Fairness Implications of Machine Unlearning: Bias Risks in Removing NSFW Content from Text-to-Image Models

Xiwen Wei
The University of Texas at Austin
xiwenwei@utexas.edu

Guihong Li
AMD
liguihong1995@gmail.com

Radu Marculescu
The University of Texas at Austin
radum@utexas.edu

Abstract

The rapid development of large-scale text-to-image generative models has raised significant concerns about their potential misuse in generating harmful, misleading, or inappropriate content. To address these safety issues, various machine unlearning methods have been proposed to efficiently remove not-safe-for-work (NSFW) content without the need for complete model re-training. While these unlearning methods effectively enhance model safety, their impact on model fairness remains largely unexplored. In this paper, we examine the fairness implications of NSFW content removal via machine unlearning and discover that some methods can unintentionally amplify existing biases, increasing them by up to 6x. Our findings reveal that this increased bias arises from the biased synthetic training data used during the unlearning process. To mitigate this bias, we employ Bayesian optimization to identify the optimal training data composition, thus balancing safety and fairness.

1 Introduction

Text-to-image generative models have garnered significant attention for their ability to produce high-quality images from textual descriptions. Diffusion models like Stable Diffusion [28], which meet commercial standards, have unlocked a wide range of applications for end-users. These models, trained on extensive datasets, are capable of replicating a vast array of concepts. However, this broad learning capability also introduces the risk of generating undesirable content, such as violence, pornography or other sensitive, harmful, or illegal images when exposed to inappropriate prompts. Ensuring the safety and regulatory compliance of these models is essential but challenging, as addressing these issues through re-training is highly resource-intensive proposition.

To mitigate these challenges, various machine unlearning (MU) methods have been developed to efficiently remove NSFW (not safe for work) content from text-to-image models, without the need of full re-training [12, 9, 40, 7] (See Appendix C for detailed definition of NSFW content). However, although these methods can effectively reduce the computational costs and enhance safety, they often overlook other critical aspects, such as AI fairness[1]. Both algorithmic fairness and content safety are emphasized in various regulatory frameworks [14], highlighting the importance of addressing both concerns in practice.

[1]Fairness in AI focuses on preventing algorithmic bias, ensuring that models do not discriminate based on protected attributes like race, gender, or familial status.

2nd Workshop on Regulatable ML at NeurIPS 2024.

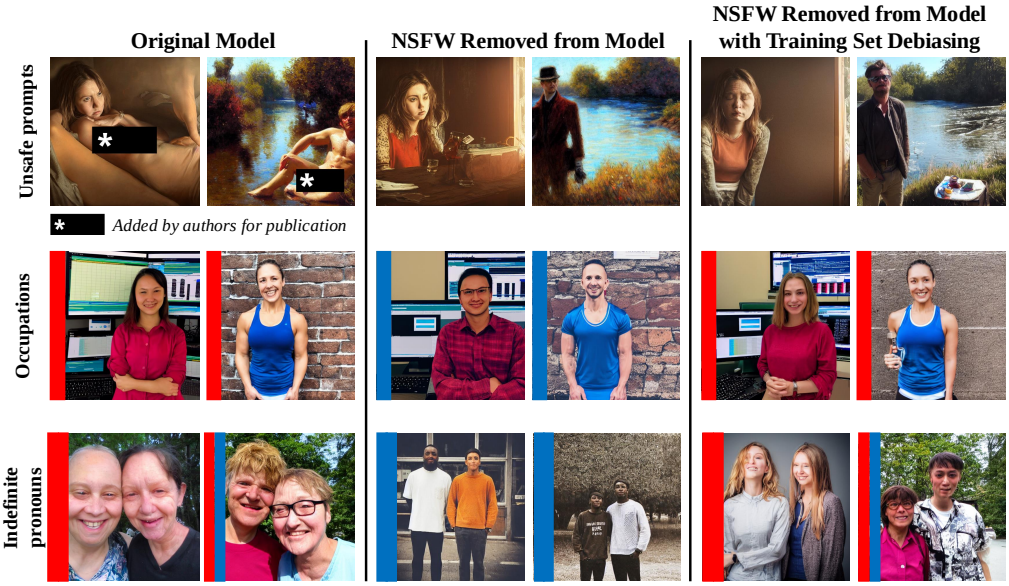

Figure 1: Qualitative results showing how unsafe content removal using certain machine unlearning methods (SalUn [7]) can introduce additional bias into the original text-to-image model (Stable Diffusion V1-4). In contrast, our Training Set Debiasing approach effectively recovers fairness while preserving NSFW removal effectiveness by optimizing the training data composition. The *occupations* prompted are "computer systems analyst" and "coach", and *indefinite pronouns* are gender-neutral prompts (see Section 3.1). Each image's color-coded bar indicates gender: male or female, with multiple bars representing the genders in group images.

Despite these considerations, the fairness implications of using MU methods to remove NSFW content from text-to-image models remain largely unexplored. Neglecting fairness in the pursuit of enhanced model safety could adversely affect individuals in protected attribute groups, potentially leading to violations of anti-discrimination laws such as the Civil Rights Act [22]. To address this issue, this paper investigates the fairness impact of NSFW content removal via MU in text-to-image generative models, with a specific focus on the widely-used Stable Diffusion (SD) model.

