

LaRE²: Latent Reconstruction Error Based Method for Diffusion-Generated Image Detection

Yunpeng Luo Junlong Du Ke Yan* Shouhong Ding
Tencent YouTu Lab

{petterluo, jeffdu, kerwinyan, ericshding}@tencent.com

Abstract

The evolution of Diffusion Models has dramatically improved image generation quality, making it increasingly difficult to differentiate between real and generated images. This development, while impressive, also raises significant privacy and security concerns. In response to this, we propose a novel **Latent RE**construction error guided feature **RE**finement method (LaRE²) for detecting the diffusion-generated images. We come up with the Latent Reconstruction Error (LaRE), the first reconstruction-error based feature in the latent space for generated image detection. LaRE surpasses existing methods in terms of feature extraction efficiency while preserving crucial cues required to differentiate between the real and the fake. To exploit LaRE, we propose an Error-Guided feature REfinement module (EGRE), which can refine the image feature guided by LaRE to enhance the discriminativeness of the feature. Our EGRE utilizes an align-then-refine mechanism, which effectively refines the image feature for generated-image detection from both spatial and channel perspectives. Extensive experiments on the large-scale GenImage benchmark demonstrate the superiority of our LaRE², which surpasses the best SoTA method by up to **11.9%/12.1%** average ACC/AP across 8 different image generators. LaRE also surpasses existing methods in terms of feature extraction cost, delivering an impressive speed enhancement of **8 times**.

1. Introduction

The rapid advancement of Diffusion Models has heralded a new era in the domain of image generation. Through concerted efforts in refining model architecture [9, 32, 33], optimizing training strategies [15, 27], and enhancing sampling methods [20, 34], contemporary Diffusion Models are now capable of generating images of unprecedented quality, surpassing the boundaries of human imagination. However, this progress raises significant concerns regarding privacy

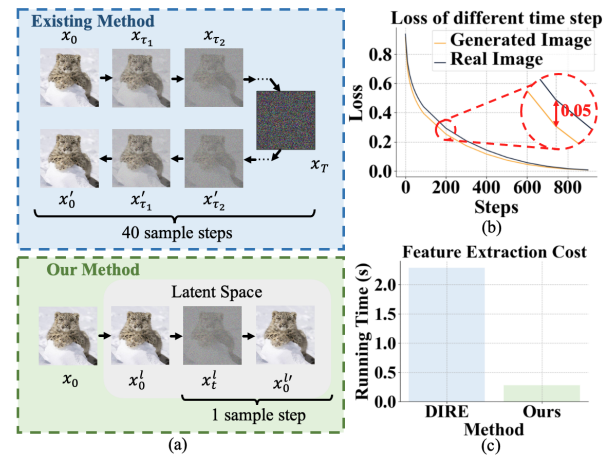


Figure 1. (a) The comparison of reconstruction-based feature extraction. Existing method [40] chooses to completely reconstruct an image by first gradually adding noise to the image and then denoising it, which involves dozens of sampling steps. Our method can directly calculate noisy images and denoise them with a single sample step. (b) Statistical analysis of the relationship between the single-step reconstruction loss (1000 images are used) and time step. The obvious gap between the two lines indicates that single-step reconstruction can also reflect the differences between real and generated images. (c) Comparison of the cost of per image feature extraction. Our method is **8x** faster than DIRE [40].

and security associated with the generated images [16]. The potential for the dissemination of toxic content and misinformation through these images poses a threat to society and could mislead the public. Consequently, there is an urgent need to develop techniques to detect the images generated by these models.

A flurry of works [6, 7, 21, 31, 40, 41] has been proposed to detect or study the properties of images generated by diffusion models. Recently, DIRE [40] has been proposed to leverage the reconstruction error as a discriminative feature for diffusion-generated image detection. It is based on the assumption that the diffusion-generated images are more easily to be reconstructed by a diffusion

*Corresponding author

model compared with real images. As a result, DIRE shows great cross-model generalizability towards different diffusion models. Based on DIRE, SeDID [21] is also proposed to harness the inherent distributional disparities between naturally occurring and diffusion-synthesized visuals for diffusion-generated image detection.

However, both DIRE and SeDID require multi-step DDIM [34] sampling processes in the feature extraction stage, which results in low efficiency for real-world applications. As shown in Fig 1(c), it takes more than 2 seconds to extract the DIRE feature for an image. Besides, the accumulation loss in the multi-step sampling also introduces uncertainty into the extracted features. Though DDIM [34] provides a deterministic inversion method to transform an image into noise, the reliability of this inversion process is not consistently upheld. In addition, DIRE and SeDID use the reconstruction error as the only feature, ignoring the correspondence between the error and the raw image. We therefore ask: (1) Do we need to completely reconstruct the image to get the discriminative feature? (2) Are there better ways to incorporate reconstruction error into image generation detection?

To address the two questions mentioned above, we conducted a series of exploratory experiments and analyses. For the first question, we investigated the training process of the diffusion model. We found that given any time step t , the forward Markov process has a closed-form solution, allowing us to directly transform x_0 into x_t . Additionally, during the training phase, the model takes x_t and t as inputs and computes the denoising loss, enabling us to obtain the loss through single-step denoising. Therefore, we perform single-step denoising on both 1000 real and 1000 generated images respectively by LDM [32]. The results, as shown in Fig. 1(b), indicate that even with only single-step reconstruction, the loss of real images is consistently greater than that of generated images. This not only demonstrates that the real images are harder to reconstruct than the generated ones but also suggests that the loss from single-step reconstruction can reflect the differences between real and generated images. To address the second question, we visualize the reconstruction loss on the original image, as shown in Fig. 2. Upon observation, we have found that the reconstruction loss is positively correlated with the local information frequency of the original image. In the first image, for instance, the reconstruction loss is relatively lower on the low-frequency background, while it is relatively higher on the high-frequency foreground. The same pattern can be observed in other images as well. The reconstruction loss exhibits spatial correlations with the original image, thereby presenting a potential to serve as a valuable cue for the generated image detection.

