
TOWARD ROBUST IMPERCEPTIBLE PERTURBATION AGAINST UNAUTHORIZED TEXT-TO-IMAGE DIFFUSION-BASED SYNTHESIS

A PREPRINT

Yixin Liu*

Lehigh University
yila22@lehigh.edu

Chenrui Fan

Huazhong University of Science and Technology
fancr0422@gmail.com

Yutong Dai

Lehigh University
yud319@lehigh.edu

Xun Chen

Samsung Research American
Xun.chen@samsung.com

Pan Zhou

Huazhong University of Science and Technology
panzhou@hust.edu.cn

Lichao Sun

Lehigh University
lis221@lehigh.edu

ABSTRACT

Text-to-image diffusion models allow seamless generation of personalized images from scant reference photos. Yet, these tools, in the wrong hands, can fabricate misleading or harmful content, endangering individuals. To address this problem, existing poisoning-based approaches perturb user images in an imperceptible way to render them "unlearnable" from malicious uses. We identify two limitations of these defending approaches: i) sub-optimal due to the hand-crafted heuristics for solving the intractable bilevel optimization and ii) lack of robustness against simple data transformations like Gaussian filtering. To solve these challenges, we propose MetaCloak, which solves the bi-level poisoning problem with a meta-learning framework with an additional transformation sampling process to craft transferable and robust perturbation. Specifically, we employ a pool of surrogate diffusion models to craft transferable and model-agnostic perturbation. Furthermore, by incorporating an additional transformation process, we design a simple denoising-error maximization loss that is sufficient for causing transformation-robust semantic distortion and degradation in a personalized generation. Extensive experiments on the VG-*G*Face2 and CelebA-HQ datasets show that MetaCloak outperforms existing approaches. Notably, MetaCloak can successfully fool online training services like Replicate, in a black-box manner, demonstrating the effectiveness of MetaCloak in real-world scenarios. Our code is available at <https://github.com/liuyixin-louis/MetaCloak>.

1 Introduction

Diffusion models achieve significant success in a wide range of applications, including image generation [15; 47; 7], image editing [21; 46; 5], and text-to-image synthesis [38]. Subject-driven text-to-image synthesis, an emerging application of diffusion models, in particular, has attracted considerable attention due to its potential to generate personalized images from a few reference photos. Among the approaches proposed to achieve this goal [34; 42], DreamBooth [40] and Text Inversion [12] are two prominent training-based methods that offer impressive personalize generation ability. With an additional lightweight model fine-tuning process for capturing the subject or a concept-related embedding training phase, personalized diffusion models can retain the generation capacity from the pre-training stage and conduct vivid personalized generation. While these methods empower high-quality personalized generation, they also raise privacy concerns as they can fabricate misleading or harmful content in the wrong hands, endangering individuals. For example, recent news [19] indicates that AI tools like diffusion models have been employed to generate fake profiles of individuals for launching a new wave of fraud.

To tackle these issues, some poisoning-based approaches [24; 26] have been recently proposed to perturb user images in an imperceptible way to render them "unlearnable" from malicious uses. Specifically, these approaches aim to craft perturbations that can mislead the personalized diffusion model generation process to compromise the generation ability. For instance, PhotoGuard [43] proposes to attack the VAE encoder or UNet decoder for crafting perturbation that aligns the perturbed latent code or denoised image to the ones of a dummy target. Similarly,

*This work was done during Yixin's internship at Samsung Research American.

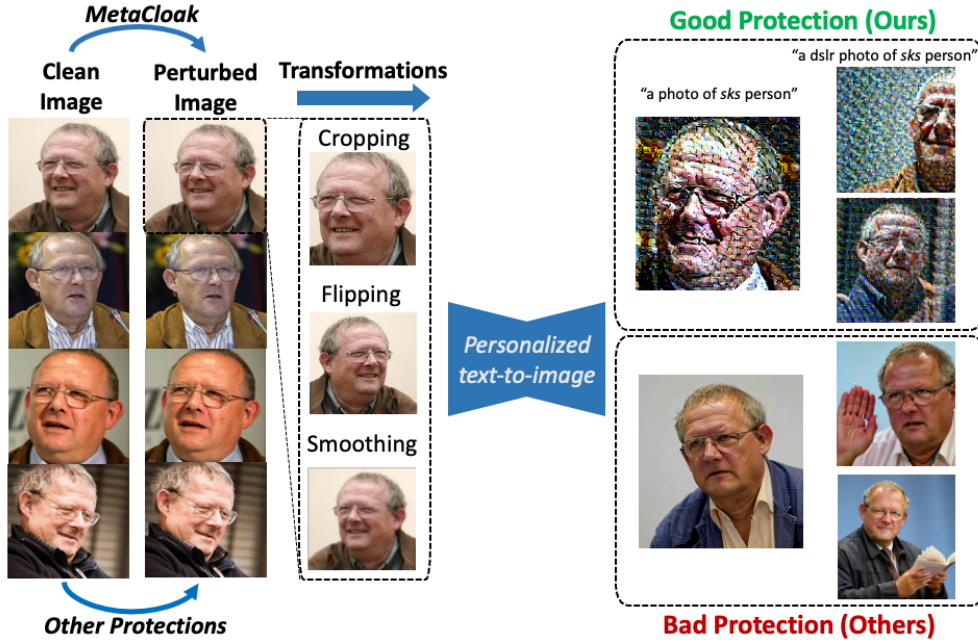


Figure 1: Image protected by existing methods fails to fool personalized text-to-image approaches after applying data transformations. In contrast, our MetaCloak is still robust in such adversity.

Liang et al. [26] proposes AdvDM to craft protected images with a pre-trained fixed surrogate model with adversarial perturbation. Its latter version, Mist [25], considers an additional targeting loss for degrading the texture quality. Targeting art style transferring, Glaze [45] proposes to minimize the feature distance between the perturbed image and a targeting image with a pre-trained feature extractor. However, these works all focus on attacking text-to-image synthesis methods that leverage *fixed* diffusion models, where protection is much easier due to the inevitable adversarial vulnerability of the DNNs.

