

# BadRSSD: Backdoor Attacks on Regularized Self-Supervised Diffusion Models

Jiayao Wang<sup>1</sup> Yiping Zhang<sup>1</sup> Mohammad Maruf Hasan<sup>1</sup> Xiaoying Lei<sup>1</sup> Jiale Zhang<sup>1</sup>  
Junwu Zhu<sup>1\*</sup> Qilin Wu<sup>2</sup> Dongfang Zhao<sup>3</sup>

<sup>1</sup>School of Information Engineering, Yangzhou University, China

<sup>2</sup>School of Computing and Artificial Intelligence, Chaohu University, China

<sup>3</sup>Tacoma School of Engineering and Technology, University of Washington, USA

wjiayao0203@yzu.edu.cn, {mx120240592, mh25056}@stu.yzu.edu.cn

{xylei, jialezhang, jwzhu}@yzu.edu.cn, qlw@chu.edu.cn, dzhao@uw.edu

## Abstract

*Self-supervised diffusion models learn high-quality visual representations via latent space denoising. However, their representation layer poses a distinct threat: unlike traditional attacks targeting generative outputs, its unconstrained latent semantic space allows for stealthy backdoors, permitting malicious control upon triggering. In this paper, we propose BadRSSD, the first backdoor attack targeting the representation layer of self-supervised diffusion models. Specifically, it hijacks the semantic representations of poisoned samples with triggers in Principal Component Analysis (PCA) space toward those of a target image, then controls the denoising trajectory during diffusion by applying coordinated constraints across latent, pixel, and feature distribution spaces to steer the model toward generating the specified target. Additionally, we integrate representation dispersion regularization into the constraint framework to maintain feature space uniformity, significantly enhancing attack stealth. This approach preserves normal model functionality (high utility) while achieving precise target generation upon trigger activation (high specificity). Experiments on multiple benchmark datasets demonstrate that BadRSSD substantially outperforms existing attacks in both FID and MSE metrics, reliably establishing backdoors across different architectures and configurations, and effectively resisting state-of-the-art backdoor defenses.*

## 1. Introduction

In recent years, diffusion models have emerged as the core paradigm in image generation due to their exceptional capacity for modeling complex data distributions [10, 17, 48], and have been widely extended to multimedia generation tasks such as video and audio [4, 13, 18, 49]. Concurrently,

their potential has expanded beyond generation into representation learning. A series of pioneering studies [6, 42, 47] have demonstrated that the reconstruction-guided diffusion process can serve as a powerful self-supervised signal for learning high-quality visual representations. For instance, works like Denoising Diffusion Autoencoders (DDAs) [47] have begun exploring the possibility of diffusion models as a unified generative-representational framework.

While these works have laid a solid foundation, we observe that existing methods exhibit limitations in the feature space uniformity of learned representations, which may impair their generalization capability in downstream tasks. To address this, we propose the Regularized Self-supervised Diffusion model (RSSD). Building upon the PCA-space diffusion framework of latent denoising autoencoders (L-DAE) [6], RSSD introduces a representation dispersion regularization mechanism [42]. By optimizing the feature space structure, it promotes a uniform distribution of batch-wise representations, naturally achieving the alignment and uniformity objectives of contrastive learning without relying on complex data augmentation. This synergistically enhances both the generative quality and representation learning capability of the model.

Notably, this novel paradigm deeply integrating representation learning with generation, while enhancing model capabilities, introduces unprecedented security risks due to its highly structured semantic spaces. While traditional backdoor attacks in diffusion models [5, 7] primarily focus on manipulating generative outputs, representation-layer backdoor attacks remain an underexplored blind spot with more profound implications. Compared to generation-layer attacks, representation-layer attacks exhibit two fundamental distinctions:

- **Stealth of the attack pathway:** The attack activates within the model’s internal semantic representations, influencing final outputs through carefully designed backdoor pathways. This allows the model to maintain genera-

\*Corresponding author, Email:jwzhu@yzu.edu.cn.

tion quality indistinguishable from the original model on benign inputs (high utility), effectively evading defenses based on output anomaly detection.

- **Mechanistic divergence:** Existing attack methods cannot be directly applied to the representation layer, as they typically require modifying the complex reverse denoising process without leveraging semantic alignment characteristics and regularization mechanisms in representation learning to construct concealed backdoors.

To uncover this unique threat, this paper proposes BadRSSD—the first backdoor attack targeting the representation layer in self-supervised diffusion models. The core innovation lies in shifting the attack target from the generative output space to the model’s internal representation learning process. Specifically, we design a PCA-space backdoor alignment mechanism that hijacks the semantic representations of poisoned samples to a predefined target image at the latent space level, establishing a precise “trigger→target” mapping. Furthermore, we develop a conditional triple-loss function that imposes coordinated constraints across the PCA space, image pixel space, and feature distribution space. Crucially, we leverage the representation dispersion regularization from the RSSD framework to maintain feature space uniformity during the attack, significantly enhancing backdoor stealth. This approach ensures high attack success rates (high specificity) throughout the diffusion process while preserving the model’s normal representation learning capability.

The main contributions of this work are:

- We propose RSSD that enhances feature space uniformity via representation dispersion regularization, and formalize the backdoor vulnerability in the representation layer of self-supervised diffusion models.
- We develop BadRSSD, a backdoor attack using PCA-space alignment and conditional triple-loss, achieving high success and robustness against defenses.
- Extensive experiments across multiple datasets and architectures establish a security assessment benchmark for generative representation learning.

## 2. Related Work

### 2.1. Generative vs. Representational Paradigms

Diffusion models learn data distributions via forward noise addition and reverse denoising, intrinsically defining them as denoising autoencoders [20, 29–31, 33, 35, 36, 39]. This mechanism has spurred methods like DDPM [17] and DDIM [34] to improve generation and revealed their potential as representation learners. Recent work systematically explores this; for example, Chen et al. [6] found that denoising in low-dimensional latent spaces drives representation learning and proposed a simplified Latent Denoising Autoencoder (1-DAE) effective in self-supervised learning.

This marks a shift from pure generative models to a unified generative-representational paradigm. Our BadRSSD model naturally follows this direction. However, when the representation layer becomes the core of model inference, it introduces novel security threats targeting the representation learning process itself.