Based on the findings above, we propose a novel Latent REconstruction error guided feature REfinement method




Figure 2. Visualization of reconstruction loss on raw images (i.e. Image + Loss). Though the randomly sampled noises are added to the whole image, there is a trend that the loss in high-frequency regions is typically greater than that in low-frequency regions.

(LaRE²) for diffusion-generated image detection. Our LaRE² consists of two parts, the Latent Reconstruction Error (LaRE) and the Error-guided Feature REfinement module (EGRE). LaRE is a more efficient reconstruction feature, which has two improvements compared with existing methods: (1) LaRE is extracted in a single step of the diffusion reverse process, which is significantly more efficient than completely reconstructing the image through dozens of denoising steps. (2) We conduct the reconstruction in the latent space, which further improves the efficiency. Besides, we find that the reconstructed loss is positively correlated with the local information frequency of the original image. Based on that, we come up with the Error-Guided Feature Refinement module (EGRE). In EGRE, we first align LaRE with the image feature map for better correspondings for feature refinement. Then the aligned LaRE is used to refine the image feature in both spatial and channel perspectives. With the guidance of LaRE, EGRE refines the feature map in both spatial and channel perspectives to better reveal the discriminativeness for detecting the generated images.

To evaluate the effectiveness of our method, we conduct extensive experiments on the GenImage [47] benchmark, which comprises 2,681,167 images, segregated into 1,331,167 real and 1,350,000 fake images. The fake images are from 8 different generators. As shown in Fig. 1(c), our LaRE is **8 times faster** than existing method [40]. In addition, LaRE² achieves a significant performance gain by up to **11.9%/12.1% ACC/AP** compared with the best SoTA. The results demonstrated the superiority of our method, which is both effective and generalizable.

The contributions of our work are three-fold:

- **Novel feature:** We are the first to propose the reconstruction error in latent space for generated-image de-

tection. Compared with the existing method, we remarkably reduce the cost of feature extraction while preserving the essential information required for the detection of diffusion-generated images.

- **Novel module:** We qualitatively analyze the reason for the effectiveness of the reconstruction loss. Based on that, we come up with a novel module EGRE, which conducts an Error-guided feature refinement to enhance the discriminativeness of image features.
- **Superior performance:** Extensive experiments demonstrate the effectiveness of our method. We achieve **11.9%/12.1%** ACC/AP gain on the large-scale GenImage benchmark, significantly outperforming the SoTA methods.

2. Related Works

2.1. Diffusion Model

Diffusion Models [14, 33, 43] have become new state-of-the-art deep generative models. Inspired by nonequilibrium thermodynamics, Denoising Diffusion Probabilistic Models (DDPM) demonstrate promising generative quality, which inspires a flurry of studies on diffusion models [3, 5, 9, 32, 34, 35, 42, 46]. DDIM [34] generalize DDPMs [14] via a class of non-Markovian diffusion processes to accelerate sampling. ADM [9] firstly obtain better generation quality than GANs [11] on ImageNet [8]. Latent Diffusion Models [32] apply the diffusion process in the latent space to improve the efficiency of diffusion models, and enable diffusion models with text conditioning inputs through the cross-attention conditioning mechanism. There are also works that employ pre-trained diffusion models for various downstream tasks [3, 5, 35, 42, 46]. The rapid advancement of deep generative models has given rise to concerns regarding the potential for malicious utilization of the generated images [16]. Consequently, it is urgent to develop robust techniques to detect generated images.

2.2. Deep Generated Image Detection

Deep-generated image detection has achieved significant improvements due to the contributions of previous research. Initially, researchers attempted to use hand-crafted features including color cues [23], saturation cues [24], blending [17] artifacts, co-occurrence features [25]. Then CNN is leveraged to detect the generated images [18, 22, 38]. Some works [10, 44] also find that there are obvious visual artifacts in the GAN-generated images and detect these images from the frequency view [28]. Most of the above works are developed and designed for GAN-generated images. With the advance of diffusion models, many methods are also provided for the study of diffusion-generated image detection. Corvi *et. al* [7, 31] find that state-of-the-art detec-