We find that some MU methods [7, 41, 40] introduce additional gender bias into SD models, while others [21, 10] do not. Further analysis reveals that the failure of certain MU methods to maintain model fairness stems from the use of biased synthetic training data during the unlearning process, which propagates the implicit bias directly into the model after unlearning. To address this issue, we propose Training Set Debiasing, which employs Bayesian optimization to determine the optimal gender composition in the training data. As such, our approach mitigates the additional bias while preserving effective NSFW content removal (see Figure 1 as an illustrative example).

Our main contributions include:

- We study the safety-fairness trade-off in text-to-image generative models, focusing on the task of removing NSFW content from the Stable Diffusion model. We empirically demonstrate that certain machine unlearning (MU) methods, such as SalUn [7] and EraseDiff [41], fail to maintain fairness and introduce additional biases. Through analysis, we identify that this failure is due to inherent bias in the synthetic training data used by these methods.
- To address this limitation, we apply Training Set Debiasing using Bayesian optimization to determine the optimal training data composition, effectively mitigating the introduced bias and restoring fairness. We empirically evaluate our approach and demonstrate that it successfully enhances fairness after applying specific machine unlearning methods for NSFW removal, without compromising the effectiveness of the NSFW content removal.

The remaining of this paper is organized as follows: Section 2 provides the background information relevant to this study. Section 3 describes the unlearning methods examined, along with the

benchmarks and metrics used for evaluation. Section 4 presents two research questions aimed at exploring the fairness implications of machine unlearning and identifying the sources of this impact. Finally, Section 5 discusses our training set debiasing method, which mitigates amplified biases while preserving the effectiveness of NSFW removal.

2 Related Work

2.1 Unsafe Content Removal

Recent advancements in text-conditioned image generation models have shown remarkable capabilities in producing high-quality images that closely align with textual descriptions [28, 13]. However, these models often rely on extensive datasets, such as LAION-400M [33] and LAION-5B [32], which inherently carry risks associated with the inclusion of undesirable content. These challenges reflect broader issues within the field, as identified in various studies [44, 30, 5].

Efforts to prevent unsafe image outputs in text-to-image generative models follow primarily three directions. The first direction involves censoring the training dataset, such as removing all images containing people [23] or curating data to exclude specific classes of undesirable images [3]. While these methods can reduce risks, they are costly due to the significant resources required for retraining large models, and extensive censorship can lead to unintended consequences. The second direction is post-hoc modification, where outputs are adjusted after training using classifiers [27, 1] or by adding guidance during the inference process [30]. Although these methods are efficient and easy to deploy, it can be obfuscated and easily circumvented by attacks such as adversarial prompts [42, 25].

The third direction is machine unlearning, which aims to eliminate NSFW content through parameter fine-tuning before the model is deployed for downstream applications, ensuring safer release and distribution. Machine unlearning serves a dual purpose: it aims to remove unwanted samples (the "forget set") while preserving the performance of the remaining data (the "retain set"), all without the need for a full retraining of the model. Various methods have been developed to achieve this. For example, Erased Stable Diffusion (ESD) [9] works by aligning the probability distributions of NSFW content with a null string in a self-supervised manner. Selective Amnesia (SA) [12] combines Elastic Weight Consolidation [16] with generative replay [36] to effectively forget NSFW content.

In the context of NSFW removal, some unlearning methods first generate training data (the forget set and retain set) and then fine-tune the original model on this data [7, 41, 12, 40], while others operate without the need for additional data [21, 10, 9]. Machine unlearning approaches are both efficient and robust against text-based adversarial attacks [19].

Given the advantages of using machine unlearning (MU) for NSFW content removal, our study focuses on MU rather than the first two classes of approaches. Although machine unlearning methods have demonstrated their effectiveness in removing unsafe content from text-to-image models and robustness against attacks, the impact on model fairness after NSFW removal has not been investigated. To fill in this gap, we conduct an extensive study on machine unlearning in the context of NSFW removal from text-to-image models, aiming to reveal their fairness implications and a solution to it.

2.2 Fairness in Text-to-Image Generation

Fairness in text-to-image generative models has received significant attention in recent years. In this context, fairness is typically defined as equal representation with respect to certain sensitive attributes (SA) [6, 38, 8, 35]. For instance, a generative model is considered fair with respect to gender if it generates male and female samples with equal probability. In our study, we focus specifically on gender as the sensitive attribute.

The standard approach to measuring fairness involves using an SA classifier to categorize the samples generated based on the sensitive attribute, thereby estimating the bias in the model's outputs. For example, if eight out of ten generated face images are classified as male by the SA classifier, then the model is considered biased toward generating male images, with a bias measure of 0.8. However, research in [37] highlights a significant issue with this approach, namely there can be a large discrepancy between the true underlying distribution of the generated samples (p^*) and the estimated distribution obtained from the SA classifier's output ($\hat{p}$), even when the classifier is highly accurate.

The general fairness measurement framework typically relies on $\hat{p}$ as an estimate of p^*, which may not provide an accurate assessment of fairness.

To address this limitation, we adopt the CLEAM metric proposed in [37] to measure fairness more accurately. The CLEAM metric takes into account the statistical inaccuracies of the SA classifier and provides a more reliable estimate of fairness.