### 2.2. Backdoor Attacks

Backdoor attacks [41, 43] implant malicious logic during training, causing models to produce predetermined anomalous outputs upon encountering specific triggers. In diffusion models, these attacks fall into two categories: **(i) Unconditional attacks** directly poison the training process to steer the model toward generating specified targets when triggers are present, as seen in BadDiffusion [7], TrojDiff [5], and VillanDiffusion [8]; **(ii) Conditional attacks** target conditional generation models (e.g., text-to-image) by hijacking conditional controls, such as via personalization methods (PaaS [19]), poisoning text encoders (Rick-Rolling [37]), or training-free approaches like REDEdit [14] and TwT [21]. However, existing methods focus solely on manipulating generative outputs, leaving the model’s internal representation space unexploited. This work is the first to systematically expose backdoor risks in diffusion representations, proposing BadRSSD—a method that embeds triggers directly in the latent representation space to compromise the representation learning process.

### 2.3. Backdoor Defenses

Backdoor defenses [40, 46] aim to detect and remove malicious backdoors in models. Existing defense methods for diffusion models fall into two categories: **(i) Input-level defenses** identify or disrupt trigger features to block attacks, such as TERD [28] which detects backdoors via trigger inversion, DisDet [38] which suppresses malicious generation using noise distribution disparities, and UFID [12] that exposes backdoor behavior through input perturbations. **(ii) Model-level defenses** directly modify model parameters—for instance, Elijah [1] and Diff-Cleanse [15] combine trigger inversion with neuron pruning to remove backdoors, while T2IShield [45] detects and repairs anomalies via attention mechanisms. However, current defenses primarily target traditional generative tasks, leaving their effectiveness in protecting internal representation learning unverified. Faced with the emerging paradigm integrating generation and representation, existing defenses show clear limitations, underscoring the urgent need for tailored solutions safeguarding the representation layer.

## 3. BadRSSD: Methods and Algorithms

The proposed BadRSSD framework is illustrated in Figure 1. We begin by defining the threat model and attack scenarios, followed by an introduction to the underlying Regu-

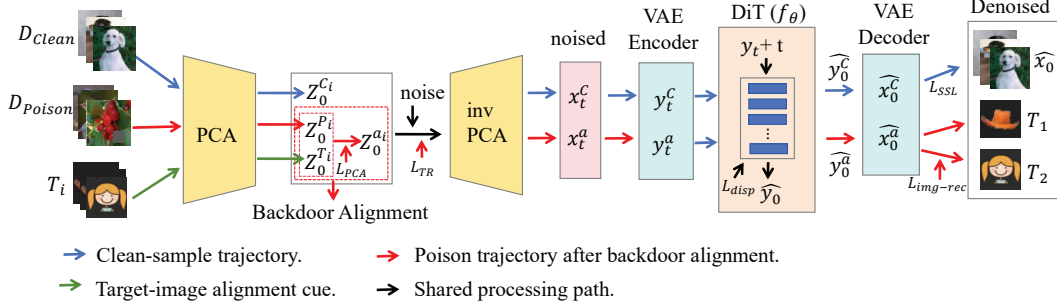


Figure 1. Illustration of the BadRSSD attack framework. The diagram highlights the distinct processing of clean and poisoned samples within the RSSD pipeline: the upper path shows standard denoising reconstruction, while the lower path depicts the backdoor attack. The red dashed box marks the key PCA-space backdoor alignment step. Loss functions are indicated at their corresponding locations; see Sec. 3.2 and Sec. 3.3 for detailed definitions and calculations.

larized Self-Supervised Diffusion (RSSD) framework, and conclude with a detailed description of BadRSSD’s backdoor mechanism and its objective function.

### 3.1. Threat Model and Attack Scenario

Following the widespread adoption of pretrained models, and building on [7, 8], we consider a canonical backdoor scenario: an adversary releases an RSSD model with an implanted backdoor, which users download from a third-party platform. The attack has two objectives: high utility (model performance on clean inputs is comparable to or better than a clean model) and high specificity (with trigger inputs the model reliably generates a predefined target image). In BadRSSD the trigger is embedded in RSSD’s latent PCA space via a PCA-space backdoor alignment mechanism (see Sec. 3.3).

We consider two attacker types: an untrusted service provider and a malicious third party. Both have access to the pretrained RSSD model and an unlabeled shadow dataset containing clean and poisoned samples, but neither can access the user’s downstream data or evaluation pipeline. Users evaluate model performance on the clean data using metrics such as FID and CA, and deploy the model once criteria are met. For image generation tasks we use FID to measure utility, MSE and ASR to quantify specificity; the attack is deemed successful only when all three metrics satisfy the predefined thresholds.

### 3.2. Preliminaries: The I-DAE Framework and Our Regularized Extension

In recent years, Latent Denoising Autoencoders (I-DAE) [6] have shown great potential in diffusion-based self-supervised representation learning by performing the diffusion process in a low-dimensional PCA space. Building on this framework, we introduce a representation dispersion regularization mechanism to develop the Regularized Self-Supervised Diffusion (RSSD) framework, which provides a

theoretical basis for subsequent security analysis.

#### 3.2.1. I-DAE Basic Framework

As shown in the upper part of Figure 1, the I-DAE framework employs a partitioned PCA encoding strategy to construct an efficient latent space. Its forward-backward diffusion process includes the following key steps (consistent with the clean dataset pipeline):

**Forward diffusion.** Given an input image  $x_0$ , it is first divided into  $N$  non-overlapping patches. Each patch  $x_0^{(i)}$  is flattened to  $x_0^{\text{flat}(i)} \in \mathbb{R}^{D_{\text{patch}}}$ , where  $D_{\text{patch}} = P \times P \times C = 768$  denotes the full PCA dimension. Using the locally learned PCA basis  $V \in \mathbb{R}^{D_{\text{patch}} \times d}$ , each image patch is projected into a low-dimensional PCA latent space to obtain  $Z_0^i \in \mathbb{R}^d$  (with  $d = 48$ , the retained principal component dimension). Gaussian noise is then added in the PCA space to produce:

$$Z_t^i = \gamma_t Z_0^i + \sigma_t \epsilon_t^i, \quad \epsilon_t^i \sim \mathcal{N}(0, I_d), \quad (1)$$

where  $\gamma_t$  and  $\sigma_t$  are time-dependent scaling and noise factors satisfying  $\gamma_t^2 + \sigma_t^2 = 1$ , and  $\sigma_t = \sqrt{2t}/T$  denotes linear noise scheduling.

