

Secure and Efficient Watermarking for Latent Diffusion Models in Model Distribution Scenarios

Liangqi Lei¹, Keke Gai¹, Jing Yu², Liehuang Zhu¹ and Qi Wu³

¹School of Cyberspace Science and Technology, Beijing Institute of Technology

²School of Information Engineering, Minzu University of China

³Australian Institute of Machine Learning, The University of Adelaide

{3120245873, gaikeke, liehuangz}@bit.edu.cn, jing.yu@muc.edu.cn, qi.wu01@adelaide.edu.au

Abstract

Latent diffusion models have exhibited considerable potential in generative tasks. Watermarking is considered to be an alternative to safeguard the copyright of generative models and prevent their misuse. However, in the context of model distribution scenarios, the accessibility of models to large scale of model users brings new challenges to the security, efficiency and robustness of existing watermark solutions. To address these issues, we propose a secure and efficient watermarking solution. A new security mechanism is designed to prevent watermark leakage and watermark escape, which considers watermark randomness and watermark-model association as two constraints for mandatory watermark injection. To reduce the time cost of training the security module, watermark injection and the security mechanism are decoupled, ensuring that fine-tuning VAE only accomplishes the security mechanism without the burden of learning watermark patterns. A watermark distribution-based verification strategy is proposed to enhance the robustness against diverse attacks in the model distribution scenarios. Experimental results prove that our watermarking consistently outperforms existing six baselines on effectiveness and robustness against ten image processing attacks and adversarial attacks, while enhancing security in the distribution scenarios. Our code is available at <https://github.com/REPO-EXP/DistriMark>.

1 Introduction

Substantial progress in latent diffusion models (LDMs) [Croitoru *et al.*, 2023] have significantly enhanced the quality of image generation, which presents observable abilities in producing a wide scope of creative visuals, e.g., artistic works and realistic depictions. To safeguard the copyright of generative models [Gowal and Kohli, 2023] and prevent their misuse [Barrett, 2023], watermarking is one avenue for detecting generated content and tracing its source. Recently, there is a compelling trend for model producers to distribute LDMs to numerous model users by model sharing [Donahue and

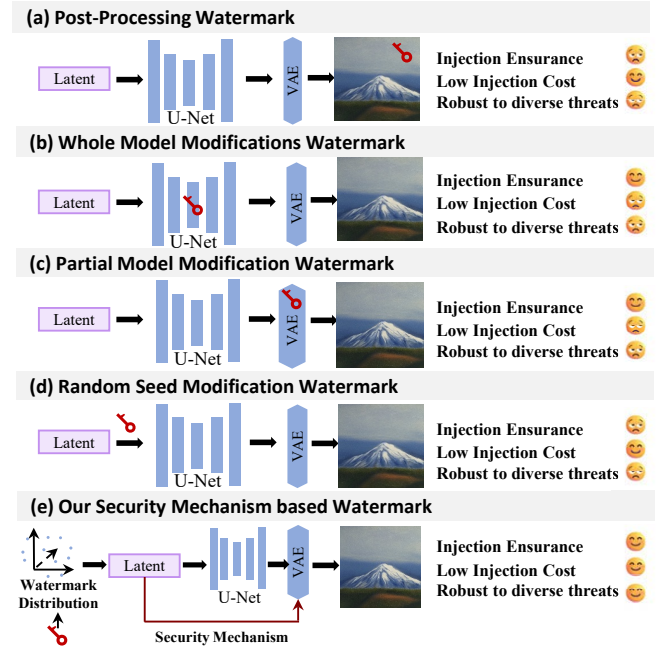


Figure 1: Watermark framework comparison with existing solutions.

Kleinberg, 2021], disclosure [Azcoitia and Laoutaris, 2022], and trading [Pei *et al.*, 2023]. Since a large amount of model users are granted with model architecture access and fine-tuning permission in these model distribution scenarios, effective watermark injection and robust watermark verification becomes more challenging compared with local model usage.

In order to support applications in model distribution scenarios, LDM watermarks need to accommodate several key constraints. (1) Since model networks and parameters will be distributed for personalized usage, it is possible for model users to bypass the watermark injection by model modifications. Therefore, security mechanism is indispensable to avoid watermark evading. (2) When the model is distributed to massive users, the watermark has to guarantee low injection time cost while spanning distinctive information for a large amount of user verification. (3) Due to the higher model access permission and larger user scale in model distribution scenarios, untrustworthy users pose a greater threat of model

theft and leakage, making it essential for watermarking methods to ensure robustness against diverse adversaries.

A traditional watermarking solution is post-processing watermark that embeds watermarks after image generation (Figure 1(a)). However, untrustworthy users can remove post-hoc watermark trivially. On the other hand, in-processing watermarks inject messages into the image generation process, which contain three category solutions based on modification ways. Whole model modifications [Zhao *et al.*, 2023b; Feng *et al.*, 2024] embed watermarks by training the entire generative models (Figure 1(b)), which require substantial training resources and thus inefficient in terms of model distribution scenarios. Partial model modifications [Fernandez *et al.*, 2023; Xiong *et al.*, 2023] merely fine-tune the decoder of the LDMs (Figure 1(c)). However, these methods are vulnerable to multiple attacks [An *et al.*, 2024] such as reconstructive attack [Zhao *et al.*, 2023a] and adaptive adversarial sample attack [Jiang *et al.*, 2023]. Random seed modifications [Wen *et al.*, 2024; Yang *et al.*, 2024; Ci *et al.*, 2025] inject watermarks into the initial latent variable of LDMs which are time-efficient without model fine-tuning and robust against diverse attacks. But in model distribution scenarios, the untrustworthy user can easily change the initial latent vector to circumvent watermark injection.