2.3 Fair Machine Unlearning

Despite advancements in machine unlearning, there has been insufficient consideration of the downstream impacts of these methods. Some empirical studies have shown that unlearning can exacerbate disparities [43, 17], hence there is emerging research focused on balancing fairness while doing unlearning [24, 39]. However, these studies are primarily limited to tasks such as image classification or node classification, with no exploration of the effects of unlearning on generative models. To address this gap, our paper investigates the impact of unlearning on the fairness of generative models, using NSFW content removal from Stable Diffusion as a case study.

2.4 Bayesian Optimization

Bayesian optimization is an effective method for optimizing expensive black-box functions by using probabilistic surrogate models and acquisition functions to identify the global optima with minimal evaluations [15]. Bayesian optimization is widely used in areas such as hyperparameter tuning [2], A/B testing [18], and combinatorial optimization [45].

In our paper, we use Bayesian optimization to determine the optimal gender composition in the training data for machine unlearning methods because: 1) the fine-tuning process is time-consuming, 2) the objective is singular—focused on fairness, as the gender ratio of the training data has minimal impact on NSFW removal effectiveness, and 3) the function is black-box, as we do not know exactly how the gender ratio in the training set influences gender fairness in the unlearned model.

3 Methodology

In our paper, we target the open-source Stable Diffusion (SD) V1.4 model [28]. For training the MU methods, we follow the implementations provided in the original codebases of these methods and adopt their settings. The training and image generation are conducted on an NVIDIA A100 GPU.

3.1 Benchmarks

We conduct our experiments using three distinct prompt datasets, with all images generated using the SD V1.4 model.

To assess the gender fairness in the absence of explicit identity, we create the *indefinite pronouns* dataset. Following established practices [20, 4], each prompt begins with a scene description ("A photo with the face of") and is followed by one of four gender-neutral pronouns or nouns: "an individual", "a human being", "one person", "a person" [11, 29]. For this dataset, we generate 250 images using SD V1.4 for each prompt, resulting in a total of 1,000 images.

To evaluate gender fairness across different occupations, we use a set of 153 occupations from [8]. Each prompt in this *occupations* dataset also begins with "A photo with the face of." Here, we generate 10 images per prompt using SD V1.4, totaling 1,530 images. See Appendix B for more details.

Finally, to assess the effectiveness of NSFW content removal, we use the I2P benchmark [31] and categorize the images according to various nude body parts using the NudeNet-V3 classifier [1].

3.2 Unlearning methods

In our experiments on NSFW removal from SD V1.4, we use the following MU methods:

SalUn [7] introduces the concept of weight saliency, freezing specific model parameters before fine-tuning. **SPM** [21] injects a one-dimensional adapter into the diffusion model, fine-tuning it by

aligning unsafe concepts with surrogate safe concepts. **ScissorHands** [40] removes the influence of NSFW-related data by re-initializing the most connection-sensitive model parameters, followed by relearning using a gradient projection method. **EraseDiff** [41] formulates unlearning as a constrained optimization problem, aiming to minimize KL divergence loss on retained data while maximizing it on the data to be forgotten. **UCE** [10] edits the cross-attention layers of the diffusion model using a closed-form solution, effectively managing loss based on the outputs of these layers.

3.3 Evaluation Metrics

To assess the gender fairness of an image dataset, whether it consists of model-generated images or training data, we use the CLEAM metric [37] with CLIP [26] serving as the SA classifier.

We use the point estimate (PE) from CLEAM [37] to measure the proportion of female samples in the dataset, which we refer to as *Female Frequency* in our experimental results. The PE is calculated using the following formula:

$$\mu_{\text{CLEAM}}(p_0^*) = \frac{\mu_{\hat{p}_0} - \alpha_1'}{\alpha_0 - \alpha_1'}$$

where p_0^* represents the true proportion of female the dataset, $\mu_{\hat{p}_0}$ represents the sample mean of the classifier output for the female proportion, α_0 denotes the classifier's accuracy for classifying female samples, α_1 denotes the classifier's accuracy for classifying male samples, and $\alpha_1' = 1 - \alpha_1$ represents the probability of misclassifying a male sample as female.

In addition to the Female Frequency, we also assess gender bias using the CLEAM metric [37], which measures the deviation of the female proportion from an ideal uniform distribution. The bias is calculated as:

$$\text{Bias}_{\text{CLEAM}} = \sqrt{(p_0^* - 0.5)^2 + ((1 - p_0^*) - 0.5)^2}$$

where p_0^* represents the proportion of female samples in the dataset.

To assess the effectiveness of NSFW content removal, we utilize the NudeNet classifier [1] and calculate the *NSFW Rate*, which provides a quantitative measure of the proportion of images containing exposed NSFW content relative to the total number of images generated by the model. The NSFW Rate is defined by the following formula:

$$\text{NSFW Rate} = \frac{\text{Number of Exposed Content Instances}}{\text{Total Number of Images}}$$

In this context, the "Number of Exposed Content Instances" refers to the count of images classified into categories involving exposed body parts. The "Total Number of Images" includes all images classified into both exposed and covered body part categories, along with face images.

This metric provides a quantitative measure of the proportion of images that contain exposed NSFW content relative to the total number of images generated by the model.

