

A Crack in the Bark: Leveraging Public Knowledge to Remove Tree-Ring Watermarks

Junhua Lin Marc Juarez
University of Edinburgh

Abstract

We present a novel attack specifically designed against Tree-Ring, a watermarking technique for diffusion models known for its high imperceptibility and robustness against removal attacks. Unlike previous removal attacks, which rely on strong assumptions about attacker capabilities, our attack only requires access to the variational autoencoder that was used to train the target diffusion model, a component that is often publicly available. By leveraging this variational autoencoder, the attacker can approximate the model’s intermediate latent space, enabling more effective surrogate-based attacks. Our evaluation shows that this approach leads to a dramatic reduction in the AUC of Tree-Ring detector’s ROC and PR curves, decreasing from 0.993 to 0.153 and from 0.994 to 0.385, respectively, while maintaining high image quality. Notably, our attacks outperform existing methods that assume full access to the diffusion model. These findings highlight the risk of reusing public autoencoders to train diffusion models—a threat not considered by current industry practices. Furthermore, the results suggest that the Tree-Ring detector’s precision, a metric that has been overlooked by previous evaluations, falls short of the requirements for real-world deployment.

1 Introduction

Recent advances in generative AI have led to the development of text-to-image models capable of creating highly realistic content. Some of these, such as Midjourney and DALL-E, are proprietary models that produce images nearly indistinguishable from those generated by humans [31, 36]. As these models become more sophisticated, discerning between AI-generated and authentic content becomes increasingly difficult, raising serious concerns about the risks of misuse, such as the potential to create *deepfakes* and spread misinformation [9].

In response, the research community has turned its attention to developing watermarking techniques that the owners of the

models can deploy to ensure that AI-generated content can be reliably identified. If these watermarks are robust against removal, they can prevent malicious actors from passing off deepfakes as authentic, thereby helping to mitigate some of the risks associated with the misuse of these technologies. Recent examples of watermarking schemes developed for this purpose are Stegastamp [54], Stable Signature [23], and Tree-Ring [56], with Tree-Ring standing out for the imperceptibility of its watermarks in the generated images and its robustness against removal attacks.

Tree-Ring was specifically designed for diffusion models. It achieves high imperceptibility by encoding a fixed string into the diffusion model’s initial latent variable, rather than into the image directly. Embedding a fixed watermark constrains the latent space to a region, altering the model’s output distribution; however, individual images generated from the watermarked region are clean images, free of unusual patterns or artifacts. To verify a watermarked image, Tree-Ring must accurately approximate the inversion of the diffusion process to recover the latent variable that generated the image and confirm the presence of the watermark.

Evaluations of Tree-Ring have shown that attacks based on simple image manipulations, such as rotation and cropping, fail to remove the watermark without substantially modifying the image [56]. Consequently, the attacked images cannot serve as substitutes for the watermarked ones, losing their value from an attacker’s perspective. However, these initial evaluations did not explore attacks that expressly target the latent space, where the watermarks are embedded.

Subsequent evaluations include more advanced attacks tailored to Tree-Ring’s latent-space embedding mechanism [4, 37, 63]. These attacks are formulated as a learning problem aimed at shifting the latent away from the watermarked region while bounding image quality degradation. Although some of these attacks are more effective than the simple attacks evaluated by Wen et al., they have been criticized for making strong assumptions about the adversary’s knowledge of the target diffusion model [4]. Some attacks even require white-box access to the diffusion model, which would eliminate

the need for an attack, as the adversary could simply use the model to generate their own non-watermarked images.

We demonstrate how an attacker can use publicly available auxiliary models to overcome some of the assumptions made in prior work. In particular, we present an attack inspired by An et al.’s surrogate-detector approach [4]. Following their attack strategy, we train a model that mimics Tree-Ring’s detector and find adversarial examples on it that transfer to the original detector. However, rather than training the detector on the pixel representation of images, we train it on intermediate latents approximated via a publicly available variational autoencoder (VAE). Since VAEs are commonly used to reduce the training cost of diffusion models, they are often shared and reused, making them readily available to attackers.

Our VAE-enhanced surrogate detector attacks outperform previous attacks, including those of An et al., drastically decreasing Tree-Ring’s detection performance while maintaining image quality. This improvement in attack performance highlights the risks associated with sharing and reusing VAEs for training diffusion models, a practice that is currently overlooked in the industry (for example, see [2, 8, 10, 14, 18, 34, 40, 49, 52]).

More specifically, our primary contributions are:

An in-depth analysis of Tree-Ring’s embedding mechanism. We identify critical flaws in Tree-Ring’s embedding mechanism, motivating the design of our attacks. The embedding mechanism violates the assumption that the latents are drawn from a Gaussian distribution. This, coupled with the persistence of the watermarks through the diffusion process exposes discriminating features that enable the training of a surrogate detector in the latent space.

Novel VAE-based surrogate attacks. We leverage public VAEs to attack Tree-Ring watermarking. Assuming access to a VAE is realistic because, in practice, a pretrained autoencoder is often used to train diffusion models [44], but we also challenge this assumption by evaluating the attacks using a different VAE. Our attacks successfully remove Tree-Ring watermarks, decreasing the detector’s TPR@1%FPR from 0.968 to 0.04 and its ROC-AUC from 0.993 to 0.15, while retaining high image quality. This is nearly a 5-fold decrease in ROC-AUC over the state-of-the-art surrogate detector attack.

An evaluation of detector precision. In contrast to previous evaluations of watermarking schemes, we measure Tree-Ring’s detection precision under base rates that capture a range of deployment scenarios. In a balanced setting, our attack reduces the PR-AUC from 0.994 to 0.385. The impact is even more pronounced at lower, more realistic base rates, suggesting that Tree-Ring’s precision may be insufficient for the intended use cases it aims to support.

2 Background and Related Work

This section provides the necessary background on diffusion and Tree-Ring watermarking to follow the rest of the paper.

2.1 Diffusion Models

A diffusion model [50] is based on a *forward diffusion process*, which involves gradually adding noise to a data point until the signal is destroyed. By learning how to reverse this process, a diffusion model estimates the underlying data distribution and provides a method for sampling new data points from it.

More formally, let x_0 be a data point drawn from the true data distribution $q(x_0)$. The forward diffusion process is a Markov chain defined by a sequence of T intermediate latent variables x_1, \dots, x_T in the same domain as x_0 , where each x_t is obtained by adding noise drawn from the noise distribution $q(x_t | x_{t-1})$ to the previous latent, x_{t-1} .

The diffusion model is a distribution $p_\theta(x_0)$, parametrized by θ , that reverses the forward process thus providing a method to sample from $q(x_0)$ and generate new data points:

$$p_\theta(x_0) = \int p_\theta(x_{0:T}) dx_{1:T}, \quad (1)$$

$$\text{where } p_\theta(x_{0:T}) := p_\theta(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t).$$

The learning objective is to find the parameters that maximize the variational lower bound on the expected log-likelihood:

$$\mathbb{E}_{q(x_0)} [\log p_\theta(x_0)] \geq \mathbb{E}_{q(x_0, \dots, x_T)} \left[\log \frac{p_\theta(x_{0:T})}{q(x_{1:T}|x_0)} \right]. \quad (2)$$

Ho et al. presented Denoising Diffusion Probabilistic Models (DDPMs) by specifying the forward diffusion process to use a fixed predefined Gaussian distribution with a variance schedule $\{\alpha_t \in (0, 1]\}_{t=1}^T$ [26], as follows:

$$q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}), \quad (3)$$

$$\text{with } q(x_t|x_{t-1}) := \mathcal{N}\left(x_{t-1}\sqrt{\frac{\alpha_t}{\alpha_{t-1}}}, \left(1 - \sqrt{\frac{\alpha_t}{\alpha_{t-1}}}\right)I\right).$$

Given this, we can express the intermediate latent variables in a closed form:

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, \quad (4)$$

which shows that x_t is a linear combination of the initial point and a Gaussian noise variable $\epsilon \sim \mathcal{N}(0, I)$. Equation 4 allows DDPM to reparameterize the learning objective to learn a predictor $\epsilon_\theta(x_t)$ of the noise to be removed to recover x_0 .

Extending upon this, Denoising Diffusion Implicit Models (DDIM) [51] propose to replace DDPM’s Markov chain with

a non-Markovian forward process, where each step depends also on the initial data point, enabling a deterministic and faster sampling mechanism compared to the original DDPM approach. The DDIM denoising process, which transforms an initial Gaussian noise vector x_T to an image x_0 is given by:

$$x_0^{(t)} = \frac{x_t - \sqrt{1 - \alpha_t} \epsilon_\theta(x_t)}{\sqrt{\alpha_t}}, \quad (5)$$

where $x_0^{(t)}$ denotes the estimation of x_0 from time step t . A single-step denoising process can thus be expressed as:

$$x_{t-1} = \sqrt{\alpha_{t-1}} x_0^{(t)} + \sqrt{1 - \alpha_{t-1}} \epsilon_\theta(x_t). \quad (6)$$

This is a recursive process starting from time step T : as t decreases, $x_0^{(t)}$ becomes a more accurate estimate of x_0 . Song et al. denote this recursive denoising process by $\mathcal{D}_\theta(x_T)$ [51].

2.2 Latent Diffusion

Diffusion models are commonly implemented using a U-Net architecture [45] that learns how to approximate the noise variable ϵ . However, training a full diffusion model directly on high-resolution images is computationally demanding. To address this, a technique known as *latent diffusion* [44] is often used to compress the image into an intermediate latent representation, reducing memory consumption at training time. Once the noise predictor is trained on this lower-dimensional space, the output of the diffusion model must be mapped back to the image space to generate an image. We denote \mathcal{E} the compression function and \mathcal{E}^{-1} the function that maps from latent to image space. In practice, \mathcal{E} and \mathcal{E}^{-1} are often parameterized by a Variational Autoencoder (VAE). For example, the popular Stable Diffusion v2.1 uses AutoencoderKL [32].

2.3 Digital Watermarking

Digital watermarking emerged from the need to protect copyright as digital media became widespread and easy to distribute [55]. Early methods considered an adversary attempting to remove watermarks to misappropriate the content. To address this threat model, watermarking techniques often rely on steganography to embed watermarks directly into the content—unlike digital signatures, which are typically appended in the content’s metadata and can be stripped out. This inseparability from the content is a design goal of watermarking that distinguishes it from other authentication methods.

The rise of generative AI has introduced new use cases for watermarking. Research in this field has focused on the use of watermarks to protect the intellectual property of model owners [11] and dataset curators [21]. More recently, watermarking has been proposed as a means to reliably identify AI-generated content [62]. For image generation, such watermarking schemes aim to distinguish digital photographs and human-made art from artificially generated images.

An et al. [4] categorize image watermarking techniques into two main classes:

Post-processing techniques embed the watermark directly into the image. This includes traditional watermarking schemes based on signal-processing methods, such as those that encode the watermark in the image’s frequency domain [3, 16], as well as recent ML-based approaches that involve training models, such as autoencoders, to learn how to insert and subsequently extract an imperceptible watermark from an image [54, 60, 64].