In contrast, fine-tuning-based personalized generation approaches, represented by DreamBooth, ensure high-quality synthesis via *actively learning* the concept of the subject by optimizing either the model parameters or a concept-related embedding. Compared to attacking *fixed diffusion model*, success in attacking these fine-tuning-based generation methods is rooted in different mechanisms, i.e., tricking the models into establishing false correlations or overfitting. To the best of our knowledge, Anti-Dreambooth [24] is the only existing work focused on this challenging setting. Although this approach has shown effectiveness in preventing unauthorized subject-driven text-to-image synthesis, it exhibits two limitations. Firstly, it uses hand-crafted heuristics to address the underlying poisoning problem, which is a challenging bilevel optimization problem, yielding sub-optimal performance. Specifically, it incorporates an additional clean set training process in the surrogate model training trajectory, which is mismatched with the actual ones trained on the poison and thus degrades the poisoning capacity. Secondly, recent studies found that these data protections are fragile and demonstrate limited robustness against minor data transformations such as filtering (see Fig. 1). Given these limitations, in this work, we ask the following question: *Can we design a more effective and robust data protection scheme that can prevent unauthorized subject-driven text-to-image diffusion-based synthesis under data transformation?*

To answer this question, we propose MetaCloak, a more effective and robust poisoning-based data protection approach against unauthorized subject-driven text-to-image diffusion-based synthesis. To tackle the first challenge, as well as to improve perturbation transferability, MetaCloak leverages meta-learning to learn model-agnostic perturbation over a pool of steps-staggered surrogates. With this meta-learning framework, we found that the perturbations are more effective than the previous approaches under the training setup without data transformations. To address the second challenge, i.e., improving the robustness of perturbation under the training setting with data transformation, we incorporate a transformation sampling process to craft transformation-robust perturbation. Besides, designing generation loss to guide poisoning is tricky since leveraging the “ground-truth” metrics might lead to an over-fitting problem. To address this problem, we design a denoising-error maximization loss to encourage the added perturbation to cause serious semantic distortion from the model perception, which aims to introduce “meaningless” and “hard-to-understand” patterns that can trick the model into learning false correlations. On the

VGGFace2 and CelebA-HQ datasets, we show that our method significantly outperforms existing approaches under data transformations. Notably, MetaCloak can successfully fool an online personalized diffusion model training service, Replicate, in a black-box manner, demonstrating the effectiveness of MetaCloak in real-world scenarios. Our main contributions are summarized as follows.

1. We propose MetaCloak, the first robust poisoning-based approach for crafting more effective perturbation that can further bypass data transformation defenses.
2. To solve the intractable bi-level poisoning, we propose to use meta-learning to learn model-agnostic and transferable perturbation by leveraging a pool of surrogates.
3. To guide the poisoning and bypass transformation defenses, we design a simple denoising-error maximization loss with an additional transformation process for transformation-robust semantic distortion.
4. Extensive experiments on the VGGFace2 and CelebA-HQ datasets show the superiority of our method compared to existing approaches under settings with and without data transformations. Notably, our method can fool online training services like Replicate in a black-box manner.

2 Related Works

Availability Attack. Availability attack aims to degrade data availability for model training by injecting imperceptible perturbations into the training data. The underlying mechanism of this attack is rooted in the shortcut learning preference of deep neural networks (DNNs), where models turn to learn easy but brittle features for prediction. Existing works on availability attacks mainly focus on image classification [16; 17; 9; 53]. Unlearnable Example [16] and L2C [9] are two of the pioneers in this field, and later research Tao et al. [48] indicates that noise can be easily removed with adversarial training and data transformations. To address this concern, many works have been proposed, with techniques of min-max optimization [11], filter-based perturbation [41], and score guiding [8]. To better understand the defense mechanism, Tao et al. [48] propose frameworks for showing some certificated purification against such attacks. To adapt to other models and applications, studies are also conducted in natural language processing [18], graph learning [27], and contrastive learning [35]. To the best of our knowledge, we are the first to study robust availability attacks to personalized generation with diffusion models.

Protection Against Unauthorized Subject-driven AI Synthesis. Unauthorized subject-driven AI synthesis including style transfer [21], personalized generation [40], and image inpainting [2], pose a serious threat to the privacy of individuals and art copyright [4; 50]. To address this concern, recent works [43; 55; 45; 24] propose to protect the data with poisoning-based approaches, which modify the data in an imperceptible way while causing severe degradation in the generation performance. PhotoGuard [43] proposes to attack the VAE encoder or UNet decoder for crafting perturbation that aligns the perturbed latent code or denoised image to the ones of the dummy target. Targeting degrading art style transferring, Glaze [45] proposes to minimize similar targeting loss with a pre-trained style-transfer model, with learned perceptual similarity as a penalty. Similarly, AdvDM [55] proposes to craft perturbation that minimizes the likelihood of the perturbed image with a pre-trained diffusion model, which equivalently maximizes the denoising loss during training. Its later version, Mist [25], further incorporates a texture targeting loss for a more robust and sharp pattern into the perturbed images. However, these works all focus on attacking the synthesis methods that leverage *fixed* diffusion models by exploiting the adversarial vulnerability of the DNNs. In contrast, our work focuses on conducting more challenging data protection against fine-tuning-based synthesis methods [40; 33]. To the best of our knowledge, Anti-DreamBooth [24] is the only work that studies data protection against fine-tuning-based synthesis approaches. However, it crafts sub-optimal perturbation and is also not robust against transformation purification. In our work, targeting fine-tuning-based synthesis methods, we study poisoning-based data protection with better effectiveness and robustness against transformations.

3 Preliminary

Text-to-Image Diffusion Models. Diffusion models are probabilistic generative models that are trained to learn a data distribution by the gradual denoising of a variable sampled from a Gaussian distribution. Our specific interest lies in a pre-trained text-to-image diffusion model denoted as $\hat{\mathbf{x}}_\theta$. This model operates by taking an initial noise map $\boldsymbol{\epsilon}$ sampled from a standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and a conditioning vector \mathbf{c} . This conditioning vector \mathbf{c} is generated through a series of steps involving a text encoder represented as Γ , a text tokenizer denoted as f , and a text prompt \mathbf{P} (i.e., $\mathbf{c} = \Gamma(f(\mathbf{P}))$). The ultimate output of this model is an image denoted as \mathbf{x}_{gen} , which is produced as a result of the operation $\mathbf{x}_{\text{gen}} = \hat{\mathbf{x}}_\theta(\boldsymbol{\epsilon}, \mathbf{c})$. They are trained using a squared error loss to denoise a variably-noised image as follows:

$$\mathcal{L}_{\text{denoise}}(\mathbf{x}, \mathbf{c}; \theta) = \mathbb{E}_{\boldsymbol{\epsilon}, t} [w_t \|\hat{\mathbf{x}}_\theta(\alpha_t \mathbf{x} + \sigma_t \boldsymbol{\epsilon}, \mathbf{c}) - \mathbf{x}\|_2^2], \quad (1)$$

where \mathbf{x} is the ground-truth image, \mathbf{c} is a conditioning vector (e.g., obtained from a text prompt), and α_t, σ_t, w_t are terms that control the noise schedule and sample quality, and are functions of the diffusion process time t .