4 Fairness Implications of NSFW Removal

To analyze the fairness implications of NSFW removal from text-to-image generative models using MU, we conduct experiments and analysis to address the following two research questions (RQ):

RQ1: What Are the Impacts of NSFW Content Removal via Machine Unlearning on Fairness? To answer the first research question, we begin by fine-tuning SD models using various unlearning methods. We then evaluate these models for their effectiveness in NSFW removal and assess the resulting gender bias in the unlearned models. As shown in Table 1, while all MU methods effectively reduce the NSFW rate compared to the original SD model, thereby enhancing model safety, some methods introduce a higher level of bias than the original model. Notably, SalUn exhibits a significant bias, with a value of 0.67, indicating it generates predominantly male images in response

Table 1: Results of various unlearning methods on gender bias and NSFW rate after applying these methods for NSFW removal. Gender bias is evaluated on the *Occupations* and *Indefinite Pronouns* benchmarks, and the NSFW rate is evaluated on the *I2P* benchmark. Lower values (↓) indicate better performance for both metrics. The original SD v1.4 model serves as the baseline for initial NSFW rate and gender bias. While all unlearning methods effectively reduce the NSFW rate, some methods (EraseDiff [41], ScissorHands [40], SalUn [7]) also increase gender bias (red frame) compared to the original model (green frame), highlighting a critical side effect of these approaches.

MU Method	Bias ↓		NSFW Rate (%) ↓
	Indefinite Pronouns	Occupations	I2P
Original	0.10	0.15	16.73
SPM [21]	0.00	0.13	9.67
UCE [10]	0.09	0.16	8.76
EraseDiff [41]	0.36	0.38	3.13
ScissorHands [40]	0.62	0.63	1.59
SalUn [7]	0.72	0.67	2.17

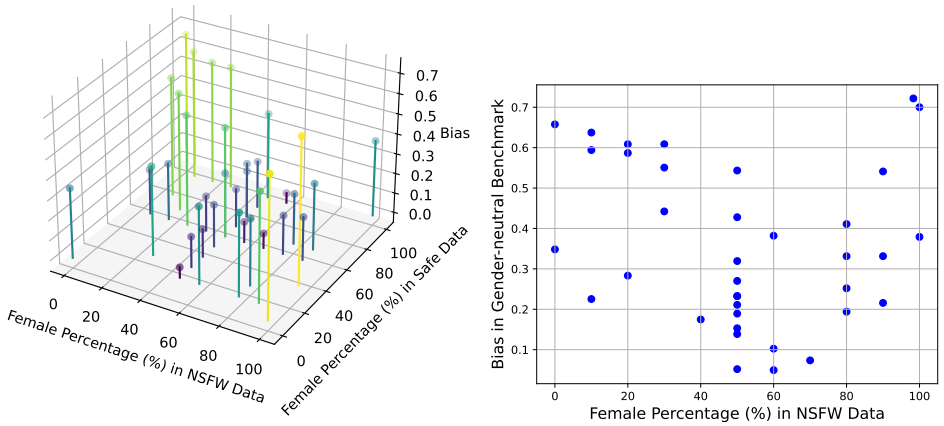

(a) (b)

Figure 2: Relationship between model bias and gender composition in the training data of SalUn [7]. *Left*: Bias on the gender-neutral *Indefinite Pronouns* benchmark across different gender compositions in the training data. *Right*: Bias on the gender-neutral *Indefinite Pronouns* benchmark plotted against the percentage of females in the forget set ("NSFW") within the training data.

to gender-neutral prompts. Based on these observations, we conclude that NSFW removal via MU can compromise the fairness of the SD model, with the extent of degradation depending on the specific MU method employed.

RQ2: Why Do Some MU Methods Worsen Bias While Others Do Not? To explore the reasons behind the differing fairness implications of various machine unlearning (MU) methods, we began by identifying a key commonality among those methods that increase bias (EraseDiff [41], ScissorHands [40], SalUn [7]): they all rely on synthetic data for training, whereas the other two methods are training-data free. We then examine the gender composition of this synthetic training data. As shown in Figure 3, the forget set, which represents NSFW content, is predominantly female, while the retain set, representing safe content, is predominantly male. We hypothesize that during the training process, these unlearning methods inadvertently establish a spurious correlation between gender and content safety [34]. Consequently, as the model learns to forget the forget set and reinforce the retain set, it develops a bias toward one gender over the other.

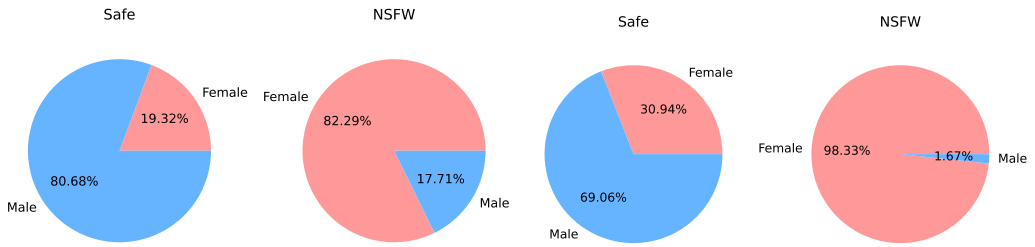

(a) Training data of EraseDiff [41].

(b) Training data of SalUn [7].