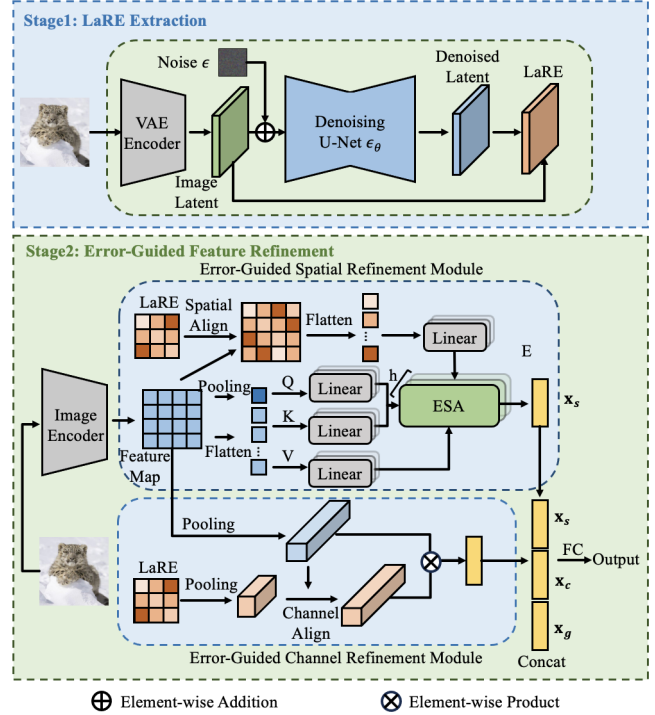


Figure 3. Overview of our method. In the first stage, we extract LaRE in the latent space through single-step reconstruction. In the second stage, to exploit LaRE, we propose the Error-guided Feature Refinement Module, which consists of the Error-guided spatial refinement module and the Error-guided Channel Refinement module. From both spatial and channel perspectives, LaRE is used to enhance the discriminativeness of the image feature for generated image detection.

tors developed for GAN suffer from a severe performance drop when applied to Diffusion-generated images. Corvi *et. al* [6, 7, 31] also find that the spectrum artifacts can also rise up in the Diffusion-generated images. Wu *et. al* [41] formulate the synthetic image detection as an identification problem and achieve generalizable detection through language-guided contrastive learning. DIRE [40] and SeDIE [21] are proposed to leverage reconstruction error from the diffusion model to achieve generalizable diffusion-generated image detection. Different from these methods, we are the first to show that reconstruction error in the latent space can also benefit diffusion-generated image detection, which is more efficient and generalizable.

3. Preliminaries

Diffusion Models has achieved remarkable image generation performance. Typically, it involves two Markov processes. In the forward process, Gaussian noise is gradually added to the raw image x_0 until the image is asymptotically transformed to pure noise, which is defined as:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = N(\mathbf{x}_t; \sqrt{\frac{\alpha_t}{\alpha_{t-1}}}\mathbf{x}_{t-1}, (1 - \frac{\alpha_t}{\alpha_{t-1}}\mathbf{I})), \quad (1)$$

where x_t is the noisy image at step t and α_t is predefined noise schedule. According to the properties of Markov process and Gaussian distribution, we can get \mathbf{x}_t from \mathbf{x}_0 directly by:

$$q(\mathbf{x}_t|\mathbf{x}_0) = N(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}_0, (1 - \alpha_t)\mathbf{I}). \quad (2)$$

In the reverse process, the noisy image is gradually denoised to get the raw image, which is defined as:

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = N(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)), \quad (3)$$

where p_θ is typically parameterized by neural networks. During training, a neural network ϵ_θ is trained to predict the added noise ϵ , given the noisy image x_t and corresponding time step t :

$$L_\theta(\mathbf{x}_0, t) = \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|^2. \quad (4)$$

During training, t and ϵ are randomly sampled.

4. Methods

In this section, we elaborate on each component of our LaRE², as illustrated in Fig. 3.

4.1. Latent Reconstruction Error

Recent methods [40] propose to use reconstruction error as the feature for generated-image detection. Through DDIM [34] inversion, an image is first inverted to noise and then re-generated based on the noise. However, this method has several limitations: (1) It is relatively slow to inverse and re-generate an image. For example, it takes more than 2 seconds to extract DIRE for an image on a Tesla V100 GPU as shown in Fig. 1(c), which limits its realistic applications; (2) The reconstruction of the image is not always reliable. Error could accumulate during the Markov forward and reverse process, making it hard to tell whether the bad reconstruction is from the difference between real and fake images, or the reconstruction itself.

Here we come up with a new reconstruction feature, named Latent Reconstruction Error (LaRE). LaRE is based on the assumption that if the generated images can be completely reconstructed with relative ease [40], then they can also be more easily reconstructed at every single step of the reverse diffusion process.

For an image x , we first get its latent code x_0 by VAE. Then we calculate the LaRE by:

$$\text{LaRE} = L_\theta(\mathbf{x}_0, t) = \mathbb{E}_\epsilon \left[\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|^2 \right]. \quad (5)$$

As shown in Fig. 3, Our LaRE takes advantage of two special properties of Diffusion Models. (1) The forward process has a close form. According to Eq. 2, we can get \mathbf{x}_t directly from \mathbf{x}_0 ; (2) Diffusion Model is trained to denoise a noisy image from any time step t , as shown in Eq. 4. Therefore, given \mathbf{x}_t , t , and ϵ , we can directly get the reconstruction error by only single-step denoising, which is much more efficient than completely reconstructing an image. Compared with complete image reconstruction, our LaRE has several advantages: (1) Only single-step denoising is required to get LaRE, which significantly improves its efficiency. (2) Reconstruction is conducted in the latent space, which is also more efficient. (3) No need for complete inversion and reconstruction of the image, which eliminates the potential accumulation of errors that may occur during multiple inversion and reconstruction steps.

With efficiency, the easier reconstruction assumption still holds for our LaRE. Recall that a diffusion-generated image is more easily reconstructed by diffusion models [40]. So the reconstruction error should be relatively small when completely reconstructing an image through the whole reverse process. That implies the reconstruction error is also relatively small during each step of the reverse process, which is also validated in Fig. 1(b). We leverage the Monte Carlo method to estimate LaRE in Eq. 6:

$$\text{LaRE} = \frac{1}{e} \sum_i^e \left[\|\epsilon_i - \epsilon_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_i, t)\|^2 \right], \quad (6)$$

where e , t are pre-defined hyperparameters. Specifically, we conduct denoising for e times and average the denoising loss to extract LaRE.