In-processing watermarking schemes are embedded within the generation process. This typically requires retraining the model [38, 58, 59], fine-tuning it [23], or modifying its sampling procedure. An approach specific to diffusion models, first proposed by Tree-Ring, is to embed the watermark in the model’s latent space [24, 56, 57]. Latent-space watermarking has gained popularity due to its imperceptibility and the added difficulty it poses for adversaries: removal attacks require approximating the inversion of the diffusion process without knowledge of the model, as it is assumed to be proprietary.

Tree-Ring is a representative and meaningful target for study. As the first scheme to propose embedding the watermark in the model’s latent space, Tree-Ring has been highly influential, with subsequent schemes closely following its embedding mechanism. Given the growing adoption of latent-space watermarking, a detailed analysis of Tree-Ring has broader implications for the design and evaluation of future schemes. For example, our use of VAEs to mount attacks in the latent space highlights a broader potential vulnerability applicable to other latent-space watermarking schemes.

2.4 Tree-Ring Watermarking

The Tree-Ring watermarking scheme, proposed by Wen et al. relies on one key property of DDIM models [56]. Dhariwal et al. discovered that the DDIM sampling process can be inverted, mapping an image x_0 back towards the initial noise vector x_T [20]. Given a learned noise function ϵ_θ , we can re-estimate x_T with:

$$x_{t+1} = \sqrt{\alpha_{t+1}} x_0^{(t)} + \sqrt{1 - \alpha_{t+1}} \epsilon_\theta(x_t). \quad (7)$$

According to Dhariwal et al., this inversion process reasonably estimates the initial latent variable if two conditions hold: (i) the inversion process reintroduces noise in small steps, i.e., the difference between x_t and x_{t+1} is small; and, (ii) $x_{t-1} - x_t \approx x_{t+1} - x_t$. If these two hold, the inversion process can accurately reconstruct the initial noise vector, enabling the key detection scheme of Tree-Ring watermarking. Wen et al. denote this inversion process by $\mathcal{D}_\theta^\dagger(x_0)$.

Figure 1 illustrates the embedding process. A predefined key is encoded as a pattern in the central region of the Fourier space of the initial noise vector. This region, which contains the low-frequency components, corresponds to global image

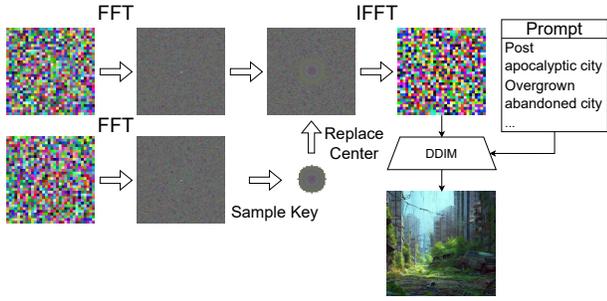


Figure 1: Embedding mechanism. The radius of the ring determines the number of pixels we sample from the frequency domain. The ring is embedded in the center of the initial Gaussian noise vector in Fourier space. After applying Inverse Fast Fourier Transform (IFFT), the resultant noise vector is fed to the diffusion model for image generation.

characteristics. Embedding a watermark here diffuses its information across these high-level features, making the watermark not only more imperceptible but also more robust to image manipulations such as rotation and cropping. As the watermark becomes entangled with the image’s semantic content, pixel-level operations that remove it inevitably destroy the image’s meaning, rendering the image unusable from the adversary’s perspective. Wen et al. experimented with various key patterns and found that the most robust to rotations is a circular pattern composed of multiple rings, each ring with a constant value sampled from the frequency domain of a Gaussian vector. To generate a new image, the modified vector is transformed back to the latent domain via IFFT and denoised into an image with standard DDIM backward diffusion.

Detecting the watermark is more challenging, especially for conditional models that take a text prompt as input. For perfect reconstruction, both the output image and the original prompt are needed, as the prompt conditions the diffusion process to generate an image that aligns with the user’s intent. However, for privacy reasons, model owners must allow users to delete their prompt history. Consequently, DDIM inversion without access to the original prompt can only approximate the initial noise vector (see Figure 2). That said, Wen et al. demonstrate that this approximation is sufficiently accurate to enable watermark detection even without the original prompt.

2.5 Previous Attacks

Tree-Ring’s security has been evaluated in several studies, including the original Tree-Ring paper [56]. This evaluation showed that Tree-Ring is robust to rotation, cropping, rescaling, and JPEG compression [56]. However, these image transformations lack fine-grained control in the latent space necessary to target the watermark without substantially altering the image. Next, we summarize follow-up works that have proposed more sophisticated attacks against Tree-Ring watermarking.

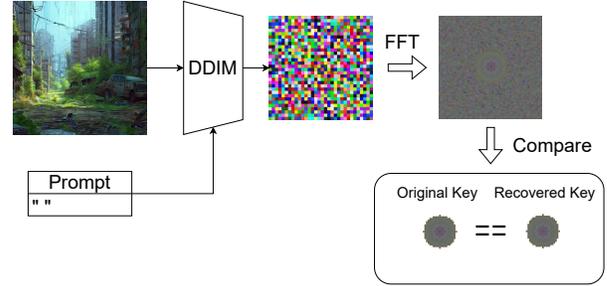


Figure 2: Detection scheme. Using forward diffusion, Tree-Ring estimates the initial noise vector used to generate the image with an empty prompt. It then extract the key in the center and thresholds its distance to the original key to make a final decision.

Regeneration attacks. Zhao et al. propose *regeneration attacks*, which re-encode the image with an off-the-shelf encoder, inject noise, and subsequently decode to reconstruct the image [63]. The authors found that these attacks are effective against pixel-level watermarking schemes like Stable Signature, but they have limited success against Tree-Ring. A more extensive evaluation by An et al. demonstrates that these attacks can be further tuned to increase their performance but then significantly degrade image quality [4].

Zhao et al. instantiate the encoder–decoder pair of a regeneration attack with public auxiliary models, such as a VAE or a diffusion model. These attacks are closest to our attacks in that they rely on a public auxiliary model. However, our attack strategy is more precise, as our attacks exploit weaknesses specific to the Tree-Ring embedding mechanism, rather than indiscriminately adding noise to the latents.

Adversarial attacks. In the most extensive survey of watermark removal attacks to date, An et al. also explore *adversarial attacks*, which are characterized by crafting adversarial examples against the watermark detector. The success rate of adversarial attacks is overall lower than that of regeneration attacks [4], but they retain most of the original image, which is a desirable property for a removal attack. An et al. distinguish between two types of adversarial attacks: *surrogate detector* and *embedding* adversarial attacks.

Embedding attacks. They rely on an existing encoder to map images to their corresponding latent representation. The adversary uses the encoder to find a small perturbation of the image that maximizes the distance between original and perturbed latents [4]. By maximizing this distance, the attacker aims to push the latent representation of the perturbed image outside the “watermarked” region of the latent space. An et al. evaluated embedding attacks extensively and showed that they are only effective if the attack’s encoder exactly matches the one used by the victim’s diffusion model. Although this attack could be improved, embedding attacks do not target specific

parts of the embedding but rather randomly shift the entire embedding, causing significant changes in the final image.

Surrogate detector attacks. The adversary first trains a surrogate classifier to detect the presence of a watermark, mimicking the behavior of the actual watermark detector. Next, they aim to find a small perturbation in the image that causes the surrogate detector to misclassify it, in the hope that this misclassification will transfer to the watermark detector. Unlike embedding attacks, these attacks do not explicitly assume access to the victim’s encoder; however, gathering the data to train the surrogate can be challenging in practice.

An et al. explore various methods to obtain training data for the surrogate detector. One of the training procedures requires white-box access to the victim’s diffusion model to generate a dataset of image pairs, each pair consisting of a watermarked image and its non-watermarked counterpart. This is an unrealistic assumption: it implies full access to the victim’s model, eliminating the need for an attack, as the adversary could directly sample from the model to generate non-watermarked images. A more realistic setting is to train the surrogate classifier with Tree-Ring images and an auxiliary publicly-available image dataset. However, this attack did not evade the Tree-Ring detector [4]. An et al. hypothesize that the reason is the difference between the public dataset distribution and that of the generated images, which may have led to spurious features that are independent of the watermarks [4].

Adaptive attacks. Lukas et al. propose *adaptive attacks* [37], which can be seen as a combination of embedding and surrogate detector attacks [4]. The instantiation of these attacks for Tree-Ring relies on a less capable version of the victim’s model. To deploy the attack, the adversary can use this surrogate diffusion model in two ways: *adversarial noising*, where the adversary finds an adversarial example through a perturbation that maximizes the distance between decoded keys (embedding attack); and *adversarial compression*, which compresses the image while minimizing the LPIPS loss. Lukas et al. evaluate these attacks assuming that the adversary has a diffusion model with the same architecture as the victim’s. This is a strong assumption: if the victim’s model is proprietary, the adversary is unlikely to have access to a diffusion model with the same architecture, and it is unclear how much attack performance degrades if this assumption does not hold.

We present novel adversarial attacks with fewer assumptions on the adversary’s access and knowledge. In particular, we use an auxiliary model, such as a public VAE, to formulate An et al.’s surrogate detector attacks in the latent space, eliminating the need for a dataset of non-watermarked images.

3 Attack Methods

In this section, we define our threat model and describe our attack methods in detail.

3.1 Threat Model

We follow the threat model in previous watermarking studies, considering two parties: the *model owner* and the *adversary*. The model owner maintains a proprietary diffusion model and monetizes it by offering subscription-based access to prompt it (e.g., through an API). Upon receiving a prompt, the model owner uses Tree-Ring to embed a *key*, k , and generate a watermarked image. We only consider proprietary models because Tree-Ring watermarks are embedded during generation and thus cannot be enforced on open-source models.

Adversary’s goal. The adversary’s goal is to use the images generated by the hosted model for their own purposes without revealing that they were generated with *that* model. To achieve that, the adversary’s strategy is to apply an image transformation such that the watermark is no longer detected.

Although certain images might be more susceptible to attacks than others, the attacker’s aim is to identify a general transformation that succeeds for an arbitrary image, implying ineffectiveness of the watermarking scheme. To be more precise, the adversary’s objective is to find an image transformation \mathcal{F} such that, for any image x generated with the queried model, it satisfies two conditions:

- i. $\|k - \tilde{k}\|_1 > \tau$, where \tilde{k} is Tree-Ring’s approximation of the original key k and τ is the decision threshold, and
- ii. $\|\mathcal{F}(x) - x\|_\infty \leq \delta$, for a small $\delta > 0$.

Although the maximum norm is a common metric to bound an image perturbation, the specific metric will depend on the adversary’s use case and their tolerance to changes in the image. In Section 5, we quantify how our attacks achieve the above objectives by measuring the decrease in detector performance under multiple common image similarity metrics.

Adversary’s capabilities. Since the target model is proprietary, the adversary has only black-box access: they can query it but have no knowledge of its parameters. Specifically, they observe only the output image and the input prompt that produced it. Additionally, the adversary is assumed to have limited computational resources; if they had sufficient resources, they could train a diffusion model from scratch instead of resorting to an attack. We therefore define a watermarking scheme as *robust* if such a computationally bounded adversary cannot successfully remove the watermark without spending more resources than would be required to collect data and train a comparable diffusion model. However, the adversary can still leverage publicly available models, such as pretrained autoencoders, readily accessible on model-sharing platforms, such as Civitai [39] and Hugging Face [28].