Adversarial Attacks to Text-to-Image Diffusion Models. Adversarial attacks aim to perform an imperceptible perturbation on the input image in order to mislead machine learning models’ predictions. In the classification scenario, for a given classifier f_{cls} , a perturbed adversarial image \mathbf{x}' is generated from the original image \mathbf{x} to misguide the model into incorrect classification. Constraints on the perceptibility of changes are often imposed through ℓ_p norms (with $p \geq 1$), such that the perturbed image \mathbf{x}' is bounded within a ℓ_p -ball centered at \mathbf{x} with radius $r > 0$, i.e., $\mathbf{x}' \in B_p(\mathbf{x}, r) = \{\mathbf{x}' : \|\mathbf{x}' - \mathbf{x}\|_p \leq r\}$. Given a classification loss \mathcal{L}_{cls} , untargeted adversarial examples are crafted by solving $\max_{\mathbf{x}' \in B_p(\mathbf{x}, r)} \mathcal{L}_{\text{cls}}(f_{\text{cls}}(\mathbf{x}'), y_{\text{true}})$, where y_{true} is the true label of image \mathbf{x} . For the text-to-image generation scenario, given a pre-trained text-to-image diffusion model $\hat{\mathbf{x}}_\theta$, the adversarial attack aims to perturb the image to hinder the model from reconstructing the image, i.e., $\mathbf{x}' \leftarrow \arg\max_{\mathbf{x}' \in B_p(\mathbf{x}, r)} \mathcal{L}_{\text{denoise}}(\mathbf{x}', \mathbf{c}; \theta)$. In this paper, we consider the ℓ_∞ -norm for its alignment with perception [13]. To solve this constrained optimization, the Projected Gradient Descent (PGD) [29] technique is commonly utilized by iteratively updating the poisoned image \mathbf{x}' . Formally, the adversarial example \mathbf{x}' is updated as

$$\mathbf{x}'_i = \Pi_{B_\infty(\mathbf{x}, r)}(\mathbf{x}'_{i-1} + \alpha \text{sign}(\nabla_{\mathbf{x}'_{i-1}} \mathcal{L}_{\text{denoise}})), \quad (2)$$

where $\mathbf{x}'_0 = \mathbf{x}$, $\text{sign}(\cdot)$ is the sign function, i is the step index, and the step size $\alpha > 0$.

During this generation process, the adversarial examples gradually progress in a direction that would increase the denoising loss while maintaining imperceptible perturbations. Recent works [24; 25] have demonstrated that perturbed images from this attack can effectively deceive text-to-image generation models [12; 40] to produce low-quality images.

Personalized Diffusion via DreamBooth Fine-tuning. DreamBooth is a method aimed at personalizing text-to-image diffusion models for specific instances. It has two main objectives: first, to train the model to generate images of the given subject with generic prompts like “a photo of sks [class noun]”, where sks specifies the subject and “[class noun]” is the category of object (e.g., “person”). For this, it uses the loss defined in Eq. 1 with x_{it} as the user’s reference image and conditioning vector $\mathbf{c} := \Gamma(f(\text{“a photo of } sks \text{ [class noun]”}))$. Similar to the classification model, this guides the model to create the correlation between the identifier and the subject. Secondly, it introduces a class-specific prior-preserving loss to mitigate overfitting and language-drifting issues. Specifically, it retains the prior by supervising the model with *its own generated samples* during the fine-tuning stage. With a class-specific conditioning vector $\mathbf{c}_{\text{pr}} := \Gamma(f(\text{“photo of a [class noun]”}))$ and random initial noise $\mathbf{z}_{t_1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, DreamBooth first generates prior data $\mathbf{x}_{\text{pr}} = \hat{\mathbf{x}}_{\theta_0}(\mathbf{z}_{t_1}, \mathbf{c}_{\text{pr}})$ using the pre-trained diffusion model and then minimize:

$$\begin{aligned} \mathcal{L}_{\text{db}}(\mathbf{x}, \mathbf{c}; \theta) = & \mathbb{E}_{\epsilon, \epsilon', t} [w_t \|\hat{\mathbf{x}}_\theta(\alpha_t \mathbf{x} + \sigma_t \epsilon, \mathbf{c}) - \mathbf{x}\|_2^2 + \\ & \lambda w_{t'} \|\hat{\mathbf{x}}_\theta(\alpha_{t'} \mathbf{x}_{\text{pr}} + \sigma_{t'} \epsilon', \mathbf{c}_{\text{pr}}) - \mathbf{x}_{\text{pr}}\|_2^2], \end{aligned} \quad (3)$$

where ϵ, ϵ' are both sampled from $\mathcal{N}(0, \mathbf{I})$, the second term is the prior-preservation term that supervises the model with its own generated images and λ controls for the relative importance of this term. With approximately one thousand training steps and four subject images, it can generate vivid personalized subject images with Stable Diffusion [49].