4.2. Error-Guided Feature Refinement

Existing methods exploit the reconstruction error as the only feature to detect the generated image, which ignores the relationships between the reconstruction error and the raw image. To better understand how reconstruction error helps generated-image detection, we visualize the extracted LaRE on the raw image, as shown in Fig. 2. We find several interesting conclusions. 1) The loss norm is proportional to the noise norm. It is easy to understand that the stronger the added noise is, the harder to reconstruct; 2) The loss norm is also proportional to the frequency of local image patches. That the low-frequency patches are easier to reconstruct than the high-frequency ones. Therefore, the reconstruction performance on the high-frequency parts of the image plays an important role in the generated image detection. To apply this valuable information, we come up with a novel Error-guided feature REfinement Module (EGRE). Based on the reconstruction error, our EGRE can better refine the image feature to reveal the discriminative parts for generated image detection.

Train on	Test on									Test on																	
	MJ	SDV1.4	SDV1.5	ADM	GLIDE	WuKong	VQDM	Biggan	Average	MJ	SDV1.4	SDV1.5	ADM	GLIDE	WuKong	VQDM	Biggan	Average	MJ	SDV1.4	SDV1.5	ADM	GLIDE	WuKong	VQDM	Biggan	Average
MJ	99.5	81.9	81.7	74.2	76.8	83.1	60.8	49.7	76.0	100.0	88.2	88.2	67.1	90.7	77.3	73.6	66.0	81.4	100.0	97.9	97.8	94.0	96.6	97.5	85.2	53.7	90.3
SDV1.4	65.8	99.7	99.7	54.5	58.1	99.4	54.3	49.8	72.7	74.0	100.0	99.9	61.7	88.5	100.0	97.2	68.7	86.2	92.7	100.0	100.0	84.5	88.2	100.0	82.8	55.0	87.9
SDV1.5	66.4	99.7	99.7	55.0	56.5	99.4	55.2	49.8	72.7	64.8	100.0	100.0	61.9	91.4	100.0	94.6	63.7	84.5	92.3	100.0	100.0	85.6	87.5	100.0	85.5	48.5	87.4
ADM	72.6	58.5	58.5	99.2	95.4	61.6	79.3	49.8	71.9	63.4	84.9	83.7	100.0	99.0	90.8	92.0	64.0	84.7	92.7	83.4	83.1	100.0	99.2	84.9	96.9	48.3	86.1
GLIDE	50.8	50.1	50.1	56.6	99.8	50.1	53.6	54.6	58.2	59.7	75.5	75.9	74.0	99.8	66.8	71.5	88.7	76.5	66.3	55.2	55.5	82.5	100.0	51.9	81.3	90.3	72.9
WuKong	64.6	99.4	99.3	54.0	55.8	99.7	54.3	49.9	72.1	52.8	100.0	99.9	54.0	58.2	100.0	96.3	68.0	78.7	91.8	100.0	100.0	85.3	88.3	100.0	82.3	62.5	88.8
VQDM	50.3	50.5	50.5	51.3	59.6	51.3	99.9	50.3	58.0	66.2	99.5	99.4	64.9	72.6	99.1	99.9	72.7	84.3	68.3	72.4	72.7	83.9	94.4	78.4	100.0	63.7	79.2
Biggan	50.0	50.0	50.0	50.0	51.3	50.0	100.0	56.4	56.4	50.0	50.0	50.0	50.0	50.0	50.0	50.0	100.0	56.2	51.9	49.1	49.4	58.0	95.8	43.7	72.2	100.0	65.0
Average	65.0	73.7	73.7	61.9	69.1	74.3	63.4	56.7	67.2	66.4	87.3	87.1	66.7	81.3	85.5	84.4	74.0	79.1	82.0	82.3	82.3	84.2	93.7	82.0	85.8	65.2	82.2

(a) Comparison on ACC. Left: baseline, right: ours

(b) Comparison on AP. Left: baseline, right: ours

Figure 4. Results of cross-validation on different training and testing subsets. For each generator, we train a model and test it on all 8 generators. For both DIRE [40] and our method, accuracy (ACC) and average precision (AP) are reported.

4.2.1 Error-guided Spatial Refinement Module

We propose the Error-guided Spatial Refinement module (ESR) for feature refinement from the spatial perspective. First, we spatially align LaRE and the image feature map by a simple adaptive average pooling layer, so that they have the same spatial size. Now we have the spatially aligned input feature map $\mathbf{x} \in \mathbb{R}^{HW \times C_1}$ and LaRE $\mathbf{e} \in \mathbb{R}^{HW \times C_2}$. $\bar{\mathbf{x}} \in \mathbb{R}^{1 \times C_1}$ is the global feature, which is the average of \mathbf{x} . To refine the feature from the spatial perspective, we choose an Error-guided Spatial Attention module (ESA):

$$\text{ESA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{E}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} + \mathbf{E}\right)\mathbf{V}. \quad (7)$$

To further enhance the feature representation, we adopt a multi-head mechanism from Transformer [37]. Therefore, our Multi-Head Error-guided Spatial Attention module (MHESA) is defined as:

$$\text{MHESA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{E}) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)\mathbf{W}^O, \quad (8)$$

where $\text{head}_i = \text{ESA}(\mathbf{Q}\mathbf{W}_i^Q, \mathbf{K}\mathbf{W}_i^K, \mathbf{V}\mathbf{W}_i^V, \mathbf{E}\mathbf{W}_i^E)$.