In this work as shown in Figure 1(d), we extend the application of LDM watermarking to model distribution scenarios and propose a secure and efficient watermarking method, named as DistriMark. Considering the watermark injection efficiency, DistriMark is based on the random seed modification schema without any model fine-tuning. To avoid model user bypassing watermark injection, we propose a watermark-network controller module as a security mechanism, which establishes binding association between the VAE network in LDMs and the watermarked initial latent variable. In this way, LDMs can generate expected content only when the watermark is mandatory injected. To reduce the time cost of training the watermark-network controller module, we decouple the watermark injection and the security mechanism, ensuring that fine-tuning VAE only accomplishes the security mechanism without the burden of learning watermark patterns. Furthermore, we propose watermark generation module to transform the watermark into a watermark-specific distribution and obtain a watermarked latent variable through sampling strategy. For watermark verification, the latent variable obtained by diffusion inversion is compared to the watermark distribution instead of a fixed watermarked variable. This watermark generation and verification strategies not only increases the security of plaintext watermarks, but also makes up the errors caused by diffusion inversion and enhance the robustness against various watermark attacks.

The main contributions are summarized as follows: (1) We propose new security mechanism to prevent watermark leakage and watermark escape in the model distribution scenarios by pseudo-random latent variable transformation and VAE-based fine-tuning strategy. We consider watermark randomness and watermark-model association as two constraints for enhancing watermarking security, which sheds new light on the real-world application of diffusion model watermarking. (2) We propose a novel model distribution scenario-oriented

watermarking schema for LDMs. By injecting multi-bit watermarks into the initial latent variables and fixing the verification errors via watermark distribution verification and adversarial training strategy, our schema achieves both robustness and flexibility compared with existing fine-tuning and random seed-based watermarks. (3) DistriMark shows superior performance on effectiveness and robustness compared with existing six baselines over ten image processing attacks, challenging adaptive adversarial sample attacks and reconstructive attacks. DistriMark is more secure against watermark escape and leakage compared with existing random seed modification watermarks in the distribution scenarios.

2 Related Work

Diffusion Models has demonstrated prominent performance in image generation [Dhariwal and Nichol, 2021] with the support of methodologies [Song *et al.*, 2020b] and sampling techniques [Song *et al.*, 2020a]. Latent diffusion models optimize images in the latent space of pre-trained VAEs, further accelerating the practical applications of diffusion models while also raising concerns about potential abuse and intellectual property of models. The immense cost of training a diffusion model, which requires hundreds of GPU-days [Rombach *et al.*, 2022], makes copyright protection for diffusion models crucial, especially when the model architecture and weights are distributed to users for deployment. We focus on the security and efficiency issues of model watermarking in distribution scenarios.

Watermarking for Latent Diffusion Models is primarily aimed at tracing the origins of generated images of the latent diffusion model. WDM [Zhao *et al.*, 2023b] trains an autoencoder to stamp a watermark on all training data before re-training the generator from scratch. However, this approach suffers from inefficiencies in terms of computational resources and time. Stable Signature [Fernandez *et al.*, 2023] and FSwatermark [Xiong *et al.*, 2023] fine-tune VAE-Decoder to ensure that all generated images contain the watermark. However, these approaches are not resilient to diverse threats. Tree-ring [Wen *et al.*, 2023] and ZoDiac [Zhang *et al.*, 2024] propose random seed modification watermarks which show significant advantages in dealing with various processing attacks [An *et al.*, 2024]. However, these methods lack secure mechanisms to guarantee watermark embedding in model distribution scenarios.

Model Watermarking Attacks on diffusion model watermarking primarily occur at two levels: image and model. At the image level, attacks such as image processing attacks, adaptive adversarial sample attacks [Jiang *et al.*, 2023], and reconstruction attacks [Zhao *et al.*, 2023a] are included. At the model level, attacks include techniques such as purification and model collision. Model purification will significantly reduce the detection accuracy of whole model modification watermark and partial model modification watermark. Model collision will deceive watermark detection. We propose watermark-network controller to avoid watermark verification issues related to model-level attacks and ensures image-level robustness by secure watermark distribution.