Justification of the public VAE access assumption. Auxiliary models suitable for use in attacks are publicly available on model-sharing platforms. These platforms are driven by the need to make diffusion model development—such as training and fine-tuning—more cost-efficient and accessible. Since

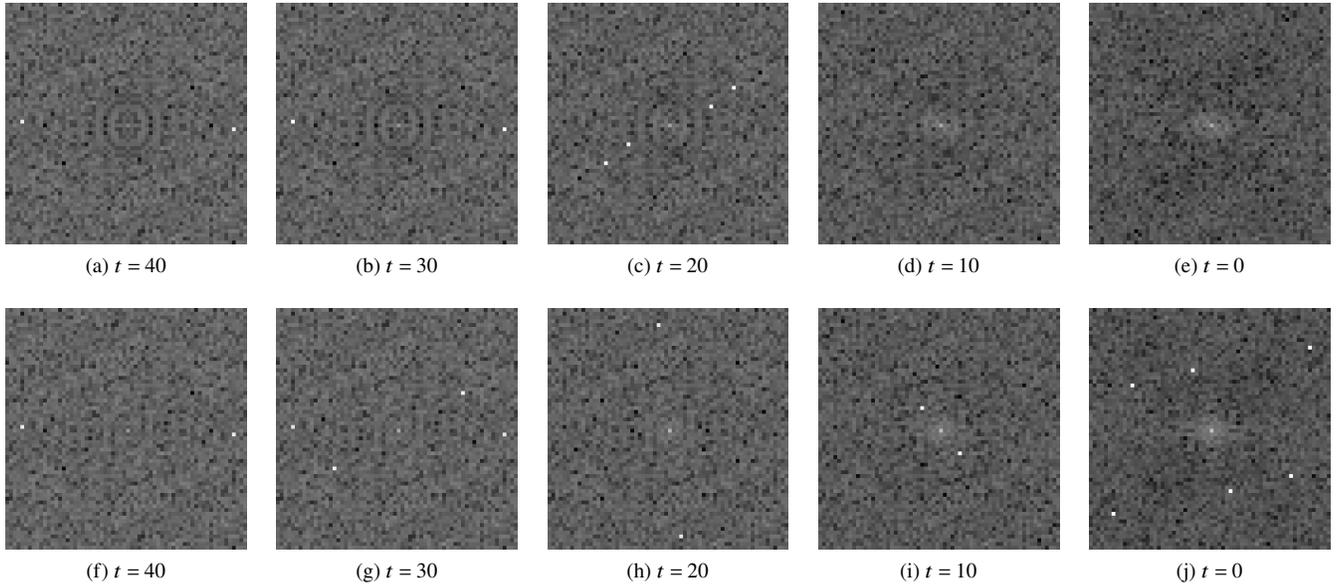


Figure 3: The Fourier transform (magnitude) of the fourth channel for the intermediate latents of the backward diffusion process. The first row of figures, (a)–(e), are the intermediate latents of an image with the watermark embedded. The second row, (f)–(j), are the intermediate latents generated with the same prompt but excluding the watermark. The plots are all in log scale and centered around the zero-frequency component. We omit the axes for emphasis.

training diffusion models from scratch is computationally expensive, practitioners often adopt latent diffusion, which reduces training costs by leveraging a pre-trained VAE. As a result, these VAEs are frequently reused and redistributed without modification.

Another common practice is to fine-tune existing diffusion models rather than retraining the full pipeline. For example, the most popular way to fine-tune a diffusion model is to apply LoRA methods [27] to fine-tune only the U-Net component, while leaving the VAE unchanged [46]. Consequently, even if the fine-tuned model is not published, the original VAE would remain public. We observe numerous real-world instances of these practices. For example, Civitai [39], a popular model-sharing platform, hosts a wide range of fine-tuned diffusion models that either retain the original VAE or continue training from publicly available checkpoints.

There is a plethora of examples of commercial diffusion models that rely on a VAE (see [1, 2, 8, 10, 14, 17, 18, 34, 40, 41, 49, 52]), and the trend of publishing the VAE used during training is also common in the industry. For instance, commercial text-to-image APIs such as those offered by Black Forest Labs [10], ShuttleAI [49], and BigModel [8] provide access to their VAEs. Even OpenAI has published the VAE for its text-to-image model, DALL-E [40]. These practices make access to the VAE associated with the target diffusion model a realistic and practical assumption. Moreover, even if the exact VAE is not publicly available, the abundance of VAEs trained for other diffusion models makes it likely that a suitable alternative is available.

3.2 Flaws of Tree-Ring Watermarking

The high accuracy of Tree-Ring’s detector despite the large amount of noise added in Wen et al.’s evaluation suggests that the approximation of the original initial latent is remarkably distinct from its non-watermarked counterpart. This inherent difference between watermarked and non-watermarked latents exposes a vulnerability in Tree-Ring that can be exploited.

Non-Gaussian nature of watermark latents. The embedding mechanism of Tree-Ring violates the assumption that the initial latent follows a Gaussian distribution. Recall that Tree-Ring draws both the initial latent vector and the values of the key ring from a Gaussian distribution. However, that does not imply that the combination of the two together will be normally distributed.

To demonstrate the difference between the watermarked and non-watermarked latent distributions, we examine the Fourier transform of the channel in which the key is embedded. Figure 3 shows the progression of the intermediate latent through the backward diffusion process, from the step when the key is embedded ($t = 40$) to the final latent ($t = 0$). As shown in the figure, a ring structure is clearly visible in the first diffusion steps and is gradually removed after each step until, in the final diffusion step ($t = 0$), the watermark is almost imperceptible. However, remnants of the original key structure remain visible at the center of the latent representations.

After applying backward diffusion, images are obtained by applying a transformation function parameterized by a neural network (e.g., a VAE) that has been pretrained to generate

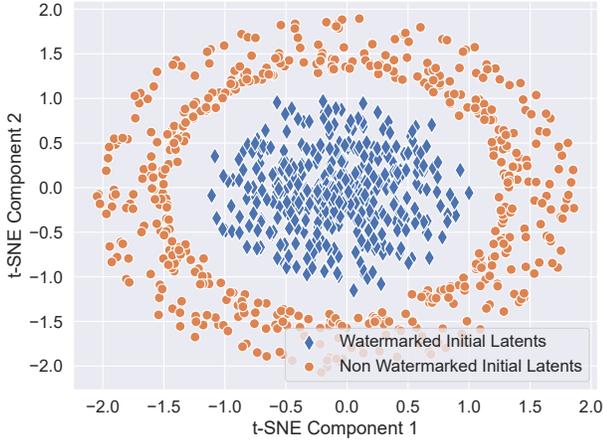


Figure 4: Two t-SNE components of the initial latents of the Wm & UnWm dataset. Data points for watermarked images are tightly clustered together around the center. A Gaussian kernel could easily separate the two classes.

images. In the Tree-Ring scheme, to detect the watermark, the image needs to be re-encoded back to latent space before forward diffusion can be used to recover the key. Although this process incurs a small amount of information loss, for key detection to work, this step must be sufficiently accurate to approximate the initial latent vector, suggesting that traces of the key must prevail to some extent.

This observation is the basis of our attacks: the difference in the distributions of the watermarked and non-watermarked latents naturally suggests attack strategies that aim to bring the watermarked latents back to their original distribution.

Separability of the latents. To further demonstrate the difference between the watermarked and non-watermarked latent distributions, we perform dimensionality reduction and plot two t-SNE (t-distributed stochastic neighbor embedding) components of the initial latents of a set of watermarked images and their non-watermarked counterparts (from the dataset “Wm & UnWm” described in Section 4.2). As Figure 4 confirms, the two classes are perfectly separable, suggesting that the initial watermarked latent variables form a hypersphere in the latent space. Moreover, Figure 4 shows this hypersphere exhibits compactness and well-defined boundaries.

This hints on the type of image transformation \mathcal{F} that would effectively bypass the detector of a Tree-Ring scheme. Such transformations require modifying the latent representation such that the forward diffusion process $\mathcal{D}_\theta^\dagger(x_0)$ does not map the image back to the original watermarked latent hypersphere. To enable such an attack, the adversary should focus on mapping the image back into latent space rather than defining transformations in the image space. Given the non-Gaussian nature of the intermediate latents, it may suffice to recover an intermediate latent. Such mapping can be achieved by using autoencoders, e.g., Stable Diffusion’s VAE [44], SDXL’s VAE [42] and other models with similar architectures [12].

3.3 Suitability of a Surrogate Detector Attack

Assuming we have the means to recover the latent space representation of an image, we ask the question of which attack strategy is most effective in pushing a point out of the watermarked hypersphere.

Regeneration attacks produce x_0^* by adding Gaussian noise to $\mathcal{E}(x_0)$, where \mathcal{E} encodes an image back to latent space. This noise completely disrupts the latent representation $\mathcal{E}(x_0)$ and changes the output image distribution. However, because the noise indiscriminately shifts the entire latent representation, the quality of the image suffers as a result.

Adversarial attacks aim to shift the latent via solving an optimization problem. For example, embedding attacks directly solve $\max_{x_0^*} \|\mathcal{E}(x_0^*) - \mathcal{E}(x_0)\|$ such that $\|x_0^* - x_0\| < \delta$, where $x_0^* := \mathcal{F}(x_0)$ is the transformed image. Although more refined than a regeneration attack, this strategy is not optimal. Due to the non-Gaussian nature of the watermark’s latent distribution, the overall goal of the transformation \mathcal{F} should not be to merely maximize the perturbation’s effect on the latent, as formulated in the optimization objective. Instead, the objective should also ensure that the latents of the transformed images are distributed like the original latents, i.e., $\mathcal{D}_\theta^\dagger(x_0^*) \sim \mathcal{N}(0; I)$, where $\mathcal{D}_\theta^\dagger(x_0^*)$ is the initial noise vector resulting from forward diffusion.

Surrogate detector attacks are better suited to achieve that. Training the surrogate detector is akin to modeling the watermarked hypersphere. In essence, such a surrogate detector is an approximation of the watermark detector independent of the implementation details of the watermark detection scheme. Thus, the loss of the surrogate detector provides a mechanism to mount an adversarial attack with precisely the desired objective: after training the surrogate detector, the adversary solves $\min_{x_0^*} L(\mathcal{E}(x_0^*), y_{\text{target}})$ s.t. $\|x_0^* - x_0\| < \delta$, where L is the loss of the surrogate detector.

This is a different objective from the surrogate detector attack evaluated by An et al. [4]. Their surrogate detector was trained on watermarked images, but, because it is unrealistic to assume the adversary can generate the non-watermarked versions of those images, they used publicly available images. Public datasets follow a different distribution than the model’s output distribution and, as a result, An et al.’s surrogate detector picked up on features that might have helped to distinguish between the two image distributions, but that were independent of the watermark, leading to an inadequate surrogate for finding adversarial examples.

Instead, we propose training the surrogate detector with the latent representation of the images, increasing the likelihood that it discriminates based on features dependent on the watermark. Our surrogate detector is therefore more likely to capture the watermarked region and generalize better than An et al.’s. This encoding step is crucial because the success of a surrogate detector attack hinges on how well the surrogate detector approximates the Tree-Ring detector.