where h is the number of heads, $\mathbf{W}_i^Q \in \mathbb{R}^{C_1 \times d}$, $\mathbf{W}_i^K \in \mathbb{R}^{C_1 \times d}$, $\mathbf{W}_i^V \in \mathbb{R}^{C_1 \times d}$, $\mathbf{W}_i^E \in \mathbb{R}^{C_2 \times d}$ are learnable projection matrices, d is the hyperparameter. The spatial refined feature x' is calculated by:

$$\mathbf{x}_s = \text{MHESA}(\bar{\mathbf{x}}, \mathbf{x}, \mathbf{x}, \mathbf{e}), \quad (9)$$

where $\bar{\mathbf{x}}$ is the only query attending to the feature map \mathbf{x} . In this stage, the aligned LaRE is used to re-weight the attention score of scaled dot-product attention. Since LaRE and the feature map are spatially aligned, LaRE can emphasize important information from the spatial perspective.

4.2.2 Error-guided Channel Refinement Module

For channel refinement, we first squeeze both \mathbf{x} and \mathbf{e} to $\bar{\mathbf{x}} \in \mathbb{R}^{1 \times C_1}$ and $\bar{\mathbf{e}} \in \mathbb{R}^{1 \times C_2}$. Then our channel refinement

is achieved by a gate mechanism:

$$\mathbf{x}_c = \text{sigmoid}(\bar{\mathbf{e}}\mathbf{W}) \odot \bar{\mathbf{x}}, \quad (10)$$

where $\mathbf{W} \in \mathbb{R}^{C_1 \times C_2}$ are learnable parameters to align LaRE and the feature map from the channel perspective. Finally, \mathbf{x}_s , \mathbf{x}_c and the original global feature \mathbf{x}_g from the last convolution block are concatenated as the final feature \mathbf{x}_{EGRE} .

$$\mathbf{x}_{\text{EGRE}} = \text{Concat}(\mathbf{x}_s, \mathbf{x}_c, \mathbf{x}_g). \quad (11)$$

\mathbf{x}_{EGRE} is followed by one FC layer. We train our model with the binary cross entropy loss.

5. Experiments

5.1. Datasets and Evaluation Metrics

We evaluate our proposed method using the GenImage [47] dataset. This dataset comprises a total of 2,681,167 images, divided into 1,331,167 real images and 1,350,000 generated images. The generated images are from 8 different generative models, namely BigGAN [4], GLIDE [26], VQDM [12], Stable Diffusion V1.4&V1.5 [32], ADM [9], Midjourney [1], and Wukong [2]. All of the fake images are generated using the template prompt ‘‘photo of CLS’’, where ‘‘CLS’’ is replaced by one of the 1000 labels from ImageNet [8]. These images are split into 8 subsets where each subset contains partial real images and all the fake images generated by one of the above generators. We adhere to the official division of the dataset in our research, allocating 2,581,167 images for training and reserving the remaining 100,000 images for validation. Following DIRE [40], we employ Accuracy (ACC) and Average Precision (AP) as our evaluation metrics. More details and results are included in our supplementary materials.

5.2. Implementation Details

All experiments are conducted using the GenImage dataset. To obtain the LaRE, we employed Stable Diffu-

Methods	Testing Subset								Avg ACC.(%)
	Midjourney	SDV1.4	SDV1.5	ADM	GLIDE	Wukong	VQDM	BigGAN	
CNNSpot [39]	58.2	70.3	70.2	57.0	57.1	67.7	56.7	56.6	61.7
Spec [45]	56.7	72.4	72.3	57.9	65.4	70.3	61.7	64.3	65.1
F3Net [29]	55.1	73.1	73.1	66.5	57.8	72.3	62.1	56.5	64.6
GramNet [19]	58.1	72.8	72.7	58.7	65.3	71.3	57.8	61.2	64.7
DIRE [40]†	65.0	73.7	73.7	61.9	69.1	74.3	63.4	56.7	67.2
Ours	66.4	87.3	87.1	66.7	81.3	85.5	84.4	74.0	79.1

Table 1. Performance comparisons on GenImage test set. For each data point, eight models are trained on eight generators respectively. Then eight models are tested on the specified test sets and the accuracy scores are averaged. † indicates our reproduction.

sion V1.5 [32] with a step size of $t = 200$ and a noise ensemble size of $e = 4$. Images are resized to 256×256 for LaRE extraction. We use the prompt “a photo” for all images. For both training and testing phases, images are resized to 224×224 . Our model utilizes the CLIP [30] pre-trained ResNet50 [13] as the backbone architecture. We set the batch size to 48 and the learning rate to $1e^{-4}$. Training is performed on a single Nvidia Tesla V100 GPU. We train eight models on eight subsets respectively, each corresponding to a different generation method. Model selection is based on performance in the validation set, which shares the same generation method as the training set.