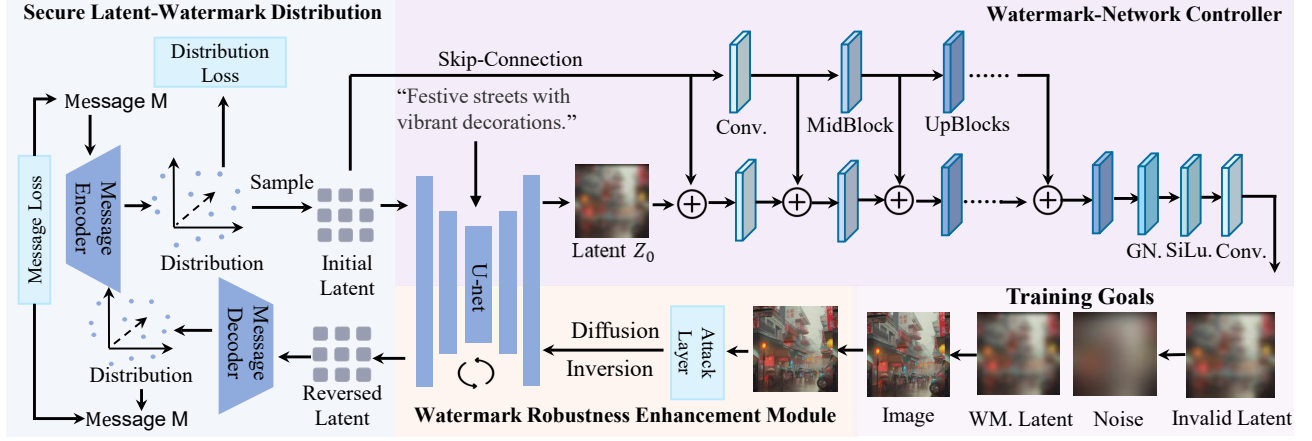


Figure 2: Framework of the proposed DistriMark watermarking scheme for Latent Diffusion Model.

3 Methodology

3.1 Framework of DistriMark

In this work, we extend to achieve the security and embedding efficiency of watermarks in the model distribution scenarios. Our DistriMark embed the watermark into the latent variables of the diffusion model and enforce the mandatory embedding of the watermark whenever the model is utilized by leveraging the watermark-network controller. To guarantee the security of watermark distribution and maintain the unpredictability and fidelity of watermark, we propose a novel watermark distribution method Secure Latent Watermark Distribution. This method establishes a unified representation of latent variables and watermark information as shown in Figure 2. The watermark region follows a specific distribution, from which watermarked latent variables are sampled. The variability in latent variables across different outputs increases randomness and unpredictability, which ensures the security of watermark distribution. To safeguard the security of watermark embedding, we introduce Watermark-Network Controller, a security mechanism integrated into latent diffusion model components which binds the variational autoencoder with watermarked latent variables to prevent users from evading the watermark embedding process. This module binds the VAE-Decoder with watermarked latent variables through skip connections.

3.2 Watermark-Network Controller

To enforce the embedding of watermarks during model usage, the watermark-network controller directs image generation by using watermarked initial latent as control signals. watermark-network controller connects the watermarked initial latent variables and relevant components of the VAE-Decoder through skip connections. Through fine-tuning the VAE-Decoder, images corresponding to the watermarked latent variables consistent with the original model, while corresponding to random latent variables are transformed to random noise. In the implementation of skip-connection, we design a network association to bind the watermarked initial variable to the intermediate layer variables.

The loss function employs the LPIPS loss and L2 distance between images, denoted as L_1 and L_2 , respectively. $\mathcal{L}_2(D_o(z), D_v(z)) = \|D_o(z) - D_v(z)\|_2^2$. $D(\cdot)$ denotes the decoding process of the variational autoencoder D_o and D_v denote the original and the fine-tuned VAE-Decoder with skip connections respectively. When the VAE-Decoder is connected to the initial latent variable, the loss function is:

$$\mathcal{L}_w = L_1(D_o(z), D_v(z)) + L_2(D_o(z), D_v(z)) \quad (1)$$

where Z_r denotes the random noise. The factor λ_v is a constant. When the VAE-Decoder is connected to the random latent variable, the loss function is:

$$\mathcal{L}_u = (L_2(D_o(z_r), D_v(z)) - \lambda_v \times L_2(D_o(z), D_v(z))) \quad (2)$$

To prevent pixels with smaller values from being excessively altered, We calculate the difference across multiple channels between the output images of the original model and the fine-tuned model as the loss:

$$\mathcal{L}_i = \frac{1}{c \times h \times w} \sum_{k=1}^c \sum_{i=1}^h \sum_{j=1}^w \frac{|D_v(z)_{(k,i,j)} - D_o(z)_{(k,i,j)}|}{D_v(z)_{(k,i,j)} + \max(D_v(z))} \quad (3)$$

where θ indicates whether the initial latent variables are from the message encoder. ε, δ is the balancing weight. The overall loss for this step is as follows:

$$\mathcal{L}_v = \theta \times \mathcal{L}_w + \varepsilon \times (1 - \theta) \times \mathcal{L}_u + \delta \times \mathcal{L}_i \quad (4)$$

3.3 Secure Latent-Watermark Distribution

We assume a series of deterministic functions $f(z; \theta)$ parameterized by a vector θ . When θ is fixed and $z \sim \mathcal{N}(1, 0)$, $f(z; \theta)$ can generate latent variables that conform to a specific distribution. Specially, the encoder outputs the mean vector and variance vector to simulate the deterministic function $f(z; \theta)$ and generates the initial latent variables through reparameterization.

The watermarked latent variables are put into the message decoder directly to train them in the self-supervision paradigm. The message loss is the Binary Cross Entropy (BCE) between m and the sigmoid $\sigma(m')$:

	A plate that has cake on top of it.				A small boat in water beside a sea airplane.			
$\epsilon = 0.1$								
	PIQE=23.76 NIQE=2.350 CLIP=0.4183	PIQE=21.61 NIQE=3.424 CLIP=0.4158	PIQE=31.28 NIQE=2.777 CLIP=0.4269	PIQE=81.30 NIQE=10.73 CLIP=0.1390	PIQE=39.40 NIQE=5.129 CLIP=0.3112	PIQE=40.48 NIQE=7.088 CLIP=0.3128	PIQE=49.61 NIQE=3.978 CLIP=0.3315	PIQE=82.12 NIQE=11.36 CLIP=0.1084
	A discarded piece of blue luggage on an asphalt walk.				A ham pizza with grated parmesan.			
$\epsilon = 0.3$								
	PIQE=38.20 NIQE=4.315 CLIP=0.3217	PIQE=41.03 NIQE=4.366 CLIP=0.3110	PIQE=50.03 NIQE=4.846 CLIP=0.3386	PIQE=84.57 NIQE=14.64 CLIP=-0.0190	PIQE=28.58 NIQE=4.287 CLIP=0.3789	PIQE=31.24 NIQE=4.821 CLIP=0.3728	PIQE=31.87 NIQE=4.481 CLIP=0.3784	PIQE=83.12 NIQE=13.94 CLIP=0.0827
	A vegetable pizza on a table.				A person who is on the motorcycle.			
$\epsilon = 0.5$								
	PIQE=45.24 NIQE=5.404 CLIP=0.3792	PIQE=49.78 NIQE=5.995 CLIP=0.3707	PIQE=54.08 NIQE=4.439 CLIP=0.3321	PIQE=87.71 NIQE=14.93 CLIP=-0.0452	PIQE=48.74 NIQE=4.389 CLIP=0.3371	PIQE=54.29 NIQE=4.482 CLIP=0.331	PIQE=28.61 NIQE=3.599 CLIP=0.3884	PIQE=88.64 NIQE=14.58 CLIP=0.0359

Figure 3: Generated image comparison under the security mechanism. Images in each sample from left to right are Watermarked Initial Latent Variables (WIL for short) without VAE-based fine-tuning (fine-tuning for short), WIL with fine-tuning, non-WIL without fine-tuning, and non-WIL with fine-tuning, representatively.

$$\mathcal{L}_m = - \sum_{k=0}^{n-1} m_k \log \sigma(m'_k) + (1 - m_k) \log(1 - \sigma(m'_k)) \quad (5)$$

Since the training samples of the diffusion model are generated by progressively adding noise until they conform to a standard normal distribution, during the inference stage, the message encoder will output initial latent variables that follow the same distribution. The Kullback-Leibler (KL) divergence between the initial latent variables and the standard normal distribution is utilized as the loss function. The output follows a normal distribution, denoted as $q(z) \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and the standard normal distribution is denoted as $p(z) \sim \mathcal{N}(0, 1)$. The distribution loss is as follows:

$$\mathcal{L}_d = D_{KL}(q(z)||p(z)) = \int_x q(x) \log \frac{q(x)}{p(x)} dx \quad (6)$$

3.4 Watermark Robustness Enhancement Module

Watermarking Verification. Watermark extraction involves diffusion inversion, an approximate process for obtaining initial hidden variables from generated images. Diffusion inversion [Dhariwal and Nichol, 2021] algorithmically retrieves the initial latent variables from images generated by a diffusion model. x_t represents the image at the timestep t . Based on the assumption $x_{t-1} - x_t \approx x_{t+1} - x_t$, diffusion inversion of the Denoising Diffusion Implicit Model (DDIM) [Song *et al.*, 2020a] is formalized as follows:

$$\hat{x}_{t+1} = \sqrt{\bar{\alpha}_{t+1}}x_0 + \sqrt{1 - \bar{\alpha}_{t+1}}\epsilon_\theta(x_t) \quad (7)$$

where $\bar{\alpha}$ is the parameter of the diffusion model. t denotes the denoising timestep. $\epsilon_\theta(x_t)$ is the estimated noise for timestep t . \hat{x}_0 represents the prediction of the image at the current timestep and is defined as:

$$\hat{x}_0 = \frac{x_t - \sqrt{1 - \bar{\alpha}_t}\epsilon_\theta(x_t)}{\sqrt{\bar{\alpha}_t}}. \quad (8)$$

To mitigate the effects of diffusion inversion and raise the robustness of image processing, we introduce the watermark robustness enhancement module which employs adversarial training to raise performance of the message decoder. The loss function is binary cross entropy between the message m and the sigmoid $\delta(m')$ which is the same as Equation 5.

Attack Simulation for Adversarial Training. Various attacks are common in practical image usage. Therefore, during the training process, we deploy an attack layer to watermarked images before employing a watermark extraction algorithm. This attack layer encompasses six common types of attacks: blur, Gaussian noise, brightness adjustment, contrast adjustment, saturation adjustment, and JPEG compression. To remain the differentiable of attack during training, we employ the differentiable simulation method to perform JPEG attack [Zhu *et al.*, 2018].

4 Experiments

Implementation details. In this paper, we focus on text-to-image generation, hence we utilized Stable Diffusion-v2 [Rombach *et al.*, 2022]. The number of inference steps is 25 for both generation and detection process. Following the settings of existing works [Wen *et al.*, 2024], we employ

Methods	Metrics							
	TPR(C.) \uparrow	TPR(Adv.) \uparrow	Bit Acc.(C.) \uparrow	Bit Acc.(Adv.) \uparrow	FID \downarrow	CLIP \uparrow	PSNR \uparrow	SSIM \downarrow
DwtDct	0.832	0.128	0.903	0.554	3.38	0.334	39.2	0.974
DwtDctSvd	1.000	0.236	0.999	0.661	9.44	0.332	39.0	0.982
RivaGan	1.000	0.714	0.999	0.829	15.3	0.333	40.5	0.980
Tree-Ring	1.000	0.995	—	—	24.6	0.336	—	—
Stable Signature	1.000	0.837	0.989	0.812	13.4	0.335	29.6	0.824
FSwatermark	1.000	0.914	0.999	0.872	21.7	0.334	31.9	0.897
DistriMark (ours)	1.000	0.989	0.983	0.939	14.6	0.334	30.8	0.856

Table 1: Watermark detection and traceability comparison. ‘C.’ refers to results without image processing attacks. ‘Adv.’ (Adversarial) refers to the average performance of a series of image processing attacks.