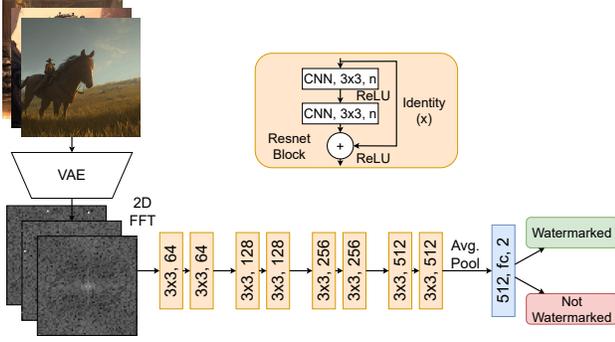


Figure 5: Surrogate detector training pipeline. The input data is a mixed batch of watermarked images generated via Stable Diffusion and either non-watermarked images generated via the same diffusion model or public images from ImageNet.

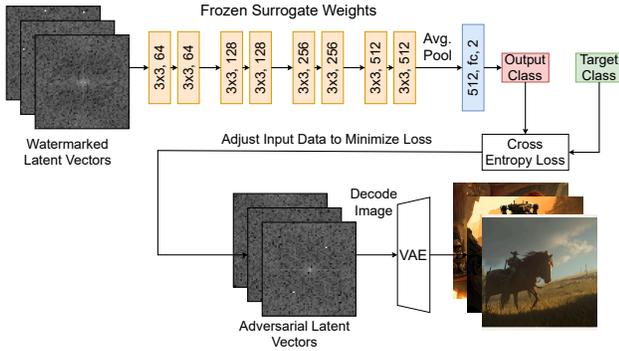


Figure 6: PGD attack pipeline. The attack uses the surrogate detector trained as shown in Figure 5 to alter the latent representation of the attacked image via Projected Gradient Descent (PGD). The target class is the “non-watermarked” class.

3.4 Overview of Our VAE-Based Attacks

Our attack implementation is divided into three main components: the *surrogate detector model*, the *autoencoder model*, and the *Project Gradient Descent (PGD) attack*. In this section, we describe how the adversary trains these models and combines them to deploy the attack.

Attack preparation. The first step of the attack is to collect a dataset of watermarked vs. non-watermarked images. Since the adversary has query access to the victim’s model, they can simply query it to generate the watermarked images. To obtain non-watermarked images, the adversary will draw them from public image datasets, ideally as semantically similar as possible to the generated images, as defined by the prompt.

To encourage the surrogate detector to learn to discriminate watermarked vs. non-watermarked images, the adversary uses a pretrained autoencoder that maps the images to latent space. In our evaluation, we start by assuming that this autoencoder

Algorithm 1 PGD Attack

Require: Surrogate Model θ_C , budget δ , image x , VAE \mathcal{E} , N steps, target class y

- 1: $\theta_{x_0} \leftarrow \mathcal{E}(x_0)$ ▷ Encode to latent space
- 2: **for** n to N **do**
- 3: $\hat{y} \leftarrow \theta_C(FFT(\theta_{x_0}))$ ▷ Fourier Transform (FFT)
- 4: $g_{\theta_{x_0}} \leftarrow \nabla_{\theta_{x_0}} BCELoss(\hat{y}, y)$
- 5: $g_{\theta_{x_0}} \leftarrow \mathcal{P}_\delta(g_{\theta_{x_0}})$ ▷ Clip gradients to δ budget
- 6: $\theta_{x_0} \leftarrow \theta_{x_0} - Adam(\theta_{x_0}, g_{\theta_{x_0}})$
- 7: **end for**
- 8: $x_0^* \leftarrow \mathcal{E}^{-1}(\theta_{x_0})$ ▷ Decode to image space
- 9: **return** x_0^*

is identical to the one of the victim and, then, evaluate the setting in which these two autoencoders are different.

Surrogate detector model. Our surrogate detector model is based on Resnet18 [25], a common neural network architecture in image classification tasks. Since our input data is in the frequency domain, we adapted its design to handle the frequency’s magnitude and phase as a complex number. A minor implementation detail is that, due to this representation, the output is also a complex number which is not fully supported by Pytorch cross-entropy loss. As such, we only use the real components of the output likelihood estimates when we train our surrogates on latent representations.

The surrogate is thus a binary classifier whose output is whether or not an image is watermarked based on its input data representation (see Figure 5). We use a batch size of 32, Adam optimizer with a learning rate of 10^{-3} , and optimize for cross-entropy loss. For model selection, we use a 7:3 evaluation split ratio, and save the model with the highest validation accuracy every 100 epochs for later use in the PGD attack.

PGD attack. The ultimate goal of the attack is to construct an image close to the original target image that our surrogate detector misclassifies as “non-watermarked.” Our surrogate detector predicts the class of the image and calculates the cross-entropy loss between the output class and the target class, i.e., the non-watermarked class (see Figure 6). Backpropagation provides the gradient step required to shift our data towards being misclassified. As shown in Algorithm 1, we repeat this process for N steps and introduce a step size parameter α to control the smoothness of the perturbations. We take $N = 200$ and $\alpha = 0.05$, a setup similar to the one by An et al. [4].

An et al. use a perturbation budget to limit the amount of pixel changes. Their budget is defined in the image domain, however, since our attack acts in the latent space, we must rescale the budget. To find the appropriate scale, we take the maximum number of watermarked latents in our dataset p and set our perturbation budget to $\frac{1}{p}$. After completing the N steps, we decode the latent into an image using the autoencoder.

To sum up, once we have trained the surrogate, the deployment of the attack on a watermarked image requires encoding

the image using the autoencoder, obtaining its Fourier transform, and performing the PGD attack with the loss of the surrogate on the image. The attacker can repeat the attack with different values of δ until the adversarial image on the surrogate transfers to the Tree-Ring detector.

4 Evaluation Methodology

4.1 Models

We take Stable Diffusion v2.1 [44] as the victim’s diffusion model. This version of Stable Diffusion uses AutoencoderKL [32] by default as its VAE to generate the image from the intermediate latent space. We use this VAE to encode and decode images for our surrogate detector attack. However, in Section 5.4, we also challenge the assumption of having the exact same VAE by evaluating the attack with different VAEs.

4.2 Datasets

To assess the effectiveness of our attacks, we used the ImageNet1K dataset [47]. Specifically, following the methodology established in previous work [4], we used 7,000 images from the ImageNet1K validation dataset. Each image has a set of annotated bounding boxes around objects each corresponding to one of the classes in ImageNet1K. We sample 7 images from each class and preprocess the image by cropping to the first annotated bounding box. Afterwards, we resize each image to 512x512 to match the output dimensions of Stable Diffusion. To match the classes from ImageNet1K, we queried Stable Diffusion v2.1 to generate an image for each class using the prompt: “A photo of a <ImageNet Class>.” Each generated image was watermarked with the Tree-Ring scheme, and 7 images were generated for each class, each with a different random seed to introduce diversity in the outputs.

To train the surrogate detector, we collect two datasets:

Wm & UnWm represents the setting where the adversary can use the diffusion model to generate both the watermarked and non-watermarked images. It includes a total of 3,000 images: 1,500 images are generated using a victim diffusion model with Tree-Ring watermarks (Wm); the remaining 1,500 are generated using the same diffusion model, random seed, and image prompt, but skipping the watermarking step (UnWm).

Wm & Pub captures the more realistic setting where the adversary uses images in a public dataset as the non-watermarked training examples. This dataset has a total of 14,000 images, 7,000 images are generated using a victim diffusion model with Tree-Ring watermarks, the other 7,000 are preprocessed ImageNet images.

The Wm & UnWm dataset is smaller since training an accurate surrogate detector with the non-watermarked version of the watermarked images requires significantly less data.

4.3 Performance Metrics

In line with existing evaluations of watermarking techniques for generative models [4, 56, 63], we have selected metrics to assess various aspects of the attack. These metrics measure the fulfillment of the two adversarial goals outlined in Section 3.1: (i) the first set of metrics quantify the success in evading detection measured by the decrease in *predictive performance* of the watermark detector after deploying the attack; (ii) the second set of metrics measure the impact of the attack on the target image by quantifying the degradation of the attacked image compared to the original.

As pointed out by An et al., these objectives are context-dependent and are subject to the adversary’s use case. For this reason, we use multiple metrics for each objective, providing a comprehensive evaluation and a discussion of their practical implications. The two sets of metrics are as follows.

4.3.1 Predictive Performance Metrics

We assume that the reader is familiar with basic binary classification predictive performance metrics, such as Accuracy, Precision, Recall (TPR), and FPR, and we will only describe metrics that are less known or that are particularly common in the evaluation of watermarking techniques.

TPR@1%FPR is the detector’s TPR at a threshold chosen to ensure a 1% FPR. A 1% FPR threshold is commonly used in prior evaluations of watermarking schemes and is considered a stringent criterion [4, 56].

ROC-AUC measures the *Area Under the Curve* (AUC) of the *Receiver Operating Characteristic* (ROC) curve, a common performance metric for binary classifiers. In this case, it quantifies the detector’s ability to detect watermarked images correctly as an aggregate measure of the detector’s TPR–FPR trade-off for all possible decision thresholds.

PR-AUC is the Area Under the Precision–Recall curve. Like the ROC-AUC, this is a composite metric that measures the classifier’s ability to ensure both high Precision and Recall (TPR), aggregating the classifier’s trade-offs at all possible decision thresholds.

We measure these metrics using the setup used in Tree-Ring’s original paper [56], setting a 1:1 ratio between the two classes. For example, for the Wm & UnWm dataset, we have 1,500 attacked watermarked images vs. 1,500 non-watermarked images. In Section 5.6, we evaluate performance while varying this ratio.

4.3.2 Image Degradation Metrics

We evaluate the visual impact of the attacks using several metrics, each reflecting different aspects of image quality.

CLIP Score is the cosine similarity between an image embedding and a textual embedding given a fixed text-to-image

model. We use it to measure how well the attacked image represents the original prompt. We use the open-source model ViT-g-14 [29, 43] pretrained on LAION-5B [48] to generate image and textual embeddings.

FID measures the distance between probability distributions of images by mapping them to a latent representation from a predefined network InceptionV3 [53], capturing similarities in the distributions rather than individual images.

LPIPS [61] computes the similarity of the activations generated by an image classification model at each layer, in our case Alexnet [33]. Unlike FID, LPIPS is defined between two images rather than between two samples of images.

In our evaluation, we are interested in the degradation effect of our attacks relative to a ground truth image (or set of images for FID). As such, for each of the image degradation metrics above, we calculate:

$$|Q(x_0, x_{\text{ref}}) - Q(x_0^*, x_{\text{ref}})|, \quad (8)$$

where Q is the quality metric (CLIP Score, FID, or LPIPS), x_0 is a generated watermarked image using Tree-Ring, x_0^* is the transformed image after applying an attack, and x_{ref} is the real image from ImageNet1k which was used to derive the prompt to generate x_0 and serves as the reference image.

4.4 Evaluation Baselines

We take various baselines as a reference point to assess the success of our attacks. These baselines include idealized settings and state-of-the-art attacks ([4, 37]).

No-Attack: Tree-Ring’s performance with no attacks applied.

Raw Pixel Values: Directly train a surrogate detector on the images without converting to latent space, i.e. the setting considered in An et al.’s study [4].