4 Problem Statement

We formulate the problem as follows. A user (*image protector*) wants to protect his images $X_c = \{\mathbf{x}_i\}_{i=1}^n$ from being used by an *unauthorized model trainer* for generating personalized images using DreamBooth, where n is the number of images. To achieve this, for some portion of images $\mathbf{x} \in X_c$, the user injects a small perturbation onto the original image to craft poisoned images set $X_p = \{\mathbf{x}'_i\}_{i=1}^n$, which is then published to the public. Later, the model trainers will collect and use X_p to finetune a text-to-image generator $\hat{\mathbf{x}}_\theta$, following the DreamBooth algorithm, to get the optimal parameters θ^* . We assume that the model trainer is aware of the poisoning to some extent, so some data transformations like filtering or cropping might be applied to the training image set X_p during the data pre-processing phase of DreamBooth training. The objective of the user is to craft a delusive and robust image set X_p to degrade the DreamBooth’s personalized generation ability, which can be formulated as:

$$X_p^* \in \arg\max_{X_p, \theta^*} \mathcal{L}_{\text{gen}}^*(X_{\text{ref}}; \hat{\mathbf{x}}_{\theta^*}, X_p) \quad (4)$$

$$\begin{aligned} \text{s.t. } \theta^* \in & \arg\min_{\theta} \{\mathcal{L}_{\text{db}}^{\text{rob}}(X_p, T; \theta) := \\ & \mathbb{E}_{\mathbf{x}'_i \sim X_p, g \sim T} \mathcal{L}_{\text{db}}(g(\mathbf{x}'_i), \mathbf{c}; \theta)\}. \end{aligned} \quad (5)$$

Here, \mathbf{c} is the class-wise conditional vector, $\mathcal{L}_{\text{gen}}^*$ is some perception-aligned loss to measure the personalization generation ability of trained model $\hat{\mathbf{x}}_{\theta^*}$ (with more details in the next section), T is a set of data transformations the expected adversary might use, X_{ref} is a clean reference set, and \mathbf{c} is the conditioning vector. Compared to vanilla \mathcal{L}_{db} in (3), $\mathcal{L}_{\text{db}}^{\text{rob}}$ ² is more robust to learning personalized diffusion models.

Overall Goals. While it’s hard to quantify a unified evaluation loss $\mathcal{L}_{\text{gen}}^*$ to measure the personalized generation quality, our overall goal is to degrade the usability of generated images, and we attempt to decompose the evaluation metric into the following two aspects: *quality-related and semantic-related distortion*. Specifically, we seek to render the generated image awful quality by tricking the victim’s model into generating an image with some artifacts. With this distortion, the model trainer can’t use those images for some quality-sensitive applications. Furthermore, the subject identity of generated images should be greatly distorted for other’s utilization. We’ll dive into the design of $\mathcal{L}_{\text{gen}}^*$ in Sec. 5.2.

5 Method

5.1 Learning to Learn Transferable and Model-agnostic Perturbation

One naive idea to solve the bilevel problem (4)-(5) is to unroll all the training steps and optimize the protected examples X_p via backpropagating. However, accurately minimizing this full bi-level objective is intractable since a computation graph that explicitly unrolls 10^3 SGD steps would not fit on most of the current machines. To address this issue, inspired by Huang et al. [17], we propose to approximately optimize the upper-level objective (4) and lower-level objective (5) in an alternative fashion. Specifically, considering the i -th iteration, when the current model weight θ_i and the protected image set X_p^i are available (with θ_0 is initialized from pre-trained diffusion model and $X_p^0 = X_c$), we make a copy of current model weight $\theta'_{i,0} \leftarrow \theta_i$ for noise crafting and optimize the lower-level problem for K steps as:

$$\theta'_{i,j+1} = \theta'_{i,j} - \beta \nabla_{\theta'_{i,j}} \mathcal{L}_{\text{db}}^{\text{rob}}(X_p^i; \theta'_{i,j}), \quad (6)$$

where $j \in \{0, 1, \dots, K-1\}$ and $\beta > 0$ is the stepsize. We term this procedure K -step method. This unrolling procedure allows us to “look ahead” in training and view how the perturbations *now* will impact the generation loss $\mathcal{L}_{\text{gen}}^*$ *after* K steps. We then leverage the unrolled model $\hat{\mathbf{x}}_{\theta'_{i,K}}$ for optimizing the upper-level problem, i.e., updating the protected images X_p as:

$$X_p^{i+1} = \Pi_{B_{\infty}(X_p^0, r)}(X_p^i + \alpha \text{sign}(\nabla_{X_p^i} \mathcal{L}_{\text{gen}}^*(X_{\text{ref}}; \hat{\mathbf{x}}_{\theta'_{i,K}}, X_p^i)). \quad (7)$$

After obtaining the updated protected images X_p^{i+1} , the surrogate model θ_i is trained with $\mathcal{L}_{\text{db}}^{\text{rob}}$ for a few SGD steps on X_p^{i+1} to get θ_{i+1} as the next iteration’s starting point

$$\theta_{i+1} = \theta_i - \beta \nabla_{\theta_i} \mathcal{L}_{\text{db}}^{\text{rob}}(X_p^{i+1}; \theta_i). \quad (8)$$

The procedure (6)-(8) is executed repeatedly until the surrogate model reaches maximum training steps to obtain the final protected images X_p^* . While this K -step method offers satisfactory results, it is not robust under various training settings with different models and initialization. Since training with a single surrogate model will turn to overfit on a single training trajectory [17]. To craft transferable perturbations, we treat the injected perturbations as a training “hyperparameter”, and use meta-learning to learn model-agnostic perturbations that degrade the performance of trained models. Different from the conventional meta-learning setting, whose goal is to transfer across datasets and tasks, in our poisoning problem, we aim to craft perturbations that transfer across different models and different training trajectories. To achieve this, we propose to learn perturbation over a pool of steps-staggered surrogates. Specifically, we train M surrogate models $\{\theta^j\}_{j=1}^M$ as the initial point for the j -th surrogate model. Given maximum training steps N_{max} , the j -th surrogate model θ^j is trained with $\mathcal{L}_{\text{db}}^{\text{rob}}$ for $\lfloor j N_{\text{max}} / M \rfloor$ steps *on clean data* from the pre-trained weight θ_0 . With this pool, we then conduct C crafting outer loop; for each loop, we sample a batch of surrogates and train them separately on the poison. Upon the completion of inner-loop training, we average the generation losses $\mathcal{L}_{\text{gen}}^*$ over ensembles and update the protected images X_p with gradient ascent. Following MAML [10], a first-order gradient approximation is used for efficient perturbation crafting.

5.2 Transformation-robust Semantic Distortion with Denoising-error Maximization

During the evaluation stage of the generated images, we can readily leverage various quality reference-based and reference-free assessment metrics like CLIP-IQA [52] and FDSR [14] for the construction of the ground-truth

²We by default omit T and simplify the notation as $\mathcal{L}_{\text{db}}^{\text{rob}}(X_p; \theta)$ in the following context.

Algorithm 1 Training robust perturbation with MetaCloak

Input: User training image set X_c , training iteration N_{\max} , the number of surrogate models M , PGD radius r , PGD step size α , the unrolling number K , the transformation distribution T , the sampling times J , learning rate β .