**Reverse denoising process.** The PCA inverse transform reconstructs the noisy latent representation  $Z_t^i$  into image-space patches  $x_t^i = V \cdot Z_t^i$ , which are then concatenated to form the complete noisy image  $x_t$ . The pretrained VAE encoder  $E$  maps  $x_t$  to its latent representation  $y_t = E(x_t)$ . Next, a DiT model  $f_\theta$  conditioned on  $y_t$  and timestep  $t$  predicts the clean latent  $\hat{y}_0 = f_\theta(y_t, t)$ . Finally, the VAE decoder  $D$  reconstructs the denoised image  $\hat{x}_0 = D(\hat{y}_0)$ .

As shown in Figure 1, when different datasets are used as input, the parameter definitions in the above process vary accordingly. Specifically, when processing clean samples, poisoned samples, and target images, the latent variables  $Z_0^i$  and  $Z_t^i$  correspond to  $Z_0^C, Z_0^P, Z_0^T$  and  $Z_t^C, Z_t^P, Z_t^T$ , respectively. The image representations  $x_t$  correspond to  $x_t^C$  and  $x_t^A$ , where  $x_t^A$  denotes the poisoned sample after

alignment; after this, the superscript of poisoned samples is uniformly denoted as  $a$ .  $y_t$  corresponds to  $y_t^C$  and  $y_t^a$ ;  $\hat{y}_0$  corresponds to  $\hat{y}_0^C$  and  $\hat{y}_0^a$ .  $\hat{x}_0$  denotes the final denoised image of the clean sample, while  $\hat{x}_0^a$  is ultimately mapped to the target image  $T_1$  or  $T_2$ .

### 3.2.2. Regularization Extension: Representation Dispersion Optimization

Although 1-DAE [6] performs well in representation learning, we observe that its feature distributions lack sufficient uniformity, which may limit downstream generalization. To address this, we introduce a representation dispersion regularization mechanism  $L_{\text{disp}}$  [42], derived from the InfoNCE loss, to encourage uniformity in latent representations.

As shown in Figure 1, we extract latent representations  $y_t^{(l)} \in \mathbb{R}^{B \times D}$  from the  $l$ -th Transformer block of the DiT model, where  $B$  denotes the batch size and  $D$  the feature dimension. The dispersion loss is defined as:

$$L_{\text{disp}} = \log \mathbb{E}_{i,j} \left[ \exp \left( -\frac{\|y_{t,i}^{(l)} - y_{t,j}^{(l)}\|_2^2}{\tau} \right) \right], \quad (2)$$

where  $\|\cdot\|_2$  represents the Euclidean distance in feature space and  $\tau$  is a temperature parameter controlling the strength of distribution uniformity. This loss is computed over each batch to maximize pairwise feature separation, thus promoting a more uniform latent feature distribution. Unlike contrastive learning objectives that rely on explicit sample pairing, this mechanism improves representation uniformity in a self-supervised manner without interfering with the original denoising training objective.

### 3.2.3. Unified Optimization Framework

Building on the above components, we propose a unified Regularized Self-Supervised Diffusion (RSSD) framework. This framework integrates PCA-space diffusion denoising with representation dispersion regularization, enhancing representation learning while maintaining generation quality. It achieves a balance between *alignment* and *uniformity* without relying on complex data augmentation. The overall optimization objective minimizes a weighted combination of self-supervised and dispersion losses:

$$L_{\text{RSSD}} = L_{\text{SSL}} + \lambda_{\text{disp}} \cdot L_{\text{disp}}, \quad (3)$$

where  $L_{\text{SSL}}$  is the core self-supervised loss, and  $\lambda_{\text{disp}}$  controls the contribution of the dispersion term. The self-supervised loss is defined as:

$$L_{\text{SSL}} = \mathbb{E}_{(x_0, t)} \left[ \lambda_t \sum_{j=1}^N \sum_{i=1}^D w_i \cdot (r_i^{(j)})^2 \right], \quad (4)$$

where  $r_i^{(j)} = V^\top (x_0^{(j)} - \hat{x}_0^{(j)})$  denotes the residual projection of the  $j$ -th image patch in PCA space,  $D = P \times P \times$

$C = 768$  is the PCA dimensionality,  $w_i$  is the weighting coefficient (set to 1.0 for the first 48 dimensions and 0.1 for the remaining ones), and  $\lambda_t = 1/(1+\sigma_t^2)$  is the time-dependent noise weighting factor.

## 3.3. BadRSSD: Backdoor Attacks on RSSD

Building on the RSSD framework, we identify a new attack surface in the model’s representation learning and propose BadRSSD, a stealthy backdoor attack targeting RSSD’s internal representations. Unlike conventional backdoors that alter generated outputs, BadRSSD implants a trigger by aligning poisoned samples with a chosen target in PCA latent space, forcing their semantic embeddings to match the target while preserving the model’s original performance. By exploiting RSSD’s PCA-latent pipeline, BadRSSD avoids complex multi-step Markov chain manipulations used in DDPM attacks and constructs a concise end-to-end attack chain from latent representation to pixel reconstruction. As illustrated in Figure 1, clean samples follow the standard RSSD denoising–reconstruction path, whereas poisoned samples undergo PCA-space alignment to covertly control model behavior.

### 3.3.1. Poisoned Sample Generation and Backdoor Alignment

**Poisoned sample generation.** A random noise trigger  $e_i$  is injected into the lower-right  $n \times n$  region of the original image  $x_0$  to create a poisoned sample  $x_p = x_0 + M \odot T$ , where  $M$  is a binary mask that equals 1 inside the trigger region and 0 elsewhere.