True Latent Vectors: Train a surrogate detector on the final ($t = 0$) latent vectors before decoding the latents into images via AutoencoderKL. This is an idealized baseline as it assumes that the adversary has full access to the model instead of using the VAE to approximate these latents.

Adversarial Noising: We reproduce the adversarial noising attack by Lukas et al. [37]. This is also an idealized baseline as it assumes access to a model with the same architecture as the victim’s, making it less practical than our attacks. We discard adversarial compression also by Lukas et al. because they show that it performs worse than adversarial noising [37].

5 Evaluation

In this section, we present the results of our evaluation. We evaluate various aspects of the attacks, including their performance compared to the baselines, and the detector’s precision under a range of deployment conditions. In addition, we conduct ablation studies to evaluate the impact of differences

Table 1: Surrogate Detector Training. Watermarked images (Wm) are generated using the victim’s diffusion model with the Tree-Ring scheme. Non-watermarked images (UnWm) are generated without Tree-Ring. Public images (Pub) are taken from the ImageNet1k ILSVRC2012 dataset.

Attack Name	Training Dataset	Training Accuracy	Validation Accuracy
Raw Pixel Values	Wm & UnWm	99.95%	96.44%
	Wm & Pub	99.86%	99.52%
True Latent Vectors	Wm & UnWm	99.00%	96.00%
	Wm & Pub	99.93%	100%
VAE-Recovered Latent Vectors	Wm & UnWm	100%	94.78%
	Wm & Pub	99.93%	98.48%

in the attack setup and the deployment setting, including a different VAE and a different diffusion model.

5.1 No-Attack Baseline

For the first experiment, we test the Tree-Ring watermark detector on our dataset when no attack is in place. The results of this experiment are presented in the first row of Table 2.

The Tree-Ring watermark detector shows near perfect detection accuracy without affecting image quality. The high detection accuracy aligns with the original reported results, and no image degradation occurs since the images remain unchanged. This establishes a baseline for evaluating how attacks impact detector performance in later experiments.

5.2 Surrogate Detector Performance

Before applying the attacks, we have evaluated the classification accuracy of the surrogate detectors we consider. We measure the training and validation accuracies to assess the potential overfitting of the surrogate detector. If the surrogate detector overfits the training data, the PGD attack is less likely to succeed in finding images that evade it.

As seen in Table 1, training and validation accuracies are high for all surrogate detectors, with none of them significantly overfitting. All surrogate detectors perform better when trained on Wm & Pub than Wm & UnWm, which can be explained by the greater difference between the watermarked and public image distributions in the latter dataset. The detector trained on true latent vectors (second row) achieves the best validation accuracy, which is rather unsurprising given that there is no loss of information in approximating the latent.

The last row in Table 1 shows the results for our surrogate detector, namely the same task as before but training on latent vectors recovered via the AutoencoderKL. The drop in Validation Accuracy indicates that this task is more challenging. We hypothesize that this is due to the information loss incurred by re-encoding the image back into latent space with the VAE.

5.3 VAE-Based Latent Vector Attacks

After training all our surrogate models, we attempt a PGD attack on the watermarked images. Figure 9 in Appendix A illustrates the visual impact of all attacks on an example image by providing the resulting attacked image (first column), and the difference between the original and resulting images in image space (second column) and Fourier space (third column). As shown in that figure, none of the attacks exhibits obvious unusual patterns in the image. However, the attacks have different impacts on the region of the Fourier space where the watermark is embedded, with our VAE-recovered attacks having a precise and concentrated effect on the key ring area.

In Table 2, we report the performance of the watermark detector under each of the attacks, as well as the impact of such attacks on image quality. We observe that our attack based on the recovery of the latents with AutoencoderKL is the most successful in disrupting watermark detection (highlighted row). Furthermore, the quality of the image is retained compared to the baseline CLIP score: it only rises by 0.008. Similarly, LPIPS and FID Score only have small increases, especially for the surrogate trained for distinguishing watermarked latents vs. public image latents. The changes in pixel space in the fifth and sixth rows of Figure 9 are imperceptible.

Contrary to the expectation that access to the true latents (fourth row of Table 2) would lead to better attack performance, we see that this attack performs worse than our surrogate attack trained on VAE-recovered latent vectors. We believe that differences in distribution and scale between the true latent vectors and the VAE-recovered ones account for the attack’s lower performance in this setting, as it learns an image perturbation that is effective in the true latent space but not necessarily in the VAE-recovered space.

As a result of the attacks, some of the ROC-AUC values are below the random guess baseline of 0.5. Note that in an adversarial setting, the model owner cannot simply invert the detector to achieve $\text{ROC-AUC} > 0.5$ because they have no reliable way of predicting the occurrence and the extent of an attack. Since the detector’s decision is based on a threshold on the error between the recovered watermark and the initial watermark, flipping the decisions implies that the greater the dissimilarity between the recovered and initial watermark, the more likely it is to be watermarked. Therefore, if the model owner preemptively deploys the inverted detector and no attack takes place, inverting it causes it to miss most watermarked images, rendering the detector ineffective.

We found a discrepancy between An et al.’s and our results when executing the PGD attack with a surrogate trained on the Raw Pixel Values of the Wm & Pub dataset: they reported an average $\text{TPR}@1\%FPR$ of 0.99, but our results show a value of 0.48. The main difference between our setups is that we trained the Resnet18 from the ground up while they fine-tuned a pretrained version. Using the pretrained Resnet18 may have caused their surrogates to incorporate broader features

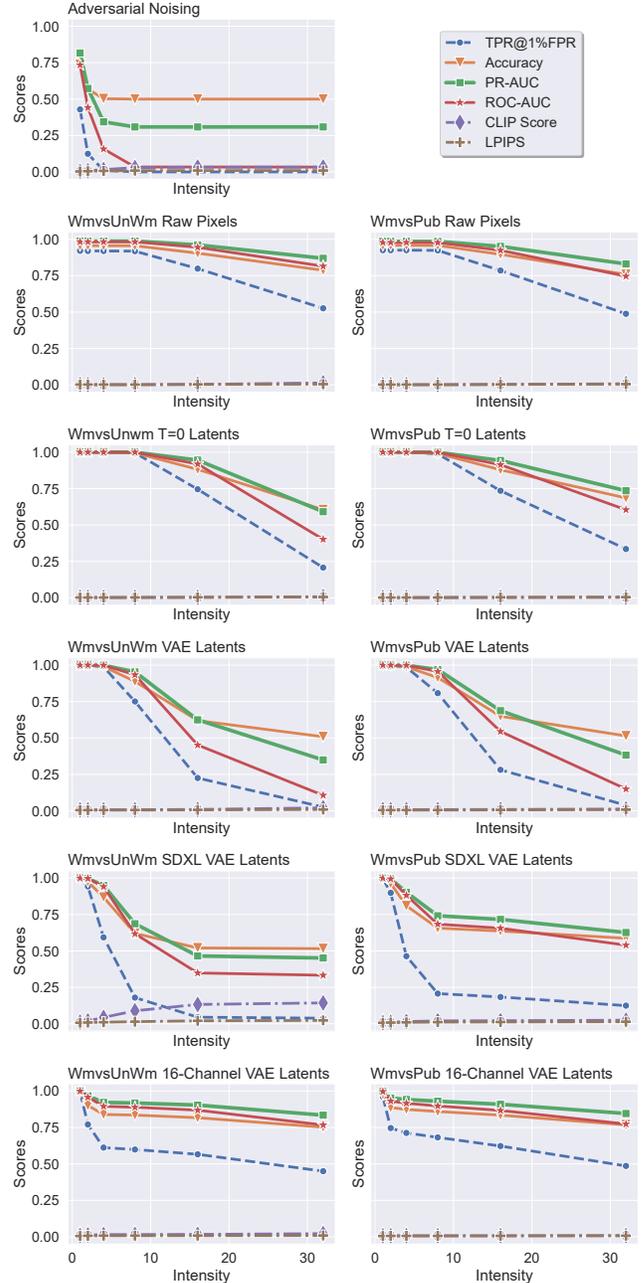


Figure 7: Results of our attack while varying δ to adjust the strength of the attack. Next to each score in the legend, we denote what direction of magnitude indicates a successful attack. We only include metrics that range within the unit interval, as such we exclude FID.

when fine-tuning for classification, which indicates that the adversary can benefit from training Resnet18 from scratch.

We reproduced the original adversarial noising paper’s results, achieving a $\text{TPR}@1\%FPR$ of 0.141 for $\delta = 2$, slightly higher than the reported 0.052, likely due to dataset differences. Although adversarial noising minimally degrades image qual-

Table 2: Average performance of the attacks across three random seeds. Standard deviations were consistently below 0.01 for all metrics and are thus omitted. For all the scores, the lower the score, the more successful the attack is.

Attack Name	Training Data	PR-AUC	ROC-AUC	Accuracy	TPR@1%FPR	CLIP Score	FID	LPIPS
No Attack	N/A	0.994	0.993	0.965	0.968	0.000	0.000	0.000
Adversarial Noising	N/A	0.584	0.459	0.575	0.141	0.005	0.602	0.018
Raw Pixel Values	Wm & UnWm	0.870	0.816	0.785	0.529	0.005	1.163	0.014
	Wm & Pub	0.828	0.743	0.756	0.485	0.006	0.303	0.008
True Latent Vectors	Wm & UnWm	0.599	0.413	0.610	0.212	0.006	0.689	0.008
	Wm & Pub	0.742	0.615	0.692	0.343	0.005	0.519	0.004
SDXL-VAE Latent Vectors	Wm & Unwm	0.452	0.333	0.516	0.039	0.023	33.71	0.143
	Wm & Pub	0.626	0.540	0.587	0.125	0.014	2.453	0.025
16-Channel VAE Latent Vectors	Wm & Unwm	0.834	0.767	0.752	0.450	0.008	3.754	0.021
	Wm & Pub	0.846	0.774	0.768	0.486	0.007	0.088	0.008
VAE-Recovered Latent Vectors	Wm & UnWm	0.350	0.108	0.508	0.023	0.009	2.684	0.021
	Wm & Pub	0.385	0.153	0.515	0.039	0.008	0.694	0.012

ity, our attacks are more effective and do not require access to the watermarking scheme.

We extend our attack evaluation by varying the perturbation budget, δ , for each attack. These results are plotted in Figure 7. The adversarial noising attack is extremely effective when using a δ of 8 and above. However, the image quality starts to deteriorate at that point. Our attacks work well with a δ of 32. Furthermore, image degradation is not affected for any of the surrogate models we have trained, demonstrating the effectiveness of surrogate detector attacks.

Finally, comparing these results with those reported in the original Tree-Ring paper, all the attacks included in our evaluation are significantly more effective than any of their tested approaches. They reported on average for each of their adversarial examples, an ROC-AUC of 0.975 and TPR@1%FPR of 0.694. Compared with our reported results, we have an average of 0.519 ROC-AUC and 0.261 TPR@1%FPR, which represents a 1.9x decrease in general detection accuracy.