Method	1	2	3	4	5	6	7	8	9	10	Comb.
DwtDct	0.587	0.929	0.497	0.479	0.598	0.512	0.493	0.498	0.524	0.494	0.507
DwtDctSvd	0.622	0.999	0.634	0.675	0.997	0.516	0.506	0.554	0.634	0.512	0.512
RivaGan	0.976	0.926	0.968	0.981	0.999	0.902	0.590	0.858	0.611	0.632	0.588
FSWatermark	0.996	0.998	0.995	0.979	0.657	0.997	0.913	0.664	0.686	0.643	0.561
Stable Signature	0.977	0.989	0.975	0.851	0.612	0.888	0.767	0.532	0.586	0.497	0.498
DistriMark (ours)	0.949	0.976	0.972	0.955	0.952	0.961	0.952	0.953	0.942	0.949	0.844

Table 2: Bit accuracy of ten image processing attack. (1) Brightness, (2) Gauss Noise, (3) Contrast, (4) Blur, (5) JPEG, (6) BM3D denoising algorithm, (7) Resize, (8-9) VAE-based compression algorithm and (10) diffusion-based reconstructive attack, respectively. ‘Comb.’ indicates a mixture of the previous attacks.

the prompts from StableDiffusionDB [Wang *et al.*, 2022] with guidance scale of 5 during inference and an empty prompt during DDIM inversion. We utilize AdamW with a learning rate of 5×10^{-4} and weight decay of 0.01 during finetuning. All experiments are conducted on a single NVIDIA L40.

Watermarking baselines. We select six typical baselines: three official watermark of Stable Diffusion [Rombach *et al.*, 2022] for cloud services called DwtDct [Cox *et al.*, 2007], DwtDctSvd [Cox *et al.*, 2007], and RivaGAN [Zhang *et al.*, 2019], two multi-bit watermarking methods named FSwatermark [Xiong *et al.*, 2023] and Stable Signature [Fernandez *et al.*, 2023], and a fine-tuning-free semantic watermarking method called Tree-Ring [Wen *et al.*, 2024].

Evaluation metrics. We measure the detection performance by the true positive rate (TPR) when the false positive rate (FPR) is at 1%. We measure the traceability performance by the bit accuracy. To measure the image generation quality, we compute the Peak Signal-to-Noise Ratio (PSNR) [Hore and Ziou, 2010] and Structural Similarity score (SSIM) [Wang *et al.*, 2004] for image distortion evaluation, the Fréchet Inception Distance (FID) [Heusel *et al.*, 2017] and CLIP score [Radford *et al.*, 2021] for image diversity and semantic evaluation, and Natural Image Quality Evaluator (NIQE) [Mittal *et al.*, 2012] and Perceptual Image Quality Evaluator (PIQE) [Venkatanath *et al.*, 2015] for image quality evaluation.

4.1 Watermark Security against Escape

In order to make the watermark flexible for distributing the LDMs to a large number of model users with strong robustness, we leverage the semantic watermarking framework to inject the watermark message into the latent variable m . However, the model users can easily escape the watermark by replacing m with other random latent variables to obtain

the non-watermarked generated images. To tackle this issue, we design the security mechanism to decrease the generated image quality when the model user escape the watermark. Representative non-watermarked images under the security mechanism are shown in Figure 3. Embedding watermark does not significantly affect the image quality metrics NIQE and PIQE, or semantic quality metric CLIP. As the quality of unauthorized images decreases further, the quality of watermarked images also slightly deteriorates.

4.2 Watermark Performance Comparison

We compare the performance of our method with existing six typical baselines over two tasks: (1) **Detection**. We consider all compared methods as single-bit watermarks with a unified watermark. We set the FPR to be 1% and test the TPR on 1,000 watermarked images. (2) **Traceability**. Each compared method, excepting for the single-bit watermark Tree-Ring, serves as a multi-bit watermark. In our experiments, we assume that there are 1,000 model users, each of them requires one watermark for model tracing. Each user generates 10 images, resulting in a dataset of 10,000 watermarked images. During test, if an image contains a watermark, we then calculate the number of matched bits (Bit Accuracy) with the watermark of each user. The user with the highest Bit Accuracy is considered the traced user and verified. The comparison results are shown in Table 1. Our watermarking achieves strong robustness and significantly outperforms baselines in both tasks. In terms of bit accuracy, it surpasses the best-performing baseline by approximately 6.7%. This can be attributed to the extensive diffusion of the watermark throughout the entire latent space, establishing a profound binding between the watermark and the image semantics.



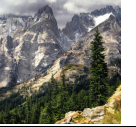
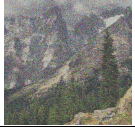
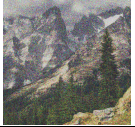
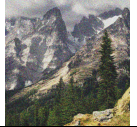
Attack	Reconstructive Attack			Adaptive Adversarial Sample Attack		
Method	DistriMark	FSwatermark	Stable Signature	DistriMark	FSwatermark	Stable Signature
Adversarial Samples						
Perturbation	0.142	0.138	0.153	0.211	0.202	0.109
Bit Accuracy	0.924	0.608	0.623	0.795	0.796	0.798

Figure 4: Adversarial samples obtained from adaptive adversarial sample attack and reconstructive attack.