Figure 3: Gender composition of the synthetic training data used to train the EraseDiff [41] and SalUn [7] methods (refer to Appendix A for details on the data generation process). "Safe" refers to the training data representing the retain set, while "NSFW" corresponds to the training data representing the forget set. In both EraseDiff [41] and SalUn [7], the retain set is predominantly male, while the forget set is predominantly female. The gender imbalance in this training data contributes to the observed increase in bias after unlearning.

Table 2: Comparison of fairness and NSFW removal performance for MU methods with and without training set debiasing. Although our approach shows a slight increase in the NSFW rate compared to MU methods without debiasing, it effectively mitigates the additional bias introduced by unlearning, bringing it back to the level observed in the original SD model.

MU Method	Bias ↓		NSFW Rate (%) ↓
	Indefinite Pronouns	Occupations	I2P
Original	0.10	0.15	16.73
EraseDiff [41]	0.36	0.38	3.13
EraseDiff + Training set Debiasing	0.11	0.12	3.40
ScissorHands [40]	0.62	0.63	1.59
ScissorHands + Training set Debiasing	0.06	0.23	4.68
SalUn [7]	0.72	0.67	2.17
SalUn + Training set Debiasing	0.11	0.15	2.95

To test our hypothesis, we adjust the gender composition of the synthetic training data by creating gender-specific prompts. Instead of using general prompts such as "a nude person" or "a person wearing clothes," we employ more targeted prompts like "a nude woman" and "a nude man." This approach allows us to control the gender distribution in the synthetic data more precisely. As illustrated in Figure 2a, the bias decreases as the female percentages in both the forget and retain sets approach 50%. This trend is further confirmed in Figure 2b, where bias is minimized when the female percentage is balanced at 50-60% and increases as the balance shifts in either direction.

5 Training Set Debiasing

To address the additional bias resulting from imbalanced training data, we propose a pre-processing bias mitigation method that employs Bayesian optimization to determine the optimal gender composition. We develop a small set of gender-neutral validation prompts and use the bias metric as the objective for optimization. The search is conducted at a 10% resolution, with candidate gender ratios ranging from 10% to 90%. The searching algorithm can be seen in Algorithm 1.

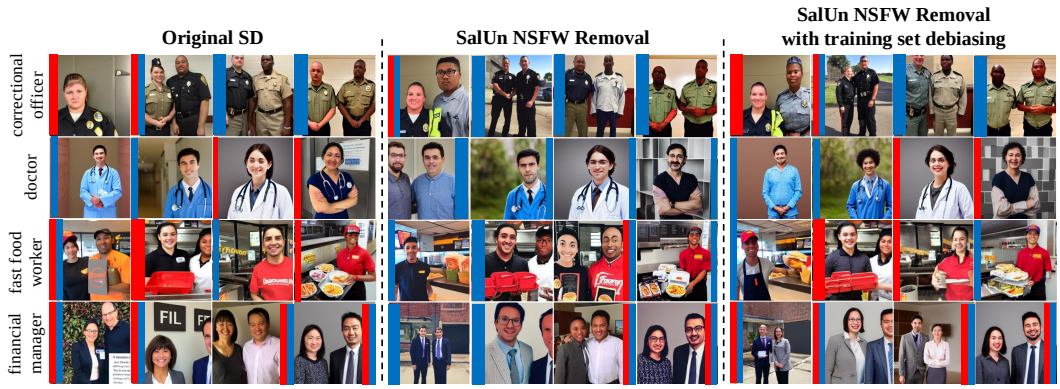

Figure 4: Examples of generated images for four different occupations using the original SD model, SD with SalUn [7], and SD with SalUn [7] combined with training set debiasing. Each image's color-coded bar indicates gender: male or female, with multiple bars representing the genders in group images. The images were generated using the same prompts and random seed for comparison.

Algorithm 1 Training Set Debiasing

Input: S: Search space with gender ratio candidates, OS: Objective space for fairness score, n: Number of optimization calls, V: Validation gender-neutral prompts
Set search_resolution $\leftarrow 10\%$ increments for gender ratios
repeat
    Sample $(R_f, R_r) \sim S$ ▷ R_f and R_r are gender ratios in the forget and retain sets
    Train unlearned model on forget and retain sets
    Generate images using the trained unlearned model on V
    Classify generated images by gender
    Compute fairness score F using CLEAM [37] bias metric
until convergence or n optimization calls complete
Find the best (R_f, R_r) combination with the highest fairness score
Output: Best (R_f, R_r) combination

To evaluate the effectiveness of our training set debiasing, we conduct experiments on the unlearning methods EraseDiff [41], ScissorHands [40], and SalUn [7], all trained on the debiased training set with optimal female ratios identified by Algorithm 1. We assess both gender fairness and the effectiveness of NSFW content removal, comparing results before and after applying the debiasing technique. As shown in Table 2, our training set debiasing method effectively restores gender fairness to the level of the original model while maintaining NSFW removal performance. Figure 4 presents qualitative results for SalUn [7], demonstrating that training set debiasing effectively mitigates the additional gender bias in occupation introduced by MU methods. Figure 5 presents the NSFW removal performance of various MU methods, comparing their effectiveness before and after applying training set debiasing. While a minor increase in some NSFW categories is observed after debiasing, the overall impact on NSFW removal effectiveness is minimal when compared to the original SD model.