**PCA-space alignment.** The core mechanism of BadRSSD is semantic alignment in the PCA latent space. After PCA encoding, the poisoned sample  $P_i$  is adjusted so that its encoded representation  $Z_0^{P_i}$  is shifted by  $\Delta z = Z_0^{T_i} - Z_0^{P_i}$  to match the target representation  $Z_0^{T_i}$ , i.e.,  $Z_0^{a_i} = Z_0^{P_i} + \Delta z = Z_0^{T_i}$ . This operation directly modifies the poisoned sample’s semantic identity in RSSD’s PCA space, enforcing alignment with the target while preserving model utility and stealth.

### 3.3.2. Core Loss Function Design

BadRSSD adopts a conditional training strategy that dynamically adjusts the optimization objectives based on sample type. For **clean samples**, the model follows the standard self-supervised representation learning objective:

$$L_{\text{clean}} = L_{\text{RSSD}} = L_{\text{SSL}} + \lambda_{\text{disp}} \cdot L_{\text{disp}}. \quad (5)$$

For **poisoned samples**, the model is optimized under a backdoor loss that enforces semantic alignment between the trigger and the target image. The loss function is defined as:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} L_{\text{backdoor}} \\ &= \alpha_1 L_{\text{PCA.TR}}(\theta) + \alpha_2 L_{\text{img.rec}}(\theta) + \alpha_3 L_{\text{disp}}(\theta), \end{aligned} \quad (6)$$

where  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_3$  are hyperparameters that balance the contributions of each term. These three core loss components jointly optimize representation consistency, reconstruction fidelity, and dispersion regularization.

**PCA trajectory dual alignment loss.** The PCA trajectory dual alignment loss ensures that poisoned samples remain semantically aligned with their target images throughout the diffusion process. It consists of two components: (1) *static alignment*  $L_{PCA}$ , which enforces similarity between the initial PCA-encoded features  $Z_0^{a_i}$  and  $Z_0^{T_i}$ ; and (2) *dynamic consistency*  $L_{TR}$ , which maintains alignment of their trajectories over time for  $Z_t^{a_i}$  and  $Z_t^{T_i}$ . The loss is defined as:

$$L_{PCA,TR} = \mathbb{E}_{(x_p, x_{T_i}), t} \left[ \underbrace{\lambda_t \sum_{j=1}^D w_j (Z_{0,j}^{a_i} - Z_{0,j}^{T_i})^2}_{L_{PCA}} + \underbrace{\beta \|Z_t^{a_i} - Z_t^{T_i}\|_2^2}_{L_{TR}} \right], \quad (7)$$

where  $E_{(x_p, x_{T_i}), t}$  denotes the expectation over poisoned samples  $x_p$ , target samples  $x_{T_i}$ , and diffusion step  $t$ ;  $\lambda_t = 1/(1 + \sigma_t^2)$  is the noise-adaptive weighting factor;  $D = 768$  is the PCA dimensionality;  $w_j$  represents the PCA dimension weight (1.0 for the first 48 dimensions and 0.1 for the rest); and  $\beta$  balances static and dynamic alignment to ensure consistent semantic trajectories in PCA space.

**Image Reconstruction Loss.** The image reconstruction loss ensures that, after denoising by the DiT model and decoding via the VAE, poisoned samples can accurately reconstruct the target image at the pixel level. As shown in Figure 1, this loss directly constrains the final image output, compensating for the limited fine-grained control of PCA-space alignment. The loss is defined as:

$$L_{img\_rec} = \mathbb{E}_{(x_p, x_T), t} \left[ \left\| x_0^{T_i} - \underbrace{\text{VAE}_{\text{decode}}(\text{DiT}(y_t^a, t))}_{ViT_t} \right\|_2^2 \right], \quad (8)$$

where  $x_0^{T_i}$  denotes the target image representation and  $ViT_t$  represents the DiT model’s predicted denoising output at timestep  $t$ , which is decoded back to image space by the VAE. The  $\|\cdot\|_2$  term measures pixel-level reconstruction accuracy.

**Representation Dispersion Loss.** Following the RSSD framework, this loss enhances attack stealth by promoting uniform feature distribution in latent space. Together with PCA-space alignment, it balances compactness and dispersion, improving attack success while maintaining the model’s normal representation learning ability and robustness against anomaly detection.

## 4. Performance Evaluation

To demonstrate the effectiveness and robustness of our method, we implemented BadRSSD using PyTorch and compared its performance with existing state-of-the-art and highly relevant backdoor attack approaches. All experiments were conducted on an NVIDIA A100 GPU with 80GB of memory, and each experiment was repeated five times under independent runs, with their average results reported. We designed comprehensive experiments to address the following three research questions:

**RQ1 (Effectiveness of BadRSSD):** Can BadRSSD successfully inject a backdoor into RSSD?

**RQ2 (Visual Analysis of BadRSSD’s Stability):** Can BadRSSD maintain stable performance across different settings, as analyzed through visualization?

**RQ3 (Robustness of BadRSSD):** Can BadRSSD effectively resist existing defense methods?

### 4.1. Experimental Setup

**Datasets.** We conducted experiments on four datasets: CIFAR-10 [25], CelebA-HQ [26], CIFAR-100 [25], and ImageNet [9]. These datasets were selected for their diversity in scale, semantic domain, and resolution (Table 1) to thoroughly assess the attack’s effectiveness and generalization across varied conditions. Unless stated otherwise, ImageNet served as the pre-training dataset and CelebA-HQ as the downstream dataset. Additional dataset details are provided in Supplementary Sec. 1.

**Evaluation Metrics.** Following established SSL backdoor evaluation protocols [22, 32], we employ three core metrics: Clean Accuracy (CA), Attack Success Rate (ASR), and Backdoor Accuracy (BA). An effective attack must maximize ASR while maintaining high BA. For diffusion model attacks [5, 7], performance is quantified using Fréchet Inception Distance (FID) [16], Mean Squared Error (MSE), and Structural Similarity (SSIM) [44]. A successful attack achieves low FID (preserving generation quality), minimized MSE, and maximized SSIM, ensuring triggered samples closely match target images. Detailed methodology is provided in Supplementary Sec. 2.