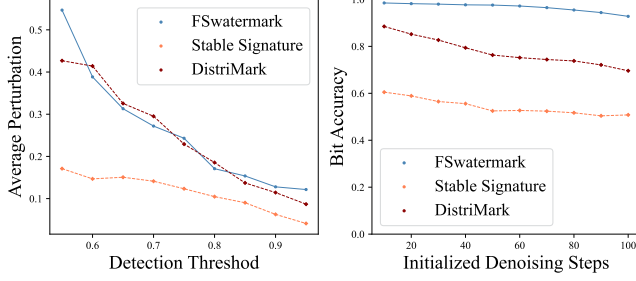


Figure 5: Results against adaptive adversarial sample attack (Left) and Reconstructive attack (Right).

4.3 Robustness against Image-Level Attacks

We examine watermark’s robustness against three typical kinds of adversarial attacks, including image processing attack [Song *et al.*, 2010] for transforming generated images, adaptive adversarial sample attacks [Jiang *et al.*, 2023] for disturbing watermark verification, and reconstructive attacks [Ballé *et al.*, 2018; Cheng *et al.*, 2020; Zhao *et al.*, 2023a] for re-generating non-watermarked images.

Image Processing Attack. We select ten representative types of image-level noise shown in Table 2. Please refer to the Supplementary Materials for detailed parameter settings.

Adaptive Adversarial Sample Attack. To further enhance the attack ability, we assume that attackers can query a black-box watermark verification interface and conduct query-based black-box attack [Jiang *et al.*, 2023]. By iterative querying the verification interface this attack compute optimal perturbations that progressively bring the watermark-free initial image closer to the original image.

Reconstructive Attack. The core idea of the reconstructive attack is to add random noise to destroy the watermark and then reconstruct the image. We utilize the implementation of the paper [Zhao *et al.*, 2023a] with denoising steps of 60.

Main Results. For each image processing attack, we report the average bit accuracy in Table 2. We see that our DistriMark watermark is indeed robust across all the transformations with the bit accuracy all above 0.9. DistriMark remarkably outperforms existing multi-bit watermarks on the average performance with more than 0.25 boost compared with the state-of-the-art (SoTA) results. Compared to the SoTA finetuning methods FSWatermark and Stable Signature, DistriMark has significant advantages in image resize, VAE-based compression algorithm, and reconstructive at-

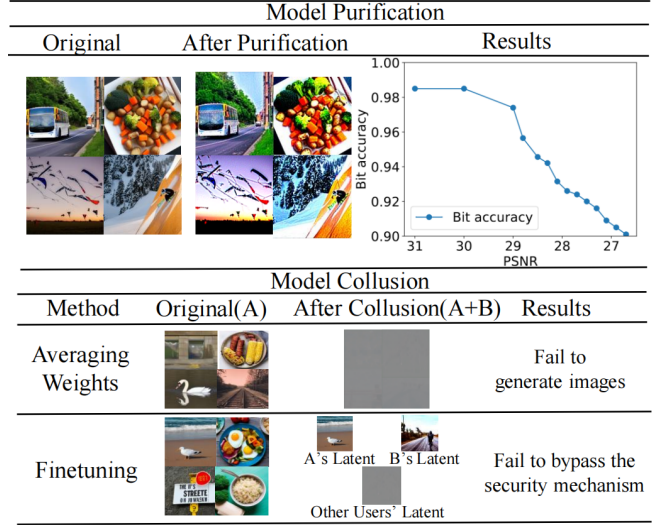


Figure 6: Results of model-level attacks.

tack. Adaptive Samples Attack utilizes the same evaluation metric ℓ_∞ -norm as described in the paper [Jiang *et al.*, 2023]. Figure 4 and Figure 5 display images and related parameters generated by Reconstructive Attack and Adaptive Adversarial Samples Attack under the same parameter conditions. Under both types of attacks, DistriMark demonstrates better robustness than the other two methods. The robustness stems from the watermark being embedded at the semantic level within the image, so more extensive attacks are needed at the pixel level to remove the watermark. Please refer to the Supplementary Materials for more results.

4.4 Robustness against Model-level Attacks

Consider the model-level attacks of the model for both single users and multiple users, we have included two types of model-level attacks: model purification and model collusion.

Model purification. The adversary fine-tunes the Variational autoencoder to circumvent watermark embedding through the same training mode as Section 3. This involves removing the message loss L_m , and shifting the focus to the perceptual loss L_w between the original image and the one reconstructed by the LDM auto-encoder.

Model Collusion. We mainly considered two types of collusion attacks: Averaging weights and Finetuning. (1) Averaging weights. User⁽ⁱ⁾ and User^(j) can average their weights

Method	PSNR	SSIM	FID	NIQE	PIQE
DwtDct	39.2	0.974	3.38	3.79	32.7
DwtDctSvd	39.0	0.982	9.44	3.79	32.9
RivaGan	40.5	0.980	15.3	3.82	32.4
Tree-ring	—	—	25.9	4.25	33.5
FSWm.	31.9	0.897	21.7	4.22	34.7
Stable Sig.	29.6	0.864	13.4	3.79	33.7
DistriMark	30.8	0.856	14.6	3.98	34.2

Table 3: Quality comparison of watermarked generated images.