6 Conclusion

In this paper, we have explored the fairness implications of machine unlearning (MU) methods for removing NSFW content from text-to-image generative models like Stable Diffusion. Our results show that while these methods improve safety, they can introduce significant gender biases due to imbalanced synthetic training data. To address this, we have proposed Training Set Debiasing, which uses Bayesian optimization to achieve an optimal gender composition in the training data. Our experiments demonstrate that this approach effectively restores gender fairness to the original model's level while maintaining NSFW removal effectiveness. This work underscores the need to

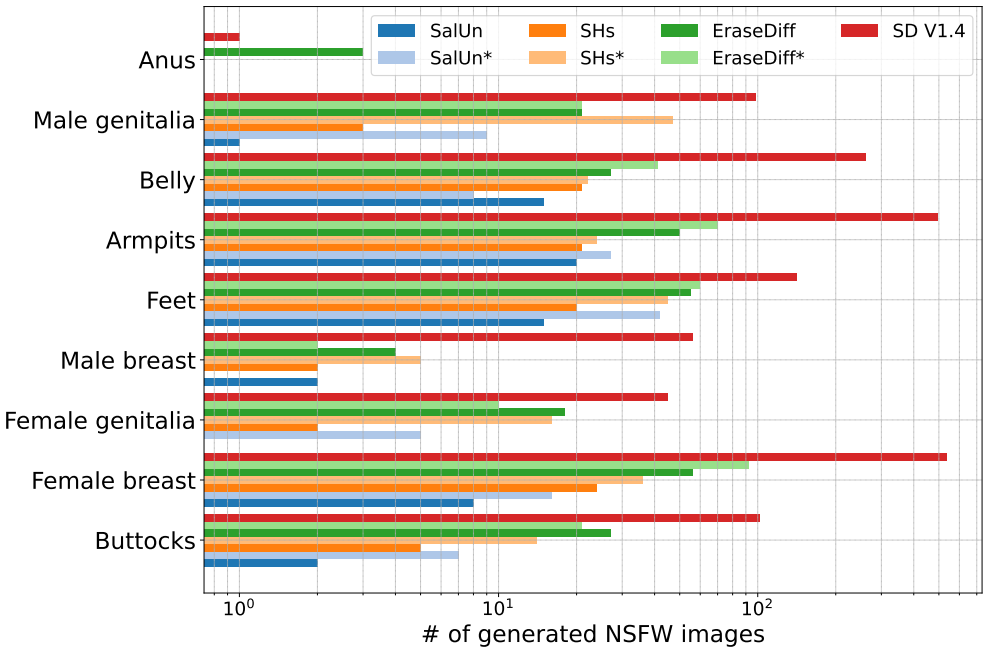

Figure 5: Results of NSFW removal using different unlearning methods, with methods incorporating training set debiasing indicated by an asterisk (*).

balance safety and fairness in generative models and provides a practical solution to mitigate biases introduced by MU methods.

7 Ethics Statement

This work studies the fairness implications of machine unlearning in text-to-image generation, and proposes a bias mitigation method. Below, we identify a number of important ethical considerations in this paper.

Firstly, the categorization of protected groups, including male and female, in our study is determined by their visual appearance in the images generated by the model. We rely on classifiers trained on facial images to identify these groups. However, this approach may overlook non-facial characteristics and risks marginalizing individuals with atypical facial features.

Secondly, our method assumes a binary gender model for debiasing purposes, which may inadequately represent individuals who do not align with conventional gender categories, such as those with non-binary identities. This limitation is prevalent in many works focused on fairness and debiasing [35, 8], underscoring an important area for future research that our current study does not yet address.

Finally, our training set debiasing method may lead to an increase in the NSFW rate compared to the original unlearned models when the model is exposed to unsafe prompts. Future work could focus on achieving a better balance between fairness and model safety.

References

[1] Bedapudi, P.: Nudenet: Neural nets for nudity classification, detection and selective censoring. URL https://github. com/notAI-tech/NudeNet (2022)

[2] Bergstra, J., Bardenet, R., Bengio, Y., Kégl, B.: Algorithms for hyper-parameter optimization. Advances in neural information processing systems **24** (2011)

[3] Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al.: Improving image generation with better captions. Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf **2**(3), 8 (2023)

[4] Bianchi, F., Kalluri, P., Durmus, E., Ladhak, F., Cheng, M., Nozza, D., Hashimoto, T., Jurafsky, D., Zou, J., Caliskan, A.: Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. pp. 1493–1504 (2023)

[5] Birhane, A., Prabhu, V.U., Kahembwe, E.: Multimodal datasets: misogyny, pornography, and malignant stereotypes. arXiv preprint arXiv:2110.01963 (2021)

[6] Choi, K., Grover, A., Singh, T., Shu, R., Ermon, S.: Fair generative modeling via weak supervision. In: International Conference on Machine Learning. pp. 1887–1898. PMLR (2020)

[7] Fan, C., Liu, J., Zhang, Y., Wei, D., Wong, E., Liu, S.: Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. arXiv preprint arXiv:2310.12508 (2023)

[8] Friedrich, F., Brack, M., Struppek, L., Hintersdorf, D., Schramowski, P., Luccioni, S., Kersting, K.: Fair diffusion: Instructing text-to-image generation models on fairness. arXiv preprint arXiv:2302.10893 (2023)