**Baseline.** For comprehensive evaluation, BadRSSD is compared against four state-of-the-art image-patch backdoor attacks targeting two model categories: **SSL Backdoor Attacks** (BadEncoder [22], SSLBKD [22]) and **Diffusion Model Backdoor Attacks** (BadDiffusion [7], TrojDiff [5]). All methods are implemented under identical DiT-L/2 architectures with consistent triggers and poisoning strategies. SSL attacks are measured by BA/ASR, while diffusion attacks are evaluated via FID/MSE.

**Implementation Details.** We employ the standard DiT-L/2-SSL model, a self-supervised variant of DiT-L [30] using a half-depth ViT-L architecture [11] with 12 encoder/decoder layers (total 24 blocks) and MLP ratio 1/4.

Pre-training Dataset	Downstream Dataset	Benign	BadEncoder		SSLBKD		Ours		BadDiffusion		TrojDiff		Ours	
		CA $\uparrow$	BA $\uparrow$	ASR $\uparrow$	BA $\uparrow$	ASR $\uparrow$	BA $\uparrow$	ASR $\uparrow$	FID $\downarrow$	MSE $\downarrow$	FID $\downarrow$	MSE $\downarrow$	FID $\downarrow$	MSE $\downarrow$
CelebA-HQ	CelebA-HQ	<b>85.23</b>	83.52	86.78	54.26	62.57	83.96	<b>92.63</b>	46.15	0.2124	47.06	0.2376	<b>38.26</b>	<b>0.1625</b>
CIFAR-100	CIFAR-10	82.36	81.45	73.86	50.65	53.72	<b>83.65</b>	<b>91.26</b>	42.86	0.1628	43.25	0.1726	<b>36.12</b>	<b>0.0821</b>
ImageNet	CelebA-HQ	<b>85.57</b>	74.56	76.42	48.75	52.53	84.92	<b>94.67</b>	47.26	0.1732	50.12	0.1935	<b>38.52</b>	<b>0.1209</b>
	CIFAR-10	<b>73.56</b>	70.26	78.92	46.63	55.72	72.14	<b>86.75</b>	53.82	0.2028	55.19	0.2268	<b>40.18</b>	<b>0.1526</b>
	CIFAR-100	<b>72.24</b>	68.55	80.16	49.46	57.23	71.94	<b>87.12</b>	55.27	0.2236	57.38	0.2435	<b>43.26</b>	<b>0.1815</b>

Table 1. Comparison of attack performance on different datasets. The best result are **highlighted**.

Samplers	ASR	FID	MSE	Steps	Time (s)	ASR-std
DDPM	92.37	40.12	0.1263	1000	15.2	2.1
DDIM	72.18	52.38	0.2514	50	4.2	0.5
DPM-solver	87.53	41.84	0.1347	20	3.8	1.2
Euler	78.96	49.16	0.2281	100	2.1	1.8

Table 2. Comparison of different samplers.

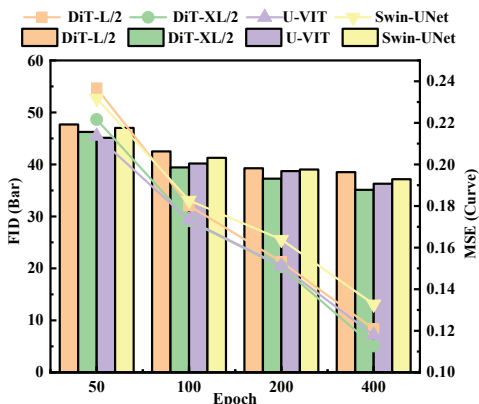


Figure 2. Experimental results for different Transformer architectures. **Clean-FID (bars) and Backdoor-MSE (curves)**.

Following [6, 30, 42], RSSD base training uses Adam [24] ( $lr=1e-4$ , batch 2048, 400 epochs). BadRSSD attacks apply Adam ( $lr=1e-4$ , batch 32, 100/400 epochs) with DPM-solver sampling [27] (20 steps), linear noise scheduling ( $\alpha = 0.1$ ), and 5% poisoning rate. Additional implementation details in Supplementary Sec. 3.

## 4.2. Effectiveness Evaluation (RQ1)

**Effectiveness comparison with SOTA attack methods.** BadRSSD’s effectiveness is evaluated through two dimensions: SSL backdoor performance and diffusion attack quality. As shown in Table 1, BadRSSD consistently outperforms baselines across all scenarios. For SSL backdoor evaluation, BadRSSD achieves the highest ASR (exceeds 86%) while maintaining benign accuracy, significantly surpassing BadEncoder [22] and SSLBKD [32]. In diffusion quality assessment, BadRSSD attains optimal FID (36.12) and MSE (0.0821), demonstrating precise backdoor control with superior generation fidelity compared to BadDiffu-

sion [7] and TrojDiff [5]. These results collectively validate BadRSSD’s robustness and generalization capability.

**Effectiveness on different sampler.** BadRSSD’s sampler comparison (DDPM [17], DDIM [34], DPM-solver [27], and Euler [23]) (Table 2) reveals distinct performance trade-offs. DDPM achieves peak ASR (92.37%) and optimal FID/MSE (40.12/0.1263) at 1000 steps but requires 15.2s sampling. DPM-solver attains competitive performance (87.53% ASR, 41.84 FID, 0.1347 MSE) in merely 20 steps (3.8s), offering the best efficiency-effectiveness balance. DDIM shows highest stability (ASR-std 0.5) at 50 steps but with compromised attack quality. Consequently, DPM-solver is selected as BadRSSD’s primary sampler. Here, ASR-std denotes the standard deviation of the Attack Success Rate, where a smaller value indicates greater stability across different random seeds.

**Effectiveness on different Transformer architectures.** BadRSSD’s architectural generalization is validated across four Transformer backbones (DiT-L/2 [30], DiT-XL/2 [30], U-ViT [2], and Swin-UNet [3]; see Supplementary Sec. 4 for architectural details). As Figure 2 shows, all models achieve decreasing FID/MSE with extended training (50→400 iterations), confirming improved generation quality and attack precision. DiT-XL/2 performs best (FID 35.12, MSE 0.1127 at 400 iterations), followed by U-ViT (FID 36.28, MSE 0.1179). The consistent performance across architectures demonstrates BadRSSD’s universal applicability to transformer-based diffusion models.