Skip Connect	Connect Strength	Image Quality		Bit Acc.
		Watermark	Random	
Without	0.1	31.7	30.3	0.986
	0.3	25.7	24.6	0.984
	0.5	24.1	23.7	0.979
Single	0.1	29.6	24.5	0.981
	0.3	25.8	17.2	0.979
	0.5	23.9	12.4	0.978
Multiple	0.1	30.8	14.7	0.985
	0.3	29.4	14.1	0.982
	0.5	26.3	12.1	0.982

Table 4: Evaluation on impact of skip connection.

like Model Soup [Wortsman *et al.*, 2022] to create a new model to deceive identification. (2) Finetuning. Another form of collusion attack is when the user B generates a large number of watermarked latent variables and watermarked generated images, and fine-tunes the VAE of A so that A can use B’s watermark latent variables to generate images.

Main results. Figure 6 shows the results of model purification attack. As for model purification, when the bit accuracy decreases, the image quality also declines. It is difficult to significantly reduce the bit accuracy without affecting the image quality. As for model collusion, because of the security mechanism, parameter averaging will cause a significant drop in image quality. As for finetuning, this could pose a threat of identity spoofing. However, this still does not break the security mechanism.

4.5 Watermarked Image Quality

Besides the qualitative examples of how the watermarked images are not sensitive for the human eyes to distinguish (see Figure 3), we further present quantitative evaluation of images generated by existing watermarking methods in Table 3. The results show that no matter the qualitative metrics, our DistriMark achieves comparable performance with existing works in the model distribution scenarios. For DistriMark, the initial watermark is only manifested in the selection of the initial latent variables, hence the image quality is consistent with the diffusion model inference. Although DistriMark shows comparable quality compared to the methods in the model distribution scenario, the method involves a trade-off between the quality of unauthorized images and watermarked images when fine-tuning the VAE-Decoder, which results in lower image quality compared to post-processing methods.

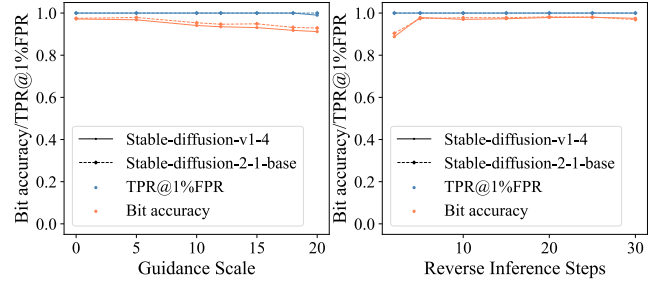


Figure 7: Ablation study results. (Left) Impact of guidance scales. (Right) Impact of reverse inference steps.

4.6 Ablation Study

Skip connection selection. The way latent variables are connected to the VAE-Decoder impacts how the VAE transforms images from the latent space to the pixel space. In Table 4, three methods were tested: no connection, single connection, and multi-level connection. Multi-level connections improve feature learning and enhance image quality.

Guidance scales. Larger guidance scales result in more faithful of the generated image adherence to prompts. Following existing works [Wen *et al.*, 2024], we cover the range of 0 to 20. In Figure 7 (left), the watermark remains robust and reliable even at a guidance scale of 18.

Number of the inversion step. The inference step is often unknown in practice, which introduces a mismatch with the inversion step. From Figure 7 (right) we can see that the number of inference steps does not significantly affect the accuracy of inversion which is beneficial in practice.

5 Conclusion

We propose a novel distribution scenario-oriented watermarking schema for diffusion models and a new security mechanism to prevent watermark leakage and watermark escape in the model distribution scenarios. We separate the watermark injection from the security mechanism, ensuring that fine-tuning the VAE focuses solely on the security mechanism without the added task of learning watermark patterns. In the future, our research directions will include adversarial methods against forge attack [Lukas *et al.*, 2023], security mechanism with higher image quality and more consideration about model distribution scenarios.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant No.s U24B20146, 62372044), and Beijing Municipal Science and Technology Commission Project (Z241100009124008).

Contribution Statement

This work was a collaborative effort by all contributing authors. Liangqi Lei and Keke Gai made equal contribution to this study and are designated as co-first authors. Jing Yu serving as the corresponding author, is responsible for all communications related to this manuscript.

