Jiancheng Liu, Yihua Zhang, Dennis Wei, Eric Wong, and Sijia Liu. Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. In *International Conference on Learning Representations*, 2024. 2, 3
- [9] Pierre-Étienne H Fiquet and Eero P Simoncelli. Video prediction using score-based conditional density estimation. arXiv preprint arXiv:2411.00842, 2024. 2
- [10] Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts from diffusion

- models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2426–2436, 2023. 1, 2, 5, 7
- [11] Rohit Gandikota, Hadas Orgad, Yonatan Belinkov, Joanna Materzyńska, and David Bau. Unified concept editing in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5111– 5120, 2024.
- [12] Chao Gong, Kai Chen, Zhipeng Wei, Jingjing Chen, and Yu-Gang Jiang. Reliable and efficient concept erasure of textto-image diffusion models. arXiv preprint arXiv:2407.12383, 2024. 2
- [13] Ian 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. 1
- [14] Alvin Heng and Harold Soh. Continual learning for forgetting in deep generative models. 2023. 2
- [15] Alvin Heng and Harold Soh. Selective amnesia: A continual learning approach to forgetting in deep generative models. Advances in Neural Information Processing Systems, 36, 2024.
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems, 30, 2017.
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 2, 5
- [18] Jeremy Howard and Sylvain Gugger. Fastai: a layered api for deep learning. *Information*, 11(2):108, 2020. 5, 7
- [19] Zheyuan Hu, Zhongqiang Zhang, George Em Karniadakis, and Kenji Kawaguchi. Score-based physics-informed neural networks for high-dimensional fokker-planck equations. arXiv preprint arXiv:2402.07465, 2024. 2
- [20] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, 2014. 1
- [21] Zhifeng Kong and Kamalika Chaudhuri. Data redaction from pre-trained gans. In 2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML), pages 638–677. IEEE, 2023. 2
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [23] Nupur Kumari, Bingliang Zhang, Sheng-Yu Wang, Eli Shechtman, Richard Zhang, and Jun-Yan Zhu. Ablating concepts in text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22691–22702, 2023. 2
- [24] 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*, pages 3730–3738, 2015. 3, 5
- [25] Shilin Lu, Zilan Wang, Leyang Li, Yanzhu Liu, and Adams Wai-Kin Kong. Mace: Mass concept erasure in diffusion

- models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6430–6440, 2024. 2
- [26] Chenlin Meng, Yang Song, Wenzhe Li, and Stefano Ermon. Estimating high order gradients of the data distribution by denoising. Advances in Neural Information Processing Systems, 34:25359–25369, 2021. 2
- [27] Saemi Moon, Seunghyuk Cho, and Dongwoo Kim. Feature unlearning for pre-trained gans and vaes. In *Proceedings of* the AAAI Conference on Artificial Intelligence, pages 21420– 21428, 2024. 2
- [28] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International* conference on machine learning, pages 8162–8171. PMLR, 2021. 2
- [29] Hadas Orgad, Bahjat Kawar, and Yonatan Belinkov. Editing implicit assumptions in text-to-image diffusion models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 7053–7061, 2023. 2
- [30] Tianyu Pang, Kun Xu, Chongxuan Li, Yang Song, Stefano Ermon, and Jun Zhu. Efficient learning of generative models via finite-difference score matching. Advances in Neural Information Processing Systems, 33:19175–19188, 2020. 3
- [31] Frank Pasquale and Gianclaudio Malgieri. Generative ai, explainability, and score-based natural language processing in benefits administration. *Journal of Cross-disciplinary Research in Computational Law*, 2(2), 2024. 2
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 5
- [33] Javier Rando, Daniel Paleka, David Lindner, Lennard Heim, and Florian Tramèr. Red-teaming the stable diffusion safety filter. arXiv preprint arXiv:2210.04610, 2022. 1, 2
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 5, 7
- [35] Hadi Salman, Alaa Khaddaj, Guillaume Leclerc, Andrew Ilyas, and Aleksander Madry. Raising the cost of malicious ai-powered image editing. arXiv preprint arXiv:2302.06588, 2023. 1, 2
- [36] Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 22522–22531, 2023. 1, 2
- [37] Thanveer Shaik, Xiaohui Tao, Haoran Xie, Lin Li, Xi-aofeng Zhu, and Qing Li. Exploring the landscape of machine unlearning: A survey and taxonomy. arXiv preprint arXiv:2305.06360, 2023. 2
- [38] Shawn Shan, Jenna Cryan, Emily Wenger, Haitao Zheng, Rana Hanocka, and Ben Y Zhao. Glaze: Protecting artists from style mimicry by text-to-image models. *arXiv preprint arXiv:2302.04222*, 2023. 1, 2

- [39] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 2
- [40] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. Advances in neural information processing systems, 33:12438–12448, 2020. 2
- [41] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pages 574–584. PMLR, 2020. 3
- [42] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, *ICLR* 2021, 2021. 1, 2, 3, 5
- [43] Hui Sun, Tianqing Zhu, Wenhan Chang, and Wanlei Zhou. Generative adversarial networks unlearning. arXiv preprint arXiv:2308.09881, 2023. 2
- [44] Anvith Thudi, Hengrui Jia, Ilia Shumailov, and Nicolas Papernot. On the necessity of auditable algorithmic definitions for machine unlearning. In *31st USENIX Security Symposium* (USENIX Security 22), pages 4007–4022, 2022. 2, 3
- [45] Jing Wu, Trung Le, Munawar Hayat, and Mehrtash Harandi. Erasediff: Erasing data influence in diffusion models. *arXiv* preprint arXiv:2401.05779, 2024. 2, 5, 7
- [46] Yi-Chiao Wu, Dejan Marković, Steven Krenn, Israel D Gebru, and Alexander Richard. Scoredec: A phase-preserving highfidelity audio codec with a generalized score-based diffusion post-filter. In ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 361–365. IEEE, 2024. 2
- [47] Jie Xu, Zihan Wu, Cong Wang, and Xiaohua Jia. Machine unlearning: Solutions and challenges. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2024. 3
- [48] Gong Zhang, Kai Wang, Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Forget-me-not: Learning to forget in text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1755–1764, 2024. 2