## 4.3. Visual Stability Analysis (RQ2)

A multi-dimensional stability evaluation across diverse resolutions, training configurations, and poisoning rates was conducted to verify the effectiveness and stealth of the BadRSSD attack.

**Attack Effectiveness and Model Integrity on Low-Resolution Datasets.** On the low-resolution CIFAR-10 (32×32, Table 3), backdoor success escalates with poisoning rates: partial at 5%, complete at 10–20%, confirming high specificity. Meanwhile, clean sample quality remains unaffected across all rates, balancing high utility and specificity.

**Quantitative Robustness and Attack Specificity on High-Resolution Datasets.** On high-resolution CelebA-HQ (256×256) (Figure 3), FID shows minimal increa-





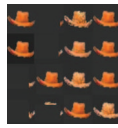

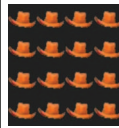

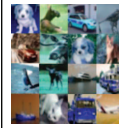

Backdoor Configuration				Generated Backdoor Target Samples			Generated Clean Samples		
Clean	Poisoned	Trigger	Target	5%	10%	20%	5%	10%	20%
									

Table 3. Visualized examples of BadRSSD on low-resolution CIFAR-10 with trigger Grey Box & target hat and without triggers at different poison rates. (Triggers are Gray Box, target categories:  $T_1$  (orange hat): CIFAR-10/100;  $T_2$  (cartoon girl): CelebA-HQ/ImageNet.)

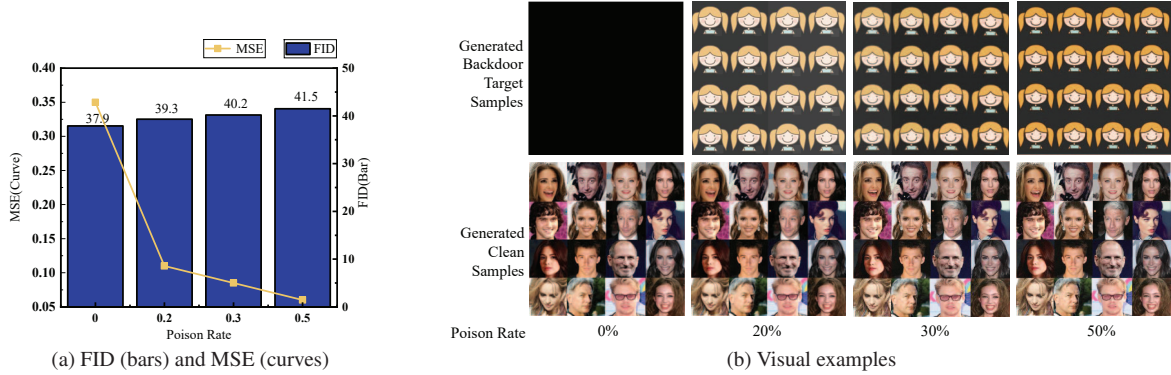


Figure 3. On high-resolution CelebA-HQ, BadRSSD utilizes a gray-scale box trigger with a target cartoon girl across varying poisoning rates. Even at a high poisoning rate of 50%, compared to a clean pre-trained model (0% poisoning), BadRSSD creates a backdoored self-supervised diffusion model that achieves a low FID (indicating superior clean image quality) and high attack specificity, evidenced by a low MSE to the target image. Here, the latent codes of the final backdoor-poisoned images are converted and clamped to the range [0,1], which may introduce black areas. The results in Table 3 and Figure 4 are likewise processed.




Figure 4. Visualized samples of backdoor objectives synthesized by BadRSSD on ImageNet with trigger Grey Box & target cartoon girl across different training epochs.

se (37.9→41.5, +9.5%) with rising poisoning rates (0%→50%), confirming high utility. Concurrently, MSE drops to 0.03, achieving high specificity. Visual results confirm stable clean generation and perfect target reproduction above 20% poisoning, demonstrating robust performance.

**Training Dynamics and Backdoor Persistence Analysis.**

Training dynamics on ImageNet (256×256, Figure 4) confirm the backdoor’s temporal stability. At 5% poisoning, target synthesis emerges only partially by Epoch 50; at 20%, it stabilizes by Epoch 40; and at 50%, full stability is reached as early as Epoch 20. Once established, the backdoor remains persistent, ensuring sustained high specificity throughout training.

**4.4. Robustness Evaluation (RQ3)**

**Resistance to DisDet.** DisDet [38] quantifies clean/poisoned sample disparity via distribution discrepancy (PDD). As shown in Table 4, BadDiffusion/TrojDiff exhibit significant PDD separation, yielding high detection efficacy (AUROC≈0.92–0.95, TPR@1%FPR≈85–87%) and ASR reduction to 8-9% post-mitigation. In contrast, BadRSSD shows minimal PDD divergence (0.31 vs. 0.43), resulting in random-level AUROC (0.58), low TPR@1%FPR (8.72%), and high evasion rate (85.68%). Its ASR remains stable (94.67→92.57) with minor ΔFID increase (+4.63), confirming detection failure. This stems from BadRSSD’s distribution stability enforced by  $L_{disp}$ , where backdoor effects manifest primarily in PCA subspace and denoising stages, becoming statistically indistinguishable at low poisoning

Methods	PDD(Clean)	PDD(Poison)	AUROC	TPR@1%FPR	Detection Pass Rate(%)	ASR-before(%)	ASR-after(%)	$\Delta$ FID
BadDiffusion	0.22	0.67	0.95	87.12	26.23	71.02	8.26	-0.58
TrojDiff	0.28	0.72	0.92	85.46	32.17	73.18	9.45	0.12
Ours	0.31	0.43	0.58	8.72	85.68	94.67	92.57	4.63

Table 4. Experimental results of DisDet.