5.4 Ablation Study for Different VAEs

While it is extremely common for text-to-image models to directly use a publicly available checkpoint, there still exist a few fine-tuned autoencoders that continue training from said publicly available checkpoint on Civitai. To account for this, we also evaluated our attack’s effectiveness when using different VAEs to train our surrogate detectors, addressing scenarios where model owners might employ private VAEs. Specifically, we considered two cases: (1) when the victim model uses a fine-tuned version of a publicly available VAE, and (2) when the adversary uses a VAE with a different architecture. For the fine-tuned VAE scenario, we utilized SDXL’s VAE, which is derived from Stable Diffusion’s VAE. To test a different architecture, we employed a 16-Channel VAE [12], which has fewer parameters than Stable Diffusion’s

VAE and produces latents of dimension (64,64,16), compared to Stable Diffusion’s (64,64,4) latent dimensions.

The results are presented in rows 5–6 of Table 2 and Figure 7. Using SDXL’s VAE, the attack remains robust, albeit less effective than when using the victim model’s VAE. However, the impact on image quality becomes more pronounced. Training on the Wm & UnWm dataset significantly degrades image quality, with the FID score increasing by 33.71 and similar increases in LPIPS and CLIP Score. We also show the attacked image in Figure 9 of Appendix A. Upon close inspection, the image has noticeable image artifacts. The 16-Channel VAE’s performance is comparable to training the surrogate detector on raw images. Interestingly, when trained on the Wm vs. Pub dataset, it causes less image degradation. We hypothesize that SDXL’s superior performance over the 16-Channel VAE is due to it being a fine-tuned version of the victim model’s VAE, resulting in re-estimated latents with distributions similar to those of the victim model’s VAE.

These findings suggest that using a customized VAE significantly enhances watermark robustness. We observed that increasing the perturbation budget δ , leads SDXL VAE to substantial image degradation, particularly for the surrogate trained on the Wm & UnWm dataset. We anticipate that further increasing δ would eventually impact the quality of attacked images derived from the surrogate trained on the Wm & Pub dataset as well. This implies a trade-off between attack effectiveness and maintaining image quality, with the choice of VAE playing a crucial role in this balance.

5.5 Ablation Study Without a VAE

To determine the importance of having access to a VAE for the success of our attack, we evaluate it on OpenAI’s Guided Diffusion [20], an open source implementation of a diffusion model that does not use a VAE. While Tree-Ring

can be directly applied to Guided Diffusion, this is not a text-to-image model but a class-conditioned diffusion model trained on ImageNet classes and thus does not expect a prompt for generation. Aside from this, the evaluation follows the same methodology: we use the model to generate images and construct the Wm & Unwm and Wm & Pub datasets to train the surrogate models and apply PGD on them.

Table 3: Attack performance on Tree-Ring applied to Guided Diffusion. Unlike Table 2, we omit CLIP Score and LPIPS as Guided Diffusion is a class-conditioned model and ground truth images are not available for assessing generation quality.

Training Data	PR-AUC	ROC-AUC	Accuracy	TPR@1%FPR	FID
Wm & UnWm	0.8503	0.7870	0.7713	0.4880	56.44
Wm & Pub	0.8506	0.7855	0.7712	0.4803	56.44

Table 3 presents the results of this evaluation. While the attack still decreases the effectiveness of Tree-Ring, it is not as successful as when the VAE is available. This result aligns with the findings by An et al. [4], who show that surrogate-based attacks on images do not perform as well as other adversarial attacks, indicating that the attack achieves lower performance if the attacker does not use a VAE to encode the images into the latent space. These results highlight the critical role of the VAE in enabling our attack and demonstrate the risk posed by making this diffusion model component publicly available.

5.6 Impact of the Base Rate

Precision has been overlooked in the evaluation of watermarking schemes. This is a limitation, as it misses an aspect of detection performance not captured by performance metrics based on TPRs and FPRs. Even with high TPR and low FPR, precision can be low if the base rate of the positive class (watermarked) is low, leading to low confidence in the detection of watermarked images. In real-world deployments, the base rate is likely to be low, as any watermarking scheme that is not broadly adopted will be overwhelmingly exposed to non-watermarked content. Excluding precision in the evaluation of a detector where a class imbalance is likely is a case of the *base rate fallacy*, an evaluation bias that has been extensively studied in domains with similar detection tasks [5, 7, 30].

Since we do not know the base rate of real-world deployments, we plot the precision curves of all attacks and baselines over a range of possible base rates for four ROC operating points obtained by imposing a minimum TPR of 0.5, 0.75, 0.95, and 0.99 (see Figure 8). The linearity of the high-TPR curves is expected, as in those cases $TPR \approx FPR$. However, for base rates below 0.1, which are plausible in a deployment scenario, precision decreases significantly, even for high TPR.

In the fourth row of Figure 8 we see the precision curves resulting from applying our attack. For both datasets, the false positives induced by our attacks make precision drop linearly

even for $TPR=0.5$. Figure 10 of Appendix B complements these results with the Precision–Recall curves where we can see that our attacks are particularly effective in decreasing the precision of the Tree-Ring detector.

Given these results, we would not recommend this scheme for deployment. Even in the non-adversarial setting, for a base rate of 0.1, the model owner would require picking an operating point close to $TPR = 0.5$ to achieve high precision (see also the first column of Figure 10 in Appendix B), missing half of the total watermarked images.

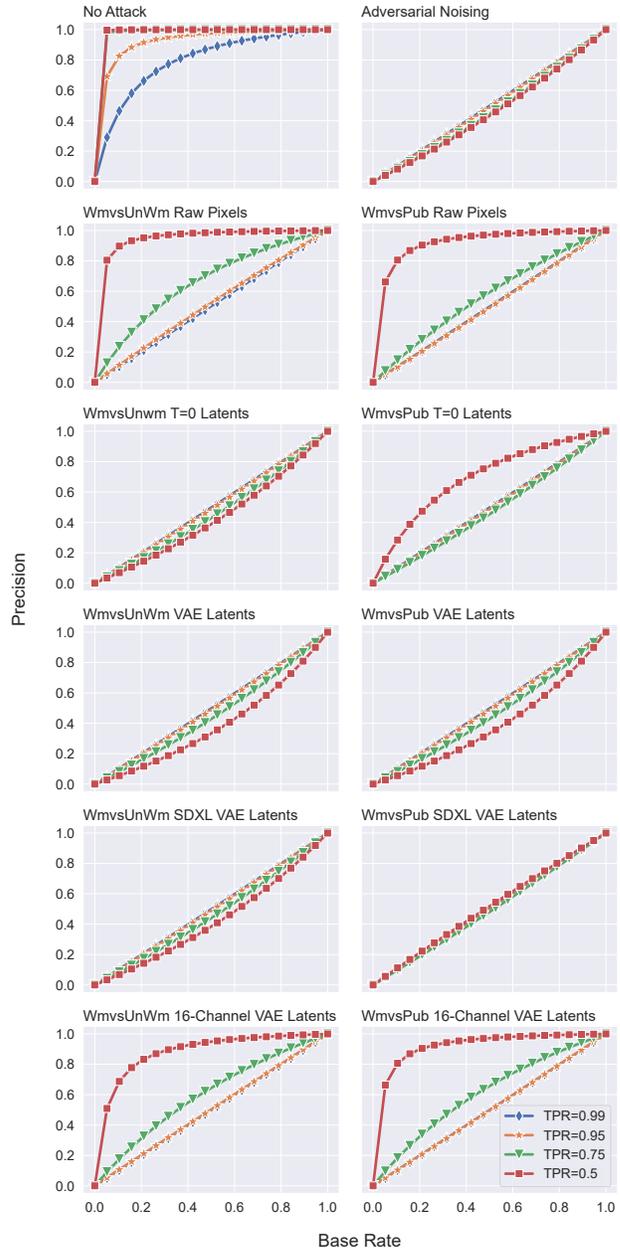


Figure 8: Detector precision over the base rate for each attack. Each line represents a different operating point in the ROC.

6 Discussion

We now discuss the implications of our findings, not only for Tree-Ring, but also for other image watermarking schemes.

Attack’s reliance on VAE access. Our results indicate that the attack heavily relies on the target VAE. We conducted two ablation studies to test this assumption: one where the attacker uses a different VAE (Section 5.4), and another where the diffusion model does not use a VAE at all (Section 5.5). In both settings, the attack is significantly less effective compared to the ideal case. Despite this limitation, our contributions remain practically relevant, as these ideal conditions are likely to be met in real-world scenarios. As argued in Section 3.1, it is common practice to reuse existing VAEs to train diffusion models through latent diffusion.

The security arms race. As watermarking schemes continue to gain traction in both academia and industry [19], ensuring their robustness against removal becomes critical. Currently, most security claims for these schemes are based on experimental evaluations against ad-hoc attacks, meaning they are only considered secure until a successful attack is discovered. The security of image watermarking schemes should instead be grounded in computational hardness assumptions or existing cryptographic primitives. There is progress in this direction for watermarking schemes for large language models [13, 22, 35] and, more recently, for image models [24].

Precision as a key metric. The evaluation of watermarking schemes often focuses on metrics that do not include precision. However, relying solely on these metrics may overestimate the performance of the detector once it is deployed. Thus, security evaluations of image watermarking should use precision-based metrics like PR-AUC, reflecting performance in low base rate scenarios, such as those expected in the future deployment of these watermarking schemes.

Our evaluation of Tree-Ring’s precision reveals that, even in the absence of attacks, the detector’s precision is low in realistic deployment scenarios. In a world where most images do not contain a watermark, the detector’s confidence in detecting watermarks decreases significantly without an attack and falls drastically under our attacks, indicating that the detector is not reliable when it detects watermarks in general settings.

These findings suggest that if watermarking schemes like Tree-Ring were deployed to monitor content on publicly accessible platforms, as it is currently proposed, their performance would be inadequate at best and completely unusable if attacks were widespread. This raises concerns about the practical viability of these watermarking techniques.

Potential countermeasures. Our attacks exploit the shift in the distribution of the latent space caused by the embedding of a fixed watermark. Therefore, a natural mitigation strategy is to ensure that the latent representations of watermarked and non-watermarked images are indistinguishable, making it harder to train effective surrogate detectors.

Additionally, the success of our attacks degrades when the attacker and the victim’s autoencoder differ. Thus, a short-term countermeasure is to train diffusion models with a custom private autoencoder, ensuring some degree of misalignment between the diffusion latents and those estimated by the attack, although this would mean that practitioners would not be able to benefit from the cost savings offered by latent diffusion [4].

Defenders may also leverage asymmetries in the deployment setting. For example, Tree-Ring assumes the attacker has black-box access to the diffusion model, as knowledge of its weights would enable direct recovery of the watermarks. This restricted access forces the adversary to rely on techniques that approximate the inversion process, which increases the cost of attack compared to that of defense. To capitalize on this asymmetry, a model owner could embed multiple watermarks per image—potentially using different embedding schemes and rotating the keys—thereby forcing an attacker to craft perturbations that evade several detectors simultaneously. In addition, access-control mechanisms could restrict the number of queries associated with a specific watermarking key, thus limiting the adversary’s ability to collect a sufficient number of watermarked samples for training a surrogate detector.

Countermeasures designed against adversarial examples, such as adversarially training the watermark detector, would improve the robustness of the detector against surrogate PGD attacks. However, prior work has consistently shown that such defenses often fail against adaptive attackers [6] or significantly degrade model accuracy [15]. Currently, no defense provides certified robustness to adversarial examples without incurring substantial accuracy trade-offs.