Output: Protected user image set X_p

- 1: Stagger the M models as a pool \mathcal{D}_θ , training the m -th model weights θ_m up to $\lfloor mN_{\max}/M \rfloor$ steps on X_c
- 2: **for** i **in** $1, \dots, C$ crafting steps **do** ▷ Outer-loop
- 3: Sample a batch of surrogate $\{\theta_m\}_{m=1}^B, \theta_m \sim \mathcal{D}_\theta$
- 4: **for** m **in** $1, \dots, B$ **do** ▷ Inner-loop
- 5: Copy model weight $\tilde{\theta} = \theta_m$
- 6: **for** $k = 1, \dots, K$ unroll steps **do**
- 7: $\tilde{\theta} = \tilde{\theta} - \beta \nabla_{\tilde{\theta}} \mathcal{L}_{\text{db}}^{\text{rob}}(X_p, T; \tilde{\theta})$
- 8: **end for**
- 9: **for** j **in** $1, \dots, J$ **do**
- 10: $\mathcal{L}_m^j = \mathcal{L}_{\text{gen}}(t_j(X_p); \tilde{\theta}), t_j \sim T$
- 11: **end for**
- 12: $\theta_m = \theta_m - \alpha \nabla_{\theta_m} \mathcal{L}_{\text{db}}^{\text{rob}}(X_p, T; \theta_m)$
- 13: **end for**
- 14: Average denoise losses $\mathcal{L}_{\text{adv}} = \sum_{m,j} \mathcal{L}_m^j / BJ$
- 15: Compute $\nabla_{X_p} \mathcal{L}_{\text{adv}}$
- 16: Update X_p using SGD and project onto r ball
- 17: **end for**
- 18: **return** X_p

generation loss $\mathcal{L}_{\text{gen}}^*$. However, during the poisoning stage, we can not simply take these “ground-truth” metric losses to serve as the loss for crafting noise: i) overfitting is prone to happen since most quality-assessment models are neural-network-based; ii) even if the metrics are rule-based, the leading distortion might over-adapt to certain assessment models. To avoid these problems, we take a different way of designing an approximated generation loss $\mathcal{L}_{\text{gen}}(X_p; \theta) \in \mathbb{R}^+$ used for crafting poison. Our design of \mathcal{L}_{gen} is motivated by the insight that introducing some “meaningless” and “hard-to-understand” patterns can trick the diffusion model into overfitting on the perturbations, leading the diffusion model to establish false correlations. Specifically, our approximated generation loss can be formulated as,

$$\mathcal{L}_{\text{gen}}(X_p; \theta) = \mathbb{E}_{\mathbf{c}, \mathbf{x}' \sim X_p} [\mathcal{L}_{\text{denoise}}(\mathbf{x}', \mathbf{c}; \theta)]. \quad (9)$$

Our empirical observation indicates that the maximization of this loss can result in chaotic content and scattered texture in the generated images. Instead of introducing an additional targeting loss term [26; 25], we found that our simple denoising-maximization loss is more effective against fine-tuning-based diffusion models. However, perturbations crafted directly with (9) are fragile to minor transformations and ineffective in bypassing (5). To remedy this, we adopt the *expectation over transformation* technique (EOT [1]) into the PGD process. Specifically, given T as a distribution over a set of transformations in (5), we apply EOT on (7) as

$$X_p^{i+1} = \mathbb{E}_{g \sim T} \left[\Pi_{B_{\infty}(X_p^i, r)} (X_p^i + \alpha \text{sign}(\nabla_{X_p^i} \mathcal{L}_{\text{gen}}(g(X_p^i); \hat{\mathbf{x}}_{\theta', i, K}))) \right], \quad (10)$$

where $g(X_p) = \{g(x_p) : x_p \in X_p\}$ is the transformed image of X_p under the transformation g , θ' is a K -step unrolled model following (6), and the expectation is estimated by Monte Carlo sampling with J samples ($J = 1$ in our setup). In summary, we present the overall framework in Alg. 1.

6 Experiments

6.1 Setup

Datasets. Our experiments are performed on human subjects using the two face datasets: CelebA-HQ [20] and VGGFace2 [3] following Anti-DreamBooth [24]. CelebA-HQ is an enhanced version of the original CelebA dataset consisting of 30,000 celebrity face images. VGGFace2 is a comprehensive dataset with over 3.3 million face images from 9,131 unique identities. Fifty identities are selected from each dataset, and we randomly pick 8 images from each individual and split those images into two subsets for image protection and reference. More results on non-face data are in the App. C.3.

Dataset	Method	SDS ↓	IMS_{CLIP} ↓	IMS_{VGG} ↓	CLIP-IQAC ↓	LIQE ↓
VGGFace2	Clean	0.903 ± 0.291	0.790 ± 0.076	0.435 ± 0.657	0.329 ± 0.354	0.984 ± 0.124
	ASPL	0.879 ± 0.321	0.712 ± 0.082	0.193 ± 0.792	-0.090 ± 0.432	0.891 ± 0.312
	EASPL	0.801 ± 0.395	0.703 ± 0.076	0.142 ± 0.814	-0.024 ± 0.406	0.820 ± 0.384
	FSMG	0.911 ± 0.280	0.718 ± 0.063	0.273 ± 0.753	-0.111 ± 0.371	0.891 ± 0.312
	AdvDM	0.903 ± 0.291	0.769 ± 0.068	0.570 ± 0.510	0.241 ± 0.341	0.984 ± 0.124
	Glaze	0.910 ± 0.279	0.774 ± 0.083	0.490 ± 0.607	0.258 ± 0.330	0.984 ± 0.124
	PhotoGuard	0.928 ± 0.255	0.793 ± 0.072	0.524 ± 0.558	0.407 ± 0.249	1.000 ± 0.000
	MetaCloak	0.432* ± 0.489	0.644* ± 0.107	-0.151* ± 0.864	-0.440* ± 0.223	0.570* ± 0.495
CelebA-HQ	Clean	0.810 ± 0.389	0.763 ± 0.119	0.181 ± 0.784	0.470 ± 0.264	0.984 ± 0.124
	ASPL	0.755 ± 0.427	0.714 ± 0.082	0.044 ± 0.823	0.054 ± 0.364	0.969 ± 0.174
	EASPL	0.841 ± 0.362	0.707 ± 0.069	-0.113 ± 0.834	0.020 ± 0.423	0.984 ± 0.124
	FSMG	0.769 ± 0.416	0.718 ± 0.081	0.073 ± 0.805	0.085 ± 0.343	0.953 ± 0.211
	AdvDM	0.866 ± 0.339	0.789 ± 0.083	0.431 ± 0.618	0.373 ± 0.268	1.000 ± 0.000
	Glaze	0.889 ± 0.312	0.778 ± 0.092	0.318 ± 0.716	0.338 ± 0.311	0.992 ± 0.088
	PhotoGuard	0.879 ± 0.320	0.790 ± 0.092	0.401 ± 0.658	0.451 ± 0.277	0.992 ± 0.088
	MetaCloak	0.305* ± 0.452	0.608* ± 0.109	-0.637* ± 0.686	-0.354* ± 0.244	0.438* ± 0.496