5.3. Cross-Generator Image Classification

We evaluate the performance of our model trained by one of the subsets and test on all eight subsets. The results are shown in Fig. 4. Several conclusions can be made from the results. (1) It is easy to detect generated images from seen generators. Both the DIRE and our proposed method demonstrate superior performance when training and testing images share the same generator, as shown by the diagonal of each image. (2) It is challenging to generalize the model to unseen generators. However, it is relatively easier to generalize the model to similar generator structures than to different structures. For example, a model trained on the BigGAN subset achieves bad performance when tested on other subsets, which are all diffusion-based generators. While the model trained on the SDV1.5 subset achieves relatively better performance on other diffusion-based models. (3) Our model is more robust and generalizable than DIRE. Our model trained on diffusion-based generators achieves better generalizability on both gan-based and diffusion-based generators. This indicates that our LaRE² can better capture discriminative features to detect the generated images.

5.4. Compare with the State of the Arts

In this section, we compare our method with several state-of-the-art generated image detection methods, including CNNSpot [39], Spec [45], F3Net [29], GramNet [19]

and DIRE [40]. In accordance with the GenImage benchmark, eight models are trained on each subset. Subsequently, these models are evaluated on the specified test set, and their accuracy scores are averaged to produce a data point for inclusion in the table. This benchmark presents a formidable challenge, as the overall accuracy is closely tied to the cross-generator generalizability of the proposed techniques. As illustrated in Tab. 1 and Fig. 4, our approach attains state-of-the-art performance across all eight generators. Notably, our overall accuracy exhibits a substantial improvement, surpassing the previous state-of-the-art by 11.9%/12.1% ACC/AP. This result demonstrates the superior generalizability of our method.

5.5. Ablation Study

In this section, we conduct several ablative studies to quantify the contribution of each design in our model and the influence of the hyperparameters.

5.5.1 Influence of EGRE

The contribution of individual components in EGRE to the overall performance of the method is reported in Tab. 2. Starting with the most simple baseline, *i.e.*, training without LaRE (Model A in the table), each component is added building up to the proposed approach. We measure the effect of the Error-guided Spatial Refinement module (ESR) and Error-guided Channel Refinement module (ECR). Specifically, ESR achieves 8.6%/5.3% ACC/AP gain and ECR achieves 5.7%/2.4% ACC/AP gain compared with the baseline. Both ESR and ECR significantly improve the performance, demonstrating the effectiveness of the two modules. Relatively, ESR brings a higher improvement in metrics compared with ECR, indicating a greater benefit from spatial feature refinement. This is also attributed to the better alignment of LaRE with the original image at the spatial scale as shown in Fig. 2. By combining ESR and ECR, our approach achieves a further improvement of 11.4%/7.2% ACC/AP. By leveraging LaRE for feature refinement at both spatial and channel scales, our model

Model	ESR	ECR	CLS	MJ	SDV1.4	SDV1.5	ADM	GLIDE	WuKong	VQDM	Biggan	Avg
A				52.9/94.4	100.0/100.0	100.0/100.0	55.0/89.5	75.5/98.7	99.7/100.0	51.8/84.4	50.4/71.7	73.1/92.3
B	✓			61.8/96.9	100.0/100.0	100.0/100.0	60.3/95.4	87.2/99.1	100.0/100.0	86.2/94.6	58.1/95.0	81.7/97.6
C		✓		57.4/95.3	100.0/100.0	100.0/100.0	57.2/91.6	81.2/98.9	100.0/100.0	78.3/89.8	56.2/82.3	78.8/94.7
D	✓	✓		64.8/98.2	100.0/100.0	100.0/100.0	61.9/98.5	91.4/99.8	100.0/100.0	94.6/100.0	63.7/99.7	84.5/99.5
E	✓	✓	✓	66.2/99.3	100.0/100.0	100.0/100.0	64.5/99.5	91.5/99.9	100.0/100.0	97.7/100.0	67.4/99.8	85.9/99.8

Table 2. Ablative studies on using different modules or prompts. ‘CLS’ indicates to use class-specific prompt for LaRE extraction.

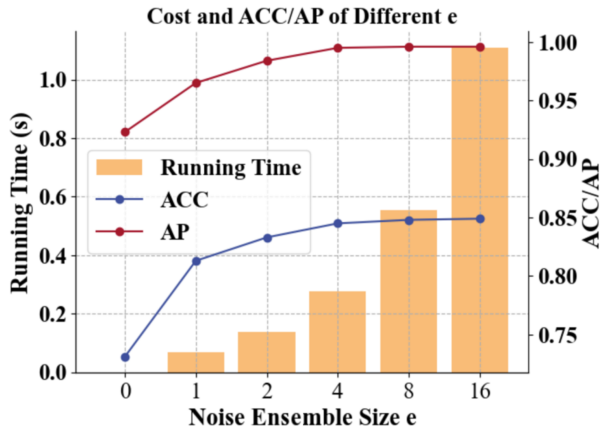


Figure 5. Trade-off between detection performance and feature extraction cost. When $e=4$, the model achieves the best trade-off.

achieves more robust and generalizable image generation detection performance.

5.5.2 Influence of Noise Ensemble

We systematically evaluated the impact of the noise ensemble parameter e on both the accuracy and feature extraction runtime. As shown in Fig. 5. The results demonstrate a consistent increase in accuracy as e increased. However, this improvement in accuracy comes at the cost of increased feature extraction runtime. Upon careful consideration of the trade-off between accuracy and computational cost, we propose $e = 4$ as the optimal choice. This selection of e provides an effective balance, maintaining the performance of our model while ensuring computational efficiency.