External validity of the results. Note that we do not require evaluating the attacks on multiple datasets to demonstrate the generality of the attacks, as the attacker can choose any publicly available dataset to train the surrogate detector.

Although we have not tested the effectiveness of our attacks against other watermarking schemes, our analysis indicates that any watermarking scheme that embeds a distinctive watermark in the latent space is likely to be susceptible to a VAE-based attack, even if the difference is not visible in the image space.

7 Conclusion

Our findings have uncovered a novel class of surrogate detector attacks based on public pretrained autoencoders. The results demonstrate that our method not only surpasses existing surrogate detector attacks but also outperforms other attacks that make strong assumptions about the threat model. Additionally, our evaluation shows that Tree-Ring’s performance is lower than initially thought under typical deployment scenarios, raising concerns about the practicality of these schemes.

Acknowledgments

We thank our shepherd and the anonymous reviewers for their valuable feedback. This research was supported by the LFCS internship program at the University of Edinburgh, which funded Junhua Lin, and by the University of Edinburgh’s Generative AI Laboratory (GAIL), which provided access to their high-performance cluster. Marc Juarez is a GAIL fellow and the recipient of a Google Research Scholar Award in Security. We are grateful to our colleagues who provided feedback on earlier drafts of the manuscript, including Kai Yao and Rik Sarkar. We extend special thanks to Nils Lukas for his assistance in reproducing adversarial noising, and to Kai Yao and Bence Szilágyi for their help in setting up the GAIL cluster and providing helper scripts.

Ethical Considerations

We have reviewed the USENIX Security Ethics Guidelines and carefully considered the ethical aspects of this research.

Regarding the impact of our findings, we concluded that the benefits of exposing these attacks clearly outweigh the risks. ML watermarking schemes are still in the prototype stage and have not been deployed. Therefore, the likelihood of malicious actors exploiting these attacks is minimal, while exposing their vulnerabilities gives a better understanding of the limitations of ML watermarking. It is crucial to raise awareness of the vulnerabilities of these schemes so that contingency measures are in place before they are deployed more broadly. Furthermore, since attacks were evaluated in a controlled and isolated environment, conducting this research did not cause any harm.

Open Science

In accordance with best practices for open science, we make our [source code](#)¹ openly available. The repository contains our attack implementations, data processing helper methods, and the datasets used. All diffusion models used and Tree-Ring were made publicly available by other researchers; we have provided instructions on how to download and integrate them into our evaluation framework. We also provide instructions on how to reproduce the results of our experiments.

References

- [1] AIFNET. Deep Dream Generator. <https://deepdreamgenerator.com>. Accessed: May 27, 2025.
- [2] AISEO. AISEO Art - AI art generation. <https://art.aiseo.ai>. Accessed: May 27, 2025.
- [3] AL-HAJ, A. Combined DWT-DCT digital image watermarking. *Journal of computer science* 3, 9 (2007), 740–746.
- [4] AN, B., DING, M., RABBANI, T., AGRAWAL, A., XU, Y., DENG, C., ZHU, S., MOHAMED, A., WEN, Y., GOLDSTEIN, T., AND HUANG, F. WAVES: Benchmarking the Robustness of Image Watermarks. In *International Conference on Machine Learning (ICML)* (2024).
- [5] ARP, D., QUIRING, E., PENDLEBURY, F., WARNECKE, A., PIERAZZI, F., WRESSNEGGER, C., CAVALLARO, L., AND RIECK, K. Dos and don’ts of machine learning in computer security. In *31st USENIX Security Symposium (USENIX Security 22)* (2022), pp. 3971–3988.
- [6] ATHALYE, A., CARLINI, N., AND WAGNER, D. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proceedings of the 35th International Conference on Machine Learning* (10–15 Jul 2018), J. Dy and A. Krause, Eds., vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 274–283.
- [7] AXELSSON, S. The base-rate fallacy and its implications for the difficulty of intrusion detection. In *Proceedings of the 6th ACM Conference on Computer and Communications Security* (1999), pp. 1–7.
- [8] BIGMODEL. CogView 4 - Large Vision Language Model. <https://open.bigmodel.cn/pricing>. Accessed: May 27, 2025.
- [9] BIRD, C., UNGLESS, E., AND KASIRZADEH, A. Typology of risks of generative text-to-image models. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (2023), pp. 396–410.
- [10] BLACK FOREST LABS. FLUX 1.1 Pro. <https://huggingface.co/black-forest-labs/FLUX.1-dev>, 2024. Accessed: May 27, 2025.
- [11] BOENISCH, F. A systematic review on model watermarking for neural networks. *Frontiers in big Data* 4 (2021), 729663.
- [12] BURKETT, J. Ostris vae - kl-f8-d16, 7 2024.
- [13] CHRIST, M., GUNN, S., AND ZAMIR, O. Undetectable watermarks for language models. In *The Thirty Seventh Annual Conference on Learning Theory* (2024), PMLR, pp. 1125–1139.
- [14] CLIPDROP. Clipdrop - AI tools powered by Stable Diffusion. <https://clipdrop.co>. Accessed: May 27, 2025.
- [15] COHEN, J., ROSENFELD, E., AND KOLTER, Z. Certified adversarial robustness via randomized smoothing. In *international conference on machine learning* (2019), PMLR, pp. 1310–1320.
- [16] COX, I., MILLER, M., BLOOM, J., FRIDRICH, J., AND KALKER, T. *Digital watermarking and steganography*. Morgan kaufmann, 2007.
- [17] CRAIYON. Craiyon - AI model based on DALL-E. <https://www.craiyon.com>. Accessed: May 27, 2025.
- [18] DAVINCI AI. Davinci AI - Powered by SDXL. <https://davinci.ai>. Accessed: May 27, 2025.
- [19] DEEPMIND, G. Synthid: Identifying ai-generated images. <https://deepmind.google/technologies/synthid>, 2023. Accessed: 2024-09-08.
- [20] DHARIWAL, P., AND NICHOL, A. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems* (2021), M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34, Curran Associates, Inc., pp. 8780–8794.
- [21] DU, L., ZHOU, X., CHEN, M., ZHANG, C., SU, Z., CHENG, P., CHEN, J., AND ZHANG, Z. Sok: Dataset copyright auditing in machine learning systems. *arXiv preprint arXiv:2410.16618* (2024).
- [22] FAIROZE, J., GARG, S., JHA, S., MAHLOUJIFAR, S., MAHMOODY, M., AND WANG, M. Publicly-Detectable Watermarking for Language Models, 2023.
- [23] FERNANDEZ, P., COUAIRON, G., JÉGOU, H., DOUZE, M., AND FURON, T. The stable signature: Rooting watermarks in latent diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 22466–22477.
- [24] GUNN, S., ZHAO, X., AND SONG, D. An undetectable watermark for generative image models. *arXiv preprint arXiv:2410.07369* (2024).

¹<https://doi.org/10.5281/zenodo.15595719>

- [25] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (6 2016).
- [26] HO, J., JAIN, A., AND ABBEEL, P. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems* (2020), H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33, Curran Associates, Inc., pp. 6840–6851.
- [27] HU, E. J., SHEN, Y., WALLIS, P., ALLEN-ZHU, Z., LI, Y., WANG, S., WANG, L., AND CHEN, W. LoRA: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations* (2022).
- [28] HUGGING FACE, I. Hugging Face. <https://huggingface.co/>, 2016. Accessed: May 21, 2025.
- [29] ILHARCO, G., WORTSMAN, M., WIGHTMAN, R., GORDON, C., CARLINI, N., TAORI, R., DAVE, A., SHANKAR, V., NAMKOONG, H., MILLER, J., HAJISHIRZI, H., FARHADI, A., AND SCHMIDT, L. Openclip, July 2021. If you use this software, please cite it as below.
- [30] JUAREZ, M., AFROZ, S., ACAR, G., DIAZ, C., AND GREENSTADT, R. A critical evaluation of website fingerprinting attacks. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security* (2014), pp. 263–274.
- [31] KAMALI, N., NAKAMURA, K., CHATZIMPARMPAS, A., HULLMAN, J., AND GROH, M. How to distinguish ai-generated images from authentic photographs, 2024.
- [32] KINGMA, D. P., AND WELLING, M. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings* (2014), Y. Bengio and Y. LeCun, Eds.
- [33] KRIZHEVSKY, A., SUTSKEVER, I., AND HINTON, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (2012), F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Eds., vol. 25, Curran Associates, Inc.
- [34] LEONARDO AI. Leonardo.Ai - Generative AI with Stable Diffusion. <https://leonardo.ai>. Accessed: May 27, 2025.
- [35] LIU, A., PAN, L., HU, X., LI, S., WEN, L., KING, I., AND PHILIP, S. Y. An unforgeable publicly verifiable watermark for large language models. In *The Twelfth International Conference on Learning Representations* (2023).
- [36] LU, Z., HUANG, D., BAI, L., QU, J., WU, C., LIU, X., AND OUYANG, W. Seeing is not always believing: benchmarking human and model perception of ai-generated images. *Advances in Neural Information Processing Systems (NeurIPS)* 36 (2024).
- [37] LUKAS, N., DIAA, A., FENAUX, L., AND KERSCHBAUM, F. Leveraging optimization for adaptive attacks on image watermarks. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024* (2024).
- [38] LUKAS, N., AND KERSCHBAUM, F. {PTW}: Pivotal tuning watermarking for {Pre-Trained} image generators. In *32nd USENIX Security Symposium (USENIX Security 23)* (2023), pp. 2241–2258.
- [39] MAIER, J. CivitAI. <https://civitai.com/>, 2022. Accessed: Jan 21, 2025.
- [40] OPENAI. openai/DALL-E. <https://github.com/openai/DALL-E>, 2021. Accessed: May 27, 2025.
- [41] OPENAI. Sora: Text-to-Video Generation. <https://openai.com/sora>, 2024. Accessed: May 27, 2025.
- [42] PODELL, D., ENGLISH, Z., LACEY, K., BLATTMANN, A., DOCKHORN, T., MÜLLER, J., PENNA, J., AND ROMBACH, R. SDXL: improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024* (2024).
- [43] RADFORD, A., KIM, J. W., HALLACY, C., RAMESH, A., GOH, G., AGARWAL, S., SASTRY, G., ASKELL, A., MISHKIN, P., CLARK, J., KRUEGER, G., AND SUTSKEVER, I. Learning transferable visual models from natural language supervision. In *ICML* (2021).
- [44] ROMBACH, R., BLATTMANN, A., LORENZ, D., ESSER, P., AND OMMER, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2022), pp. 10684–10695.
- [45] RONNEBERGER, O., FISCHER, P., AND BROX, T. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (Cham, 2015), N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds., Springer International Publishing, pp. 234–241.
- [46] RUIZ, N., LI, Y., JAMPANI, V., PRITCH, Y., RUBINSTEIN, M., AND ABERMAN, K. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Los Alamitos, CA, USA, 6 2023), IEEE Computer Society, pp. 22500–22510.
- [47] RUSSAKOVSKY, O., DENG, J., SU, H., KRAUSE, J., SATHEESH, S., MA, S., HUANG, Z., KARPATY, A., KHOSLA, A., BERNSTEIN, M., BERG, A. C., AND FEI-FEI, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115, 3 (2015), 211–252.
- [48] SCHUHMAN, C., BEAUMONT, R., VENCU, R., GORDON, C. W., WIGHTMAN, R., CHERTI, M., COOMBES, T., KATTA, A., MULLIS, C., WORTSMAN, M., SCHRAMOWSKI, P., KUNDURTHY, S. R., CROWSON, K., SCHMIDT, L., KACZMARCZYK, R., AND JITSEV, J. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track* (2022).
- [49] SHUTTLEAI. shuttleai/shuttle-3-diffusion. <https://huggingface.co/shuttleai/shuttle-3-diffusion/blob/main/vae/config.json>, 2024. Accessed: May 27, 2025.
- [50] SOHL-DICKSTEIN, J., WEISS, E., MAHESWARANATHAN, N., AND GANGULI, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning* (2015), PMLR, pp. 2256–2265.
- [51] SONG, J., MENG, C., AND ERMON, S. Denoising diffusion implicit models. In *International Conference on Learning Representations* (2021).
- [52] STABILITY AI. DreamStudio - Image generation using Stable Diffusion. <https://beta.dreamstudio.ai/generate>. Accessed: May 27, 2025.
- [53] SZEGEDY, C., VANHOUCHE, V., IOFFE, S., SHLENS, J., AND WOJNA, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (6 2016).
- [54] TANCIK, M., MILDENHALL, B., AND NG, R. Stegastamp: Invisible hyperlinks in physical photographs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (6 2020).
- [55] TIRKEL, A. Z., RANKIN, G., VAN SCHYNDEL, R., HO, W., MEE, N., AND OSBORNE, C. F. Electronic watermark. *Digital Image Computing, Technology and Applications (DICTA'93)* (1993), 666–673.
- [56] WEN, Y., KIRCHENBAUER, J., GEIPING, J., AND GOLDSTEIN, T. Tree-rings watermarks: Invisible fingerprints for diffusion images. In *37th Conference on Neural Information Processing Systems (NeurIPS)* (2023).
- [57] YANG, Z., ZENG, K., CHEN, K., FANG, H., ZHANG, W., AND YU, N. Gaussian shading: Provable performance-lossless image watermarking for diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 12162–12171.
- [58] YU, N., SKRIPNIUK, V., ABDELNABI, S., AND FRITZ, M. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *Proceedings of the IEEE/CVF International conference on computer vision* (2021), pp. 14448–14457.