Table 1: Results of different methods under the transformations training setting with the corresponding mean and standard deviation (\pm) on two datasets. The best data performances are in **bold**, and the second runners are shaded in gray. The Gaussian filtering kernel size is set to 7. * denotes a significant improvement according to the Wilcoxon signed-rank significance test ($p \leq 0.01$).

Training Settings. The Stable Diffusion (SD) v2-1-base [37] is used as the model backbone by default. For Dreambooth training, we fine-tune both the text-encoder and U-Net model with a learning rate of 5×10^{-7} and batch size of 2 for 1000 iterations in mixed-precision training mode. We consider two training settings: *standard training* (*Stand. Training*) and *training with data transformations* (*Trans. Training*). For the standard training setting, DreamBooth is trained without performing special pre-processing. For the training with data transformations scenario, we consider transformations including Gaussian filtering with a kernel size of 7, horizontal flipping with half probability, center cropping, and image resizing to 512x512. For both the settings, we leverage two inferring prompts with 50 inferring steps, “a photo of sks person” and “a DSLR portrait of sks person” during inference to generate 16 images per prompt.

Baselines and Implementation Details. We compare our method with the following adopted state-of-the-art baselines in [26; 24; 25]: i) *ASPL* [24] alternatively update the perturbations and surrogate models, where the surrogate models are updated on both poisons and clean data; ii) *E-ASPL* is an extension of ASPL that ensembles multiple types of diffusion models for better transferability; iii) *FSMG* leverages a DreamBooth trained on clean image for crafting adversarial examples; iv) *AdvDM* [26; 25] leverages a pre-trained diffusion model for crafting adversarial examples with additional targeting loss for texture distortion; v) *Glaze* [45] minimizes the representation between the perturbed image and a target dummy image with a pre-trained encoder; vi) *PhotoGuard* [43] perturb image to align its denoised image closer to a dummy target with efficient partial back-propagation; Following the common setup [24], we set the noise radius (ℓ_∞ -norm ball) to 11/255 with a step size of 1/255 and a step number of 6 by default. We set the unrolling number $K = 1$, surrogate model number $M = 5$, sample batch size $B = 1$, and crafting step $C = 4000$ with SD v2-1-base as the surrogate model, which takes about 3 GPU hours to train instance-wise noise. See App. A for more details.

Metrics. We evaluate the generated images in terms of their semantic-related quality and graphical quality. For the semantic-related score, first, we want to determine whether the subject is present in the generated image. We term this score as *Subject Detection Score* (*SDS*). For human faces, we can take the mean of face detection confidence probability using RetinaFace detector [6] as its SDS. Secondly, we are interested in how the generated image is semantically close to its subject. We term this score as *Identity Matching Score* (*IMS*) [24], the similarity between embedding of generated face images and an average of all reference images. We use VGG-Face [44] and CLIP-ViT-base-32 [32] as embedding extractors and employ the cosine similarity. Besides, we use *LIQE* [54], a vision-language multitask-learning image quality assessment model, for human scene category prediction. For the graphical quality, we design *CLIP-IQAC*, which is based on CLIP-IQA [52] by considering additional class information. See the App. B.1 for details.




Figure 2: Visualization of transformation robustness of different methods. The first row is a generated sample from DreamBooth trained on poisons with no transformation defenses. The 2-th row showcases the robustness of each method under transformation with a Gaussian kernel size of 7. Our method performs robustly under transformation defenses, while other methods fail to preserve the perturbation.




Figure 3: Results under online training-as-service settings with the Full and LoRA DreamBooth fine-tuning settings on the Replicate.

6.2 Effectiveness of MetaCloak

Effectiveness Comparison through Quantitative and Qualitative Metrics. As observed in Tab. 1, MetaCloak consistently outperforms other baselines across all the metrics. Specifically, in the most important metric, SDS, which measures whether a face appeared in the generated image, MetaCloak successfully degraded this metric by 36.9% and 45.0% compared to previous SOTA on VGGFace2 and CelebA-HQ. Regarding reference-based semantic matching metrics, the results on IMS-VGG and IMS-CLIP also show that our method is more effective than other baselines. Regarding image quality metrics, the results on CLIP-IQAC suggest that MetaCloak can effectively degrade the image quality of generated images. Furthermore, the results in Tab. 5 (see appendix) also suggest that MetaCloak also achieves better effectiveness than existing methods under Stand. Training settings. For visualization, as we can see in Fig. 2, compared to other baselines, MetaCloak can robustly fool the model to generate images with low quality and semantic distortion under both Stand. and Trans. training while others are sensitive to transformations. More visualizations are in App C.2.

Effectiveness under Online Training-as-services Scenarios. To test the effectiveness of our framework in the wild, we conduct experiments under online training-as-service settings. Unlike local training, attacking online training services is more challenging due to the limited knowledge of data preprocessing. We first showcase the performance of our method under two common DreamBooth fine-tuning scenarios, including full fine-tuning (Full-FT) and




Figure 4: Results of MetaCloak with different protection radii and ratios. The dashed line marks the previous SOTA results under $r = 32/255$ and the protection ratio of 100%, respectively.



Table 2: Generated images of Stable Diffusion fine-tuned with Textual Inversion on a randomly selected instance from VGGFace2. S^* denotes a learned token initialized from token $\langle person \rangle$.