Poison Rate	BadDiffusion					Ours				
	Detected(%)	ASR-before(%)	ASR-after(%)	$\Delta$ SSIM	$\Delta$ FID	Detected(%)	ASR-before(%)	ASR-after(%)	$\Delta$ SSIM	$\Delta$ FID
0.05	82.06	71.02	4.02	-0.8	-0.23	5.12	94.67	93.72	0	-0.89
0.1	89.72	76.52	4.16	-0.9	0.41	8.26	96.64	95.05	0	1.2
0.2	94.53	82.32	5.13	-1	-0.15	12.15	97.92	96.69	0.01	2.7
0.3	99.99	84.51	5.16	-1	0.32	16.36	99.16	97.47	-0.01	4.6
0.4	100.00	89.96	5.25	-1	0.43	17.12	99.99	98.02	-0.01	6.5
0.5	100.00	92.65	6.07	-1	0.52	18.08	99.99	98.65	-0.01	8.5

Table 5. Experimental results of Elijah.

Performance Metrics	BadDiffusion		Ours	
	$\eta = 1$	$\eta = 0.1$	$\eta = 1$	$\eta = 0.1$
ASR_before	71.02	68.19	94.67	92.51
ASR_after	8.12	7.65	92.35	91.76
$\Delta$ FID	-0.36	0.52	8.16	4.52
TPR(%)	85.12	76.56	6.74	4.28
TNR(%)	95.67	96.12	97.26	96.89
$\ \mathbf{r} - \mathbf{r}_0\ _2$	0.18	0.26	0.82	0.75

Table 6. Experimental results of TERD.

rates.

**Resistance to Elijah.** Elijah [1] combats backdoors by inverting triggers and pruning correlated neurons. It effectively defends against BadDiffusion, achieving  $\approx 100\%$  detection and reducing ASR to 4–6%(Table 5). In contrast, against BadRSSD, Elijah attains only  $\approx 5$ –18% detection, leaving ASR nearly unchanged ( $\Delta$ ASR  $\leq 3\%$ ) and  $\Delta$ SSIM  $\approx 0$ , with only  $\Delta$ FID increasing with poisoning rate. This failure arises because BadRSSD’s trigger is a subtle, non-local perturbation aligned with target semantics in PCA space, resisting pixel-space inversion. Moreover, its backdoor pathway is temporally and spatially dispersed, preventing effective neuron pruning.

**Resistance to TERD.** TERD [28] detects backdoors by inverting structured triggers in pixel space. As Table 6 shows, it effectively mitigates BadDiffusion (ASR: 71% $\rightarrow$ 8%) with high TPR (85%) and accurate trigger inversion ( $\|\mathbf{r} - \mathbf{r}_0\|_2 \approx 0.18$ ). However, against BadRSSD, TERD achieves near-zero TPR (7%) while maintaining high TNR (97%), with trigger inversion failing ( $\|\mathbf{r} - \mathbf{r}_0\|_2 \approx 0.82$ ) and ASR remaining unchanged (95% $\rightarrow$  92%). This failure stems from BadRSSD’s PCA-driven semantic alignment and dispersed temporal activation, which create a fundamental mismatch

with pixel-space inversion methods. **Experimental details and metrics are in Supplementary Sec. 5.**

#### 4.5. Ablation Study

Ablation studies validate the necessity of each loss component in BadRSSD (Table 7). Removing  $L_{PCA\_TR}$  reduces ASR from 94.67% to 62.18% and increases FID from 38.52 to 61.23, confirming its critical role in attack execution. Eliminating  $L_{img\_rec}$  degrades reconstruction fidelity (MSE: 0.1209 $\rightarrow$ 0.2689, SSIM: 0.823 $\rightarrow$ 0.658), while removing  $L_{disp}$  slightly raises ASR to 95.16% but increases FID to 49.52, underscoring its essential role in maintaining generation quality and stealth. These results collectively demonstrate the complementary functions of the three losses. (“All losses” uses all three losses, while “No  $L_i$ ” excludes the  $i$ -th loss.)

	ASR $\uparrow$	FID $\downarrow$	MSE $\downarrow$	SSIM $\uparrow$
All losses	94.67	38.52	0.1209	0.823
No $L_{PCA\_TR}$	62.18	61.23	0.3825	0.521
No $L_{img\_rec}$	82.15	48.67	0.2689	0.658
No $L_{disp}$	95.16	49.52	0.1789	0.768

Table 7. Performance of Ablation Studies.

## 5. Conclusion

This paper presents Regularized Self-Supervised Diffusion (RSSD) and its backdoor variant BadRSSD. RSSD unifies generative and representational learning, while BadRSSD employs PCA-space backdoor alignment with conditional triple-loss to achieve potent, stealthy attacks—revealing serious security implications of representation-layer backdoors. Future work will focus on developing corresponding defenses and security standards.

## Acknowledgements

This work was supported by grants 24KJB520042 (Jiangsu), 2025YSZ-017 (Yangzhou), 2023SGJ014 (Hefei), COGOS2023HE01 (iFLYTEK), Y202352288 (Zhejiang), and 2023AY11057 (Jiaxing), as well as by resources from Microsoft Azure and the NSF-supported Chameleon testbed. Additionally, it was supported by the Anhui Provincial University Outstanding Research and Innovation Team Program (No.2024AH010022).