[9] Gandikota, R., Materzyńska, J., Fiotto-Kaufman, J., Bau, D.: Erasing concepts from diffusion models. In: Proceedings of the 2023 IEEE International Conference on Computer Vision (2023)

[10] Gandikota, R., Orgad, H., Belinkov, Y., Materzyńska, J., Bau, D.: Unified concept editing in diffusion models. IEEE/CVF Winter Conference on Applications of Computer Vision (2024)

[11] Haspelmath, M.: Indefinite pronouns (1997)

[12] Heng, A., Soh, H.: Selective amnesia: A continual learning approach to forgetting in deep generative models. Advances in Neural Information Processing Systems **36** (2024)

[13] Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)

[14] House, W.: Blueprint for an ai bill of rights. Making Automated Systems Work for the American People (Ocotober 2022)(abrufbar unter https://www. whitehouse. gov/wpcontent/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights. pdf) (2022)

[15] Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. Journal of Global optimization **13**, 455–492 (1998)

[16] Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. Proceedings of the national academy of sciences **114**(13), 3521–3526 (2017)

[17] Koch, K., Soll, M.: No matter how you slice it: Machine unlearning with sisa comes at the expense of minority classes. In: 2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML). pp. 622–637. IEEE (2023)

[18] Letham, B., Karrer, B., Ottoni, G., Bakshy, E.: Constrained bayesian optimization with noisy experiments (2019)

[19] Li, X., Yang, Y., Deng, J., Yan, C., Chen, Y., Ji, X., Xu, W.: Safegen: Mitigating unsafe content generation in text-to-image models. arXiv preprint arXiv:2404.06666 (2024)

[20] Liu, V., Chilton, L.B.: Design guidelines for prompt engineering text-to-image generative models. In: Proceedings of the 2022 CHI conference on human factors in computing systems. pp. 1–23 (2022)

[21] Lyu, M., Yang, Y., Hong, H., Chen, H., Jin, X., He, Y., Xue, H., Han, J., Ding, G.: One-dimensional adapter to rule them all: Concepts diffusion models and erasing applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7559–7568 (2024)

[22] Murakawa, N.: The first civil right: How liberals built prison America. Oxford University Press (2014)

[23] Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)

[24] Oesterling, A., Ma, J., Calmon, F., Lakkaraju, H.: Fair machine unlearning: Data removal while mitigating disparities. In: International Conference on Artificial Intelligence and Statistics. pp. 3736–3744. PMLR (2024)

[25] Qu, Y., Shen, X., He, X., Backes, M., Zannettou, S., Zhang, Y.: Unsafe diffusion: On the generation of unsafe images and hateful memes from text-to-image models. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. pp. 3403–3417 (2023)

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

[27] Rando, J., Paleka, D., Lindner, D., Heim, L., Tramèr, F.: Red-teaming the stable diffusion safety filter. arXiv preprint arXiv:2210.04610 (2022)

[28] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)

[29] Saguy, A.C., Williams, J.A.: A little word that means a lot: A reassessment of singular they in a new era of gender politics. Gender & Society **36**(1), 5–31 (2022)

[30] Schramowski, P., Brack, M., Deiseroth, B., Kersting, K.: Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22522–22531 (2023)

[31] Schramowski, P., Brack, M., Deiseroth, B., Kersting, K.: Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22522–22531 (2023)

[32] Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. Advances in Neural Information Processing Systems **35**, 25278–25294 (2022)

[33] Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., Komatsuzaki, A.: Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv preprint arXiv:2111.02114 (2021)

[34] Seo, S., Lee, J.Y., Han, B.: Information-theoretic bias reduction via causal view of spurious correlation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 2180–2188 (2022)

[35] Shen, X., Du, C., Pang, T., Lin, M., Wong, Y., Kankanhalli, M.: Finetuning text-to-image diffusion models for fairness. arXiv preprint arXiv:2311.07604 (2023)

[36] Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. Advances in neural information processing systems **30** (2017)

[37] Teo, C., Abdollahzadeh, M., Cheung, N.M.M.: On measuring fairness in generative models. Advances in Neural Information Processing Systems **36** (2024)

[38] Teo, C.T., Abdollahzadeh, M., Cheung, N.M.: Fair generative models via transfer learning. In: Proceedings of the AAAI conference on artificial intelligence. vol. 37, pp. 2429–2437 (2023)

[39] Wang, C.L., Huai, M., Wang, D.: Inductive graph unlearning. In: 32nd USENIX Security Symposium (USENIX Security 23). pp. 3205–3222 (2023)

[40] Wu, J., Harandi, M.: Scissorhands: Scrub data influence via connection sensitivity in networks. arXiv preprint arXiv:2401.06187 (2024)

[41] Wu, J., Le, T., Hayat, M., Harandi, M.: Erasediff: Erasing data influence in diffusion models. arXiv preprint arXiv:2401.05779 (2024)

[42] Yang, Y., Hui, B., Yuan, H., Gong, N., Cao, Y.: Sneakyprompt: Jailbreaking text-to-image generative models. arXiv preprint arXiv:2305.12082 (2023)

[43] Zhang, D., Pan, S., Hoang, T., Xing, Z., Staples, M., Xu, X., Yao, L., Lu, Q., Zhu, L.: To be forgotten or to be fair: Unveiling fairness implications of machine unlearning methods. AI and Ethics **4**(1), 83–93 (2024)