LoRA-fine-tuning (LoRA-FT). We sample data from VGGFace2 and upload its clean and poisoned images to Replicate [36] for DreamBooth training. From the results in Fig. 3, we can see that MetaCloak performs well under the Full-FT setting; for instance, it successfully degrades the SDS from 98.9% to 21.8%. Under the LoRA-FT setting, the model seems to be more resilient, but MetaCloak can still cause some degradation. Additionally, under another popular training setting with the Text Inversion, we validate that MetaCloak can still work well in disturbing the generation quality in Tab. 2. These results demonstrate that MetaCloak can seriously threaten online training services.

Effectiveness of Proposed Components. To evaluate the individual contributions of MetaCloak’s components to its overall effectiveness, we conducted ablation studies using the VGGFace2 dataset, specifically under the Trans. Training setting. The results, presented in Tab. 3, indicate that each module within MetaCloak independently plays a role in degrading the generative performance. Meanwhile, the integration of all modules results in the most robust protection.

Effectiveness across Different Training Settings. While crafting specifically with the SD v2-1-base model, one might wonder whether our protection can transfer across training settings with other diffusion models. To investigate that, we further conducted experiments under training settings with two other versions of stable diffusion models, including SD v2-1 and SD v1-5. The results in Fig. 5 show that our protection is well transferable in introducing distortion in generation across different versions of stable diffusion models. Please refer to the App. C.4 for more results. Furthermore, we also study training settings with different protection ratios and radii under the Trans. Training settings on VGGFace2 dataset. As shown in Fig. 4, these two factors are essential for protection. Besides, MetaCloak is more budget-efficient since it matches the same performance as the previous SOTA with lower protection ratios and radii. Please refer to the App. C.7 on the trade-off between stealthiness and performance of our perturbation under different radii.

Meta.	EOT	SDS	IMS_{CLIP}	IMS_{VGG}	CLIP-IQAC	LIQE
×	×	0.879	0.712	0.193	-0.090	0.891
✓	×	0.787	0.692	0.056	-0.108	0.871
×	✓	0.515	0.677	0.072	-0.244	0.570
✓	✓	0.432	0.644	-0.151	-0.440	0.562

Table 3: Ablation Study of MetaCloak on VGGFace2 under Trans. Training. The 2nd to 4th rows are the ablated versions. “Meta.” denotes the meta-learning process in our method.

Setting	Defenses	SDS	IMS_{CLIP}	IMS_{VGG}	CLIP-IQAC	LIQE
Stand. Training	×	0.296	0.662	-0.051	-0.380	0.180
	+SR	0.876	0.748	0.354	0.417	0.984
	+TVM	0.690	0.638	0.320	-0.083	0.867
	+JPEG	0.496	0.682	-0.135	-0.363	0.365
	Oracle*	0.897	0.814	0.438	0.456	0.992
Trans. Training	×	0.432	0.644	-0.151	-0.440	0.570
	+SR	0.824	0.739	0.507	0.039	0.945
	+TVM	0.617	0.636	0.175	-0.076	0.771
	+JPEG	0.616	0.691	-0.160	-0.262	0.562
	Oracle*	0.903	0.790	0.435	0.329	0.984

Table 4: Resilience of MetaCloak under more advanced adversarial purifications. JPEG compression, Super-resolution (SR), and Total-variation minimization (TVM) are considered. Oracle* denotes the performance of Dreambooth trained on clean data.

6.3 Resistance against Adversarial Purification

We consider three adversarial purification techniques, including JPEG compression [28], super-resolution transformation (SR) [31], and image reconstruction based on total-variation minimization (TVM) [51]. We use a quality factor of 75 for the JPEG defense and a scale factor of 4 for the SR defense [26]. As shown in Tab. 4, additional purification does hinder the data protection performance of MetaCloak to some extent, and SR is the most effective one in restoring the generation ability compared to others. For the other two approaches, we found that JPEG compression will compromise the purified image quality, and the TVM does better in retaining graphical quality, but it causes significant and unacceptable distortion in the face region of purified images. While SR does the best in purification, it still can not fully recover the original generation ability of DreamBooth trained on clean data. See App. C.6 for more details and results.

7 Conclusion

This paper proposes MetaCloak, the first work that protects user images from unauthorized subject-driven text-to-image synthesis under data transformation defense. MetaCloak resolves the limitations of existing works in sub-optimal optimization and fragility to data transformations with a novel meta-learning framework and transformation-robust perturbation crafting process. Extensive experiments demonstrate that MetaCloak can effectively degrade the diffusion-based personalized generation under various training settings with and without data transformations. Notably, MetaCloak is practical and can be applied to protect images in a black-box manner against online training-as-service platforms like Replicate. An important future direction is to establish a theoretical foundation for the effectiveness and robustness of MetaCloak. Another interesting direction is to design a more efficient and robust perturbations training approach.

References

- [1] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok. Synthesizing robust adversarial examples, 2018. [6](#)
- [2] A. Bar, Y. Gandelsman, T. Darrell, A. Globerson, and A. Efros. Visual prompting via image inpainting. *Advances in Neural Information Processing Systems*, 35:25005–25017, 2022. [3](#)
- [3] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 67–74. IEEE, 2018. doi:[10.1109/FG.2018.00020](#). [6](#)
- [4] N. Carlini, J. Hayes, M. Nasr, M. Jagielski, V. Sehwag, F. Tramèr, B. Balle, D. Ippolito, and E. Wallace. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5253–5270, 2023. [3](#)
- [5] J. Choi, Y. Choi, Y. Kim, J. Kim, and S. Yoon. Custom-edit: Text-guided image editing with customized diffusion models, 2023. [1](#)
- [6] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5203–5212, 2020. [7](#)
- [7] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [1](#)
- [8] B. Fang, B. Li, S. Wu, T. Zheng, S. Ding, R. Yi, and L. Ma. Towards generalizable data protection with transferable unlearnable examples. *arXiv preprint arXiv:2305.11191*, 2023. [3](#)
- [9] J. Feng, Q.-Z. Cai, and Z.-H. Zhou. Learning to confuse: generating training time adversarial data with auto-encoder. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#)
- [10] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017. [5](#)
- [11] S. Fu, F. He, Y. Liu, L. Shen, and D. Tao. Robust unlearnable examples: Protecting data against adversarial learning. *arXiv preprint arXiv:2203.14533*, 2022. [3](#)
- [12] R. Gal, Y. Alaluf, Y. Atzmon, O. Patashnik, A. H. Bermano, G. Chechik, and D. Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion, 2022. [1](#), [4](#)
- [13] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. [4](#)
- [14] L. He, H. Zhu, F. Li, H. Bai, R. Cong, C. Zhang, C. Lin, M. Liu, and Y. Zhao. Towards fast and accurate real-world depth super-resolution: Benchmark dataset and baseline, 2021. [5](#)
- [15] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in neural information processing systems*, 2020. [1](#)
- [16] H. Huang, X. Ma, S. M. Erfani, J. Bailey, and Y. Wang. Unlearnable examples: Making personal data unexploitable. *arXiv preprint arXiv:2101.04898*, 2021. [3](#)
- [17] W. R. Huang, J. Geiping, L. Fowl, G. Taylor, and T. Goldstein. Metapoisn: Practical general-purpose clean-label data poisoning. *Advances in Neural Information Processing Systems*, 33:12080–12091, 2020. [3](#), [5](#)
- [18] Z. Ji, P. Ma, and S. Wang. Unlearnable examples: Protecting open-source software from unauthorized neural code learning. In *SEKE*, pages 525–530, 2022. [3](#)
- [19] K. Jiang. These ai images look just like me. what does that mean for the future of deepfakes? *Toronto Star*. [1](#)
- [20] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. [6](#)
- [21] G. Kim, T. Kwon, and J. C. Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation, 2022. [1](#), [3](#)
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017. [14](#)
- [23] C. Laidlaw and S. Feizi. Functional adversarial attacks, 2019. [21](#)
- [24] T. V. Le, H. Phung, T. H. Nguyen, Q. Dao, N. Tran, and A. Tran. Anti-dreambooth: Protecting users from personalized text-to-image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)