- [59] ZENG, Y., ZHOU, M., XUE, Y., AND PATEL, V. M. Securing deep generative models with universal adversarial signature. *arXiv preprint arXiv:2305.16310* (2023).
- [60] ZHANG, K. A., XU, L., CUESTA-INFANTE, A., AND VEERAMACHANENI, K. Robust invisible video watermarking with attention. *arXiv preprint arXiv:1909.01285* (2019).
- [61] ZHANG, R., ISOLA, P., EFROS, A. A., SHECHTMAN, E., AND WANG, O. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR* (2018).
- [62] ZHAO, X., GUNN, S., CHRIST, M., FAIROZE, J., FABREGA, A., CARLINI, N., GARG, S., HONG, S., NASR, M., TRAMER, F., ET AL. Sok: Watermarking for ai-generated content. *arXiv preprint arXiv:2411.18479* (2024).
- [63] ZHAO, X., ZHANG, K., SU, Z., VASAN, S., GRISHCHENKO, I., KRUEGEL, C., VIGNA, G., WANG, Y.-X., AND LI, L. Invisible image watermarks are provably removable using generative ai, 2023.
- [64] ZHU, J., KAPLAN, R., JOHNSON, J., AND FEI-FEI, L. Hidden: Hiding data with deep networks. In *Proceedings of the European conference on computer vision (ECCV)* (2018), pp. 657–672.

Appendix

A Attack Visualization

From Figure 9, we can see that each attack learns a correlation between the area where the watermark is embedded and the class prediction. As seen in the difference between the Fourier transform of the attacked images and the original images (third column), the surrogates trained on the Wm & UnWm dataset learn a significantly stronger correlation, which is

unexpected given that the difference between the two latent distributions is only the watermark. With surrogate detectors trained on the Wm & Pub dataset, the correlation is weaker, but compared to surrogates trained on raw pixel values, the attack is more concentrated around the key ring, rather than scattered throughout the space. When training on the true latents (at $t = 0$), the correlation between attack area and misclassification becomes stronger; however, it also starts to include features outside the key ring.

When using a fine-tuned version of the victim’s VAE, the SDXL-VAE, we see a more sparse effect of the attack in Fourier space, which is in line with the degradation in performance of the attack. Furthermore, when the victim uses a different VAE, the 16-Channel VAE, there is limited impact of the attack on the Fourier space.

B Precision–Recall Curves

Figure 10 plots the Precision–Recall curves for all the attacks at a base rate of 0.5, namely a balanced dataset. Remarkably, our attack (third column) exerts the most significant impact on the curve, making it the attack in which the defender would achieve the weakest compromise between precision and recall. This is especially relevant in scenarios where the base rate is low, as the additional false positives can have an overwhelming effect on the deployment performance of the detector.

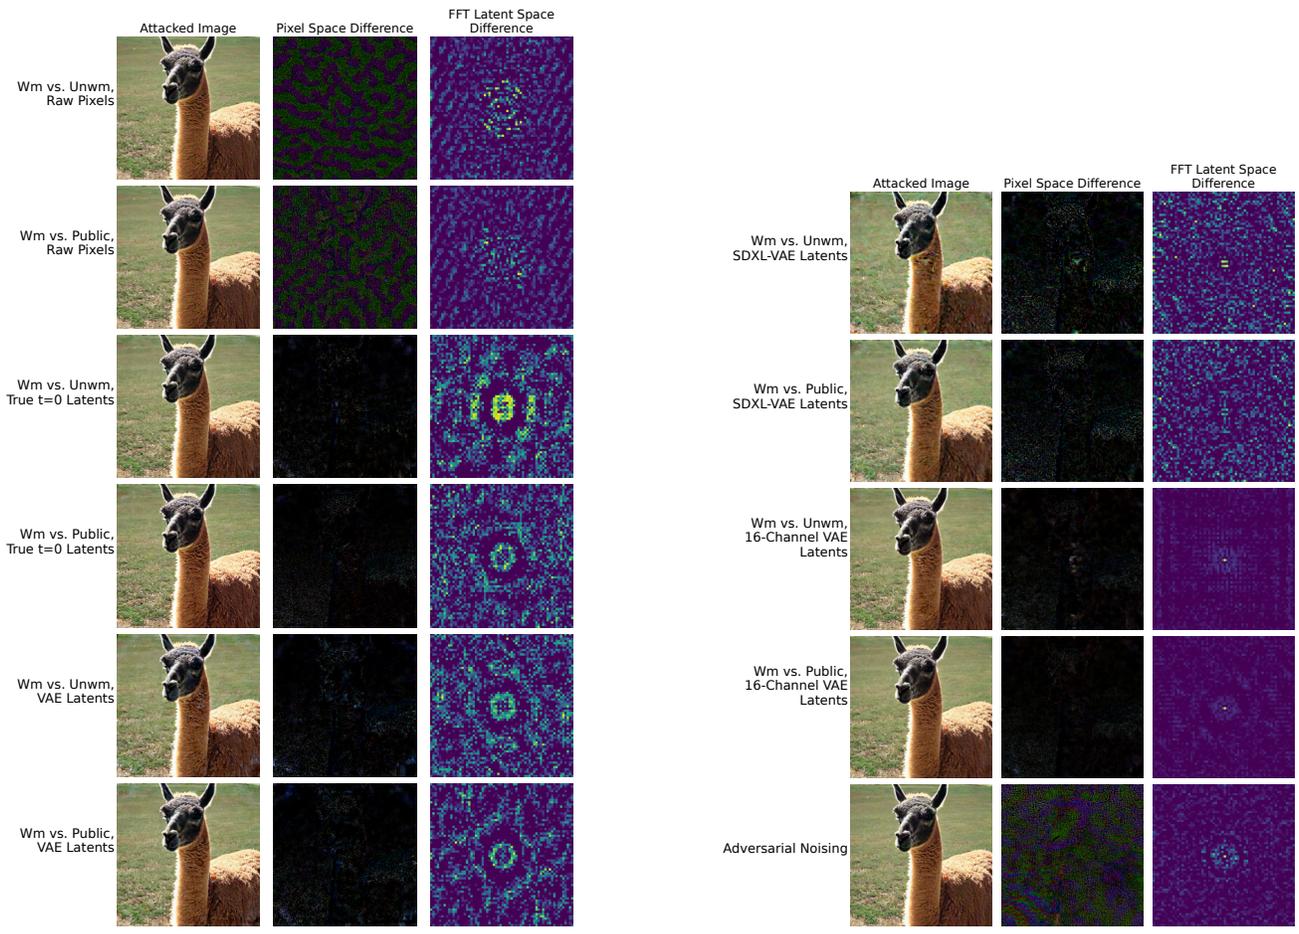


Figure 9: Attacked images. Each row corresponds to a combination of dataset and input type. The columns from left to right are: our augmented images $x_{0_{aug}}$, the changes in pixel values, and the changes in latent space for the watermark channel after FFT.

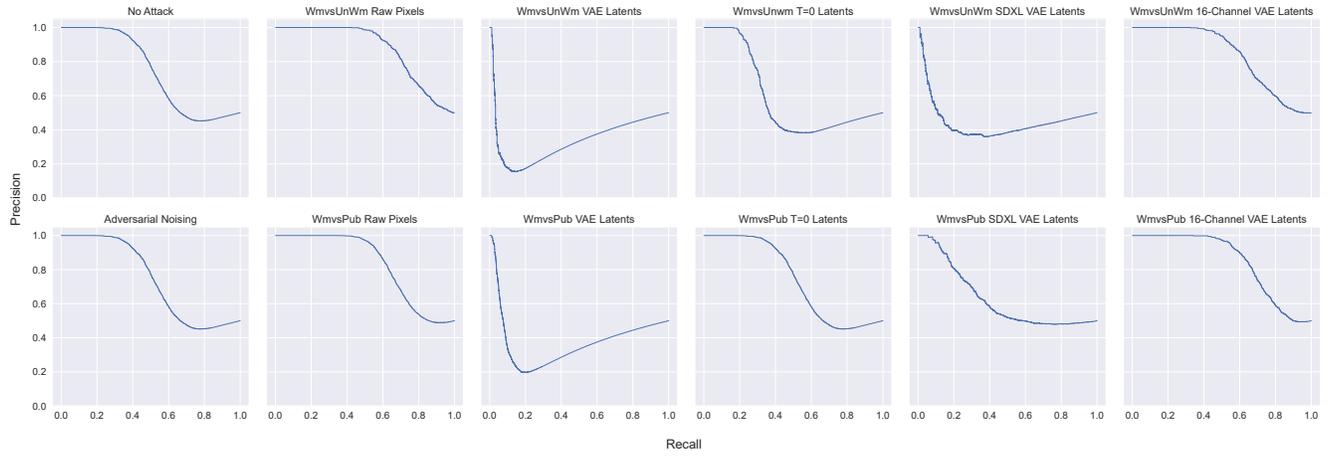


Figure 10: PR curves for all the attacks when the base rate is 0.5. $y = 0.5$ is the PR curve of random guessing.