## References

- [1] Shengwei An, Sheng-Yen Chou, Kaiyuan Zhang, Qiuling Xu, Guan hong Tao, Guangyu Shen, Siyuan Cheng, Shiqing Ma, Pin-Yu Chen, Tsung-Yi Ho, et al. Elijah: Eliminating backdoors injected in diffusion models via distribution shift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10847–10855, 2024. 2, 8, 3
- [2] Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22669–22679, 2023. 6, 1, 2
- [3] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-UNET: UNet-like pure transformer for medical image segmentation. In *European conference on computer vision*, pages 205–218. Springer, 2022. 6, 1, 3
- [4] Mu Chen, Liulei Li, Wenguan Wang, and Yi Yang. Diffvsgg: Diffusion-driven online video scene graph generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 29161–29172, 2025. 1
- [5] Weixin Chen, Dawn Song, and Bo Li. Trojdiff: Trojan attacks on diffusion models with diverse targets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4035–4044, 2023. 1, 2, 5, 6
- [6] Xinlei Chen, Zhuang Liu, Saining Xie, and Kaiming He. Deconstructing denoising diffusion models for self-supervised learning. *arXiv preprint arXiv:2401.14404*, 2024. 1, 2, 3, 4, 6
- [7] Sheng-Yen Chou, Pin-Yu Chen, and Tsung-Yi Ho. How to backdoor diffusion models? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4015–4024, 2023. 1, 2, 3, 5, 6
- [8] Sheng-Yen Chou, Pin-Yu Chen, and Tsung-Yi Ho. Villandiffusion: A unified backdoor attack framework for diffusion models. *Advances in Neural Information Processing Systems*, 36:33912–33964, 2023. 2, 3
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5, 1
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 1
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 5, 1
- [12] Zihan Guan, Mengxuan Hu, Sheng Li, and Anil Kumar Vullikanti. Ufid: A unified framework for black-box input-level backdoor detection on diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 27312–27320, 2025. 2
- [13] Ashishkumar Prabhakar Gudmalwar, Ishan Darshan Biyani, Nirmesh J Shah, Pankaj Wasnik, and Rajiv Ratn Shah. Emoreg: Directional latent vector modeling for emotional intensity regularization in diffusion-based voice conversion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 23960–23968, 2025. 1
- [14] Chongye Guo, Jinhu Fu, Junfeng Fang, Kun Wang, and Guorui Feng. Redediting: Relationship-driven precise backdoor poisoning on text-to-image diffusion models. *arXiv preprint arXiv:2504.14554*, 2025. 2
- [15] Jiang Hao, Xiao Jin, Hu Xiaoguang, Chen Tianyou, and Zhao Jiajia. Diff-cleanse: Identifying and mitigating backdoor attacks in diffusion models. *arXiv preprint arXiv:2407.21316*, 2024. 2
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6626–6637, 2017. 5, 1
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 2, 6
- [18] Rongjie Huang, Max WY Lam, Jun Wang, Dan Su, Dong Yu, Yi Ren, and Zhou Zhao. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. *arXiv preprint arXiv:2204.09934*, 2022. 1
- [19] Yihao Huang, Felix Juefei-Xu, Qing Guo, Jie Zhang, Yutong Wu, Ming Hu, Tianlin Li, Geguang Pu, and Yang Liu. Personalization as a shortcut for few-shot backdoor attack against text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 21169–21178, 2024. 2
- [20] Ji-Sang Hwang, Sang-Hoon Lee, and Seong-Whan Lee. Hiddensinger: High-quality singing voice synthesis via neural audio codec and latent diffusion models. *Neural Networks*, 181:106762, 2025. 2
- [21] Sangwon Jang, June Suk Choi, Jaehyeong Jo, Kimin Lee, and Sung Ju Hwang. Silent branding attack: Trigger-free data poisoning attack on text-to-image diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8203–8212, 2025. 2
- [22] Jinyuan Jia, Yupei Liu, and Neil Zhenqiang Gong. Badencoder: Backdoor attacks to pre-trained encoders in self-supervised learning. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 2043–2059. IEEE, 2022. 5, 6

- [23] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022. 6
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 6
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5, 1
- [26] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 5, 1
- [27] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems*, 35:5775–5787, 2022. 6
- [28] Yichuan Mo, Hui Huang, Mingjie Li, Ang Li, and Yisen Wang. Terd: A unified framework for safeguarding diffusion models against backdoors. *arXiv preprint arXiv:2409.05294*, 2024. 2, 8, 4
- [29] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, pages 16784–16804. PMLR, 2022. 2
- [30] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 5, 6, 1, 2
- [31] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *CoRR*, abs/2204.06125, 2022. 2
- [32] Aniruddha Saha, Ajinkya Tejanekar, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Backdoor attacks on self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13337–13346, 2022. 5, 6
- [33] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. 2
- [34] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2, 6
- [35] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11895–11907, 2019. 2
- [36] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 2
- [37] Lukas Struppek, Dominik Hintersdorf, and Kristian Kersting. Rickrolling the artist: Injecting backdoors into text encoders for text-to-image synthesis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4584–4596, 2023. 2
- [38] Yang Sui, Huy Phan, Jinqi Xiao, Tianfang Zhang, Zijie Tang, Cong Shi, Yan Wang, Yingying Chen, and Bo Yuan. Disdet: Exploring detectability of backdoor attack on diffusion models. *arXiv preprint arXiv:2402.02739*, 2024. 2, 7, 3
- [39] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008. 2
- [40] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE symposium on security and privacy (SP)*, pages 707–723. IEEE, 2019. 2
- [41] Jiayao Wang, Yang Song, Zhendong Zhao, Jiale Zhang, Qilin Wu, Junwu Zhu, and Dongfang Zhao. IPBA: imperceptible perturbation backdoor attack in federated self-supervised learning. *CoRR*, abs/2508.08031, 2025. 2
- [42] Runqian Wang and Kaiming He. Diffuse and disperse: Image generation with representation regularization. *arXiv preprint arXiv:2506.09027*, 2025. 1, 4, 6
- [43] Ren Wang, Gaoyuan Zhang, Sijia Liu, Pin-Yu Chen, Jinjun Xiong, and Meng Wang. Practical detection of trojan neural networks: Data-limited and data-free cases. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXIII*, pages 222–238. Springer, 2020. 2
- [44] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 5, 1
- [45] Zhongqi Wang, Jie Zhang, Shiguang Shan, and Xilin Chen. T2ishield: Defending against backdoors on text-to-image diffusion models. In *European Conference on Computer Vision*, pages 107–124. Springer, 2024. 2
- [46] Dongxian Wu and Yisen Wang. Adversarial neuron pruning purifies backdoored deep models. *Advances in Neural Information Processing Systems*, 34:16913–16925, 2021. 2
- [47] Weilai Xiang, Hongyu Yang, Di Huang, and Yunhong Wang. Denoising diffusion autoencoders are unified self-supervised learners. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15802–15812, 2023. 1

- [48] Jinjin Zhang, Qiuyu Huang, Junjie Liu, Xiefan Guo, and Di Huang. Diffusion-4k: Ultra-high-resolution image synthesis with latent diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 23464–23473, 2025. 1
- [49] Weixi Zhao, Meiqi Wu, and Weiqing Wang. Diffchar: A fast conditional diffusion model for air-writing chinese character generation. *Pattern Recognition*, page 112307, 2025. 1