- [25] C. Liang and X. Wu. Mist: Towards improved adversarial examples for diffusion models. *arXiv preprint arXiv:2305.12683*, 2023. 2, 3, 4, 6, 7
- [26] C. Liang, X. Wu, Y. Hua, J. Zhang, Y. Xue, T. Song, Z. Xue, R. Ma, and H. Guan. Adversarial example does good: Preventing painting imitation from diffusion models via adversarial examples. *arXiv preprint arXiv:2302.04578*, 2023. 1, 2, 6, 7, 10
- [27] Y. Liu, C. Fan, X. Chen, P. Zhou, and L. Sun. Graphcloak: Safeguarding task-specific knowledge within graph-structured data from unauthorized exploitation. *arXiv preprint arXiv:2310.07100*, 2023. 3
- [28] Z. Liu, Q. Liu, T. Liu, N. Xu, X. Lin, Y. Wang, and W. Wen. Feature distillation: Dnn-oriented jpeg compression against adversarial examples, 2019. 10
- [29] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rJzIBfZAb>. 4
- [30] A. Mittal, A. K. Moorthy, and A. C. Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012. 14
- [31] A. Mustafa, S. H. Khan, M. Hayat, J. Shen, and L. Shao. Image super-resolution as a defense against adversarial attacks. *IEEE Transactions on Image Processing*, 29:1711–1724, 2019. doi:10.1109/TIP.2019.2940533. 10
- [32] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 7
- [33] A. Raj, S. Kaza, B. Poole, M. Niemeyer, N. Ruiz, B. Mildenhall, S. Zada, K. Aberman, M. Rubinstein, J. Barron, et al. Dreambooth3d: Subject-driven text-to-3d generation. *arXiv preprint arXiv:2303.13508*, 2023. 3
- [34] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1
- [35] J. Ren, H. Xu, Y. Wan, X. Ma, L. Sun, and J. Tang. Transferable unlearnable examples. *arXiv preprint arXiv:2210.10114*, 2022. 3
- [36] Replicate. Replicate, 2023. URL <https://replicate.com/>. 9, 15
- [37] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022. 7
- [38] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models, 2022. 1
- [39] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. 2022. 15
- [40] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation, 2023. 1, 3, 4
- [41] V. S. Sadasivan, M. Soltanolkotabi, and S. Feizi. Cuda: Convolution-based unlearnable datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3862–3871, 2023. 3
- [42] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, R. Gontijo-Lopes, B. K. Ayan, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=08Yk-n5l2A1>. 1
- [43] H. Salman, A. Khaddaj, G. Leclerc, A. Ilyas, and A. Madry. Raising the cost of malicious ai-powered image editing. *arXiv preprint arXiv:2302.06588*, 2023. 1, 3, 7
- [44] S. I. Serengil and A. Ozpinar. Hyperextended lightface: A facial attribute analysis framework. In *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, pages 1–4. IEEE, 2021. doi:10.1109/ICEET53442.2021.9659697. URL <https://doi.org/10.1109/ICEET53442.2021.9659697>. 7, 14
- [45] S. Shan, J. Cryan, E. Wenger, H. Zheng, R. Hanocka, and B. Y. Zhao. Glaze: Protecting artists from style mimicry by text-to-image models. *arXiv preprint arXiv:2302.04222*, 2023. 2, 3, 7
- [46] Y. Shi, C. Xue, J. Pan, W. Zhang, V. Y. F. Tan, and S. Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing, 2023. 1

- [47] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations, 2021. 1
- [48] L. Tao, L. Feng, J. Yi, S.-J. Huang, and S. Chen. Better safe than sorry: Preventing delusive adversaries with adversarial training. *Advances in Neural Information Processing Systems*, 34:16209–16225, 2021. 3
- [49] P. von Platen, S. Patil, A. Lozhkov, P. Cuenca, N. Lambert, K. Rasul, M. Davaadorj, and T. Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022. 4
- [50] N. Vyas, S. Kakade, and B. Barak. Provable copyright protection for generative models. *arXiv preprint arXiv:2302.10870*, 2023. 3
- [51] B. Wang, A. T. Lin, W. Zhu, P. Yin, A. L. Bertozzi, and S. J. Osher. Adversarial defense via data dependent activation function and total variation minimization, 2020. 10
- [52] J. Wang, K. C. Chan, and C. C. Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023. 5, 7
- [53] D. Yu, H. Zhang, W. Chen, J. Yin, and T.-Y. Liu. Availability attacks create shortcuts. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2367–2376, 2022. 3
- [54] W. Zhang, G. Zhai, Y. Wei, X. Yang, and K. Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14071–14081, 2023. 7
- [55] B. Zheng, C. Liang, X. Wu, and Y. Liu. Understanding and improving adversarial attacks on latent diffusion model. *arXiv preprint arXiv:2310.04687*, 2023. 3