5.5.3 Influence of Prompts

Because Latent diffusion is a text-conditioned model, a prompt is needed to denoise an image. Here we try to use a class-specific prompt instead of an unchanged prompt. We choose the template “a photo of CLS” where “CLS” is replaced with the ImageNet label of the image. Results are shown in Tab. 2. When using class-specific prompts (Model E), we can achieve further improvements by 1.4%/0.3%

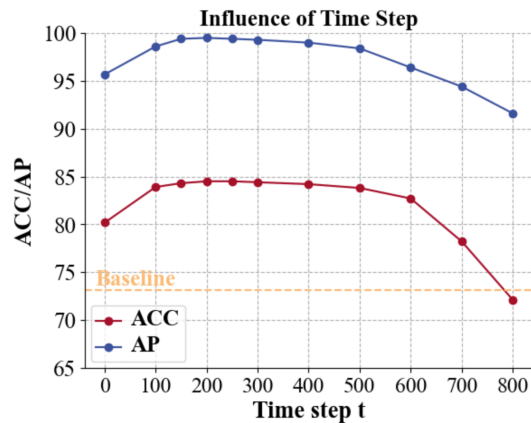


Figure 6. ACC and AP when choosing different t . The results demonstrate that our model is robust to the choice of t .

ACC/AP. This suggests that the category information of images can aid in the extraction of LaRE features. Even in the absence of category information, it is interesting that the accuracy does not decline significantly. We speculate this is because we did not fully convert the images into noise. Even with the added noise, the image still retains some of the original important semantic information, which can assist the model in denoising and thus obtaining good LaRE features. This experiment validates that when the textual information of an image is known, the extracted LaRE features can further enhance the performance. However, we still use the generic prompt in our method, as it already achieves satisfactory results. As for accurately describing images to further enhance performance, we leave that for future research work.

5.5.4 Influence of Sample Step

Here we conduct the ablative experiment on the sample step t . As illustrated in Fig. 1(b), there is a gap between the reconstruction loss of real images and the generated images. When $t \in [100, 500]$, there is a more obvious difference in the reconstruction loss between the two types of images. To verify the impact of t selection on the performance, we conducted experiments with different values of t , and the

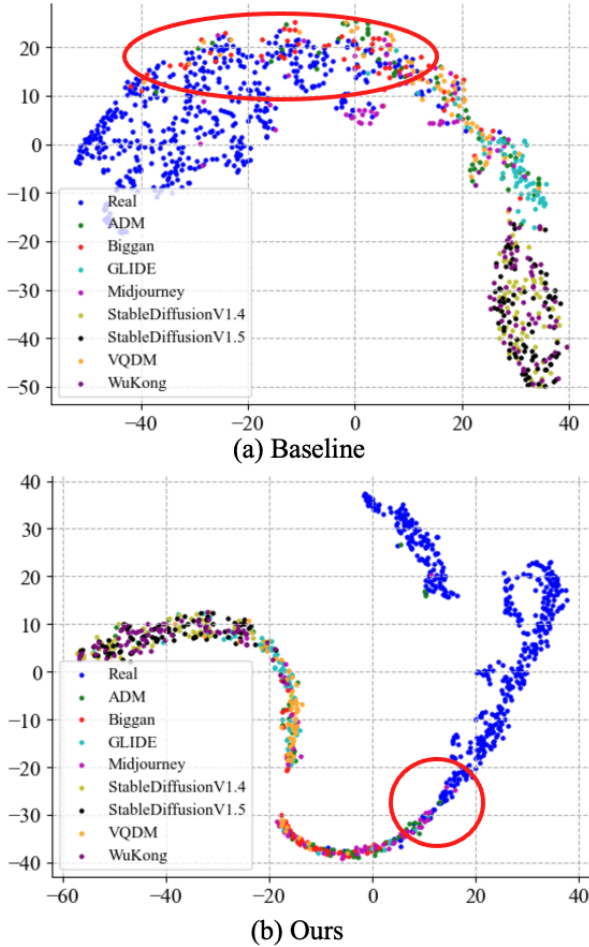


Figure 7. Visualization of the spatial representation of the baseline model (a) and our model (b). Both models were trained on the SDV1.5 subset. In the case of seen data, both models exhibit good discrimination in the feature space. However, for unseen data, the baseline model shows a significant overlap between real images and fake images (red circle), whereas our model exhibits a smaller overlap (red circle). This further indicates the better generalization capability of our model.

results are shown in Fig. 6. When $t \in [150, 300]$, the model exhibits stable performance, indicating the robustness of our model with respect to t . However, when t exceeds this range, the performance of the model tends to decline.

5.5.5 Influence of Input Information

We conduct experiments to train the detection model with different inputs. The results are shown in Tab. 3. Since our LaRE is a compressed latent representation whose size is only $32 \times 32 \times 4$. We train a Resnet-20 with LaRE as the only input (Model B). The experimental results show that both the AP and ACC metrics have decreased significantly. We speculate that there are two reasons. Firstly,

Model	Image	LaRE	Method	AVG. ACC	AVG. AP
A	✓		None	73.1	92.3
B		✓	None	66.2	68.3
C	✓	✓	Concat	76.8	93.5
D	✓	✓	EGRE	84.5	99.5

Table 3. Results of using different input. Models are trained on the SDv1.5 subset and tested on all the subsets. ‘AVG’ indicates average scores over 8 subsets. ‘Concat’ indicates concat LaRE with the image feature map after the last convolution block.