References

- [An *et al.*, 2024] Bang An, Mucong Ding, Tahseen Rabhani, Aakriti Agrawal, Yuancheng Xu, Chenghao Deng, Sicheng Zhu, Abdirisak Mohamed, Yuxin Wen, Tom Goldstein, et al. Benchmarking the robustness of image watermarks. *arXiv preprint arXiv:2401.08573*, 2024.
- [Azcoitia and Laoutaris, 2022] Santiago Andrés Azcoitia and Nikolaos Laoutaris. A survey of data marketplaces and their business models. *ACM SIGMOD Record*, 51(3):18–29, 2022.
- [Ballé *et al.*, 2018] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018.
- [Barrett, 2023] Clark Barrett. Identifying and mitigating the security risks of generative ai. *arXiv preprint arXiv:2308.14840*, 2023.
- [Cheng *et al.*, 2020] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7939–7948, 2020.
- [Ci *et al.*, 2025] Hai Ci, Pei Yang, Yiren Song, and Mike Zheng Shou. Ringid: Rethinking tree-ring watermarking for enhanced multi-key identification. In *European Conference on Computer Vision*, pages 338–354. Springer, 2025.
- [Cox *et al.*, 2007] Ingemar Cox, Matthew Miller, Jeffrey Bloom, Jessica Fridrich, and Ton Kalker. *Digital watermarking and steganography*. Morgan kaufmann, 2007.
- [Croitoru *et al.*, 2023] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.
- [Dhariwal and Nichol, 2021] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [Donahue and Kleinberg, 2021] Kate Donahue and Jon Kleinberg. Model-sharing games: Analyzing federated learning under voluntary participation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5303–5311, 2021.
- [Feng *et al.*, 2024] Weitao Feng, Wenbo Zhou, Jiyan He, Jie Zhang, Tianyi Wei, Guanlin Li, Tianwei Zhang, Weiming Zhang, and Nenghai Yu. Aqualora: Toward white-box protection for customized stable diffusion models via watermark lora. *arXiv preprint arXiv:2405.11135*, 2024.
- [Fernandez *et al.*, 2023] Pierre Fernandez, Guillaume Couairon, Hervé Jégou, Matthijs Douze, and Teddy Furon. The stable signature: Rooting watermarks in latent diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22466–22477, 2023.
- [Gowal and Kohli, 2023] Sven Gowal and Pushmeet Kohli. Identifying ai-generated images with synthid. <https://www.deepmind.com/blog/identifying-ai-generated-images-with-synthid>, 2023. Accessed: 2023-09-23.
- [Heusel *et al.*, 2017] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [Hore and Ziou, 2010] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010.
- [Jiang *et al.*, 2023] Zhengyuan Jiang, Jinghui Zhang, and Neil Zhenqiang Gong. Evading watermark based detection of ai-generated content. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 1168–1181, 2023.
- [Lukas *et al.*, 2023] Nils Lukas, Abdulrahman Daa, Lucas Fenaux, and Florian Kerschbaum. Leveraging optimization for adaptive attacks on image watermarks. *arXiv preprint arXiv:2309.16952*, 2023.
- [Mittal *et al.*, 2012] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2012.
- [Pei *et al.*, 2023] Jian Pei, Raul Castro Fernandez, and Xiaohui Yu. Data and ai model markets: Opportunities for data and model sharing, discovery, and integration. *Proceedings of the VLDB Endowment*, 16(12):3872–3873, 2023.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021.
- [Rombach *et al.*, 2022] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [Song *et al.*, 2010] Chunlin Song, Sud Sudirman, Madjid Merabti, and David Llewellyn-Jones. Analysis of digital image watermark attacks. In *2010 7th IEEE Consumer Communications and Networking Conference*, pages 1–5. IEEE, 2010.
- [Song *et al.*, 2020a] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

- [Song *et al.*, 2020b] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [Venkatanath *et al.*, 2015] N Venkatanath, D Praneeth, Maruthi Chandrasekhar Bh, Sumohana S Channappayya, and Swarup S Medasani. Blind image quality evaluation using perception based features. In *2015 Twenty First National Conference on Communications*, pages 1–6. IEEE, 2015.
- [Wang *et al.*, 2004] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [Wang *et al.*, 2022] Zijie J Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. *arXiv preprint arXiv:2210.14896*, 2022.
- [Wen *et al.*, 2023] Yuxin Wen, John Kirchenbauer, Jonas Geiping, and Tom Goldstein. Tree-ring watermarks: Fingerprints for diffusion images that are invisible and robust. *arXiv preprint arXiv:2305.20030*, 2023.
- [Wen *et al.*, 2024] Yuxin Wen, John Kirchenbauer, Jonas Geiping, and Tom Goldstein. Tree-rings watermarks: Invisible fingerprints for diffusion images. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Wortsman *et al.*, 2022] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR, 2022.
- [Xiong *et al.*, 2023] Cheng Xiong, Chuan Qin, Guorui Feng, and Xinpeng Zhang. Flexible and secure watermarking for latent diffusion model. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1668–1676, 2023.
- [Yang *et al.*, 2024] Zijin Yang, Kai Zeng, Kejiang Chen, Han Fang, Weiming Zhang, and Nenghai Yu. Gaussian shading: Provable performance-lossless image watermarking for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12162–12171, 2024.
- [Zhang *et al.*, 2019] Kevin Alex Zhang, Lei Xu, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Robust invisible video watermarking with attention. *arXiv preprint arXiv:1909.01285*, 2019.
- [Zhang *et al.*, 2024] Lijun Zhang, Xiao Liu, Antoni Viros Martin, Cindy Xiong Bearfield, Yuriy Brun, and Hui Guan. Robust image watermarking using stable diffusion. *arXiv preprint arXiv:2401.04247*, 2024.
- [Zhao *et al.*, 2023a] Xuandong Zhao, Kexun Zhang, Zihao Su, Saastha Vasan, Ilya Grishchenko, Christopher Kruegel, Giovanni Vigna, Yu-Xiang Wang, and Lei Li. Invisible image watermarks are provably removable using generative ai. *arXiv preprint arXiv:2306.01953*, 2023.
- [Zhao *et al.*, 2023b] Yunqing Zhao, Tianyu Pang, Chao Du, Xiao Yang, Ngai-Man Cheung, and Min Lin. A recipe for watermarking diffusion models. *arXiv preprint arXiv:2303.10137*, 2023.
- [Zhu *et al.*, 2018] Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei. Hidden: Hiding data with deep networks. In *Proceedings of the European Conference on Computer Vision*, pages 657–672, 2018.