[44] Zhang, Y., Jia, J., Chen, X., Chen, A., Zhang, Y., Liu, J., Ding, K., Liu, S.: To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images... for now. arXiv preprint arXiv:2310.11868 (2023)

[45] Zhang, Y., Sohn, K., Villegas, R., Pan, G., Lee, H.: Improving object detection with deep convolutional networks via bayesian optimization and structured prediction. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 249–258 (2015)

A Synthetic Training Data Generation

Based on the paper and official codebases of EraseDiff [41], SalUn [7], and ScissorHands [40], the process of generating synthetic training data follows these common steps:

1. Create a small set of prompts for NSFW (forget set) and safe content (retain set). For NSFW content, EraseDiff and ScissorHands use four prompts: "nudity," "naked," "erotic," and "sexual," while SalUn uses "a photo of a nude person." For safe content, EraseDiff and ScissorHands use the prompt "a person wearing clothes," and SalUn uses "a photo of a person wearing clothes."
2. Generate images using the SD V1.4 model. In EraseDiff and ScissorHands, both the forget and retain sets consist of 400 images, whereas in SalUn, each set contains 800 images.

In our experiments, we adhere to these settings and create both the forget and retain training sets with 400 images each. The prompts we use are: "a nude/naked/erotic/sexual man/woman" for the forget set, and "a man/woman wearing clothes" for the retain set. We adjust the gender ratio of the training data by varying the ratio of "man" and "woman" in the prompts.

B Experimental Details

The 153 occupations we use in the *Occupations* benchmark are: ['accountant', 'aerospace engineer', 'aide', 'air conditioning installer', 'architect', 'artist', 'author', 'baker', 'bartender', 'bus driver', 'butcher', 'career counselor', 'carpenter', 'carpet installer', 'cashier', 'ceo', 'childcare worker', 'civil engineer', 'claims appraiser', 'cleaner', 'clergy', 'clerk', 'coach', 'community manager', 'compliance officer', 'computer programmer', 'computer support specialist', 'computer systems analyst', 'construction worker', 'cook', 'correctional officer', 'courier', 'credit counselor', 'customer service representative', 'data entry keyer', 'dental assistant', 'dental hygienist', 'dentist', 'designer', 'detective', 'director', 'dishwasher', 'dispatcher', 'doctor', 'drywall installer', 'electrical engineer', 'electrician', 'engineer', 'event planner', 'executive assistant', 'facilities manager', 'farmer', 'fast food worker', 'file clerk', 'financial advisor', 'financial analyst', 'financial manager', 'firefighter', 'fitness instructor', 'graphic designer', 'groundskeeper', 'hairdresser', 'head cook', 'health technician', 'host', 'hostess', 'housekeeper', 'industrial engineer', 'insurance agent', 'interior designer', 'interviewer', 'inventory clerk', 'it specialist', 'jailer', 'janitor', 'laboratory technician', 'language pathologist', 'lawyer', 'librarian', 'logistician', 'machinery mechanic', 'machinist', 'maid', 'manager', 'manicurist', 'market research analyst', 'marketing manager', 'massage therapist', 'mechanic', 'mechanical engineer', 'medical records specialist', 'mental health counselor', 'metal worker', 'mover', 'musician', 'network administrator', 'nurse', 'nursing assistant', 'nutritionist', 'occupational therapist', 'office clerk', 'office worker', 'painter', 'paralegal', 'payroll clerk', 'pharmacist', 'pharmacy technician', 'photographer', 'physical therapist', 'pilot', 'plane mechanic', 'plumber', 'police officer', 'postal worker', 'printing press operator', 'producer', 'programmer', 'psychologist', 'public relations specialist', 'purchasing agent', 'radiologic technician', 'real estate broker', 'receptionist', 'repair worker', 'roofer', 'sales manager', 'salesperson', 'school bus driver', 'scientist', 'security guard', 'sheet metal worker', 'singer', 'social assistant', 'social worker', 'software developer', 'stocker', 'supervisor', 'taxi driver', 'teacher', 'teaching assistant', 'teller', 'therapist', 'tractor operator', 'truck driver', 'tutor', 'underwriter', 'veterinarian', 'waiter', 'waitress', 'web developer', 'welder', 'wholesale buyer', 'writer']

Table 3: Classification of NSFW versus Safe content is based on NudeNet categories of exposed and covered body parts.

NSFW	"BUTTOCKS EXPOSED", "FEMALE BREAST EXPOSED", "FEMALE GENITALIA EXPOSED", "MALE BREAST EXPOSED", "ANUS EXPOSED", "FEET EXPOSED", "ARMPITS EXPOSED", "BELLY EXPOSED", "MALE GENITALIA EXPOSED"
Safe	"FEMALE GENITALIA COVERED", "FACE FEMALE", "BELLY COVERED", "FEET COVERED", "ARMPITS COVERED", "FACE MALE", "ANUS COVERED", "FEMALE BREAST COVERED", "BUTTOCKS COVERED"

C NSFW: Not-safe-for-work Content

In our evaluation, NSFW content refers specifically to images classified as "exposed" by the NudeNet classifier [1], using a probability threshold of 0.6 for nudity detection. The specific categories distinguishing NSFW from safe content are provided in Table 3.