LaRE is in the latent space and has more information loss compared to the original image. It is not sufficient to serve as the only information for detecting the generated images. Secondly, as shown in the Fig. 2, LaRE mainly reflects the reconstruction quality of the original image at different spatial positions. It is better to leverage LaRE to enhance the representation of the original image features. Therefore, we conducted further experiments. In model C, we only concatenated LaRE with the feature map of the image, which also improves the performance. This validates our hypothesis. Besides, model D achieves significant performance gain compared with Model C. By using EGRE, we can further improve the performance of our model by a significant margin, which also confirms the advantages of EGRE over simple feature concatenation.

5.6. Qualitative Results and Visualizations

We utilize t-SNE [36] visualization to illustrate feature vectors derived from the final layers of our model and the baseline model, as depicted in Fig. 7. Each model is trained on the SDv1.5 subset of GenImage and evaluated across all eight generative methods. Although both models exhibit good feature discrimination on seen images, they both have limitations when it comes to unseen scenarios. For images from MJ, ADM, and BigGAN, both models show overlaps between real and fake images in the feature space. However, our model demonstrates a smaller proportion of overlap, indicating its better generalization ability.

6. Conclusion

In this paper, we propose a novel reconstruction-based diffusion-generated image detection method called LaRE². We come up with LaRE, a novel and more efficient reconstruction-based feature by reconstructing the image in the Latent space. Notably, LaRE is 8 times faster compared with existing reconstruction-based methods. By incorporating LaRE with the Error-guided Feature Refinement module (EGRE). Our LaRE² achieves superior performance on diffusion-generated image detection, demonstrating state-of-the-art performance.

References

- [1] Midjourney. <https://www.midjourney.com/home/>, 2022. 5
- [2] Wukong. <https://xihe.mindspore.cn/modelzoo/wukong>, 2022. 5
- [3] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khruikov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. 2021. 3
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. 2018. 5
- [5] Kevin Clark and Priyank Jaini. Text-to-image diffusion models are zero-shot classifiers. 2023. 3
- [6] Riccardo Corvi, Davide Cozzolino, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. Intriguing properties of synthetic images: from generative adversarial networks to diffusion models. In *CVPR*, pages 973–982, 2023. 1, 3
- [7] Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Verdoliva. On the detection of synthetic images generated by diffusion models. In *ICASSP*, pages 1–5. IEEE, 2023. 1, 3
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 3, 5
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. 34:8780–8794, 2021. 1, 3, 5
- [10] Joel Frank, Thorsten Eisenhofer, Lea Schönherr, Asja Fischer, Dorothea Kolossa, and Thorsten Holz. Leveraging frequency analysis for deep fake image recognition. In *ICML*, pages 3247–3258. PMLR, 2020. 3
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. 63(11):139–144, 2020. 3
- [12] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *CVPR*, pages 10696–10706, 2022. 5
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. 33:6840–6851, 2020. 3
- [15] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. 2022. 1
- [16] Felix Juefei-Xu, Run Wang, Yihao Huang, Qing Guo, Lei Ma, and Yang Liu. Countering malicious deepfakes: Survey, battleground, and horizon. 130(7):1678–1734, 2022. 1, 3
- [17] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face x-ray for more general face forgery detection. In *CVPR*, pages 5001–5010, 2020. 3
- [18] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In *CVPR*, pages 8060–8069, 2020. 3
- [19] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In *CVPR*, pages 8060–8069, 2020. 6
- [20] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. 35:5775–5787, 2022. 1
- [21] Ruipeng Ma, Jinhao Duan, Fei Kong, Xiaoshuang Shi, and Kaidi Xu. Exposing the fake: Effective diffusion-generated images detection. 2023. 1, 2, 3
- [22] Francesco Marra, Diego Gragnaniello, Davide Cozzolino, and Luisa Verdoliva. Detection of gan-generated fake images over social networks. In *MIPR*, pages 384–389. IEEE, 2018. 3
- [23] Scott McCloskey and Michael Albright. Detecting gan-generated imagery using color cues. 2018. 3
- [24] Scott McCloskey and Michael Albright. Detecting gan-generated imagery using saturation cues. In *ICIP*, pages 4584–4588. IEEE, 2019. 3
- [25] Lakshmanan Nataraj, Tajuddin Manhar Mohammed, Shivkumar Chandrasekaran, Arjuna Flenner, Jawadul H Bappy, Amit K Roy-Chowdhury, and BS Manjunath. Detecting gan generated fake images using co-occurrence matrices. 2019. 3
- [26] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. 2021. 5
- [27] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, pages 8162–8171. PMLR, 2021. 1
- [28] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *ECCV*, pages 86–103. Springer, 2020. 3
- [29] Yuyang Qian, Guojun Yin, Lu Sheng, Zixuan Chen, and Jing Shao. Thinking in frequency: Face forgery detection by mining frequency-aware clues. In *ECCV*, pages 86–103. Springer, 2020. 6
- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 6
- [31] Jonas Ricker, Simon Damm, Thorsten Holz, and Asja Fischer. Towards the detection of diffusion model deepfakes. 2022. 1, 3
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 1, 2, 3, 5, 6
- [33] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256–2265. PMLR, 2015. 1, 3
- [34] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. 2020. 1, 2, 3, 4

- [35] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. 2023. [3](#)
- [36] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. 9(11), 2008. [8](#)
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. 30, 2017. [5](#)
- [38] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *CVPR*, pages 8695–8704, 2020. [3](#)
- [39] Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated images are surprisingly easy to spot... for now. In *CVPR*, pages 8695–8704, 2020. [6](#)
- [40] Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, Hezhen Hu, Hong Chen, and Houqiang Li. Dire for diffusion-generated image detection. 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [41] Haiwei Wu, Jiantao Zhou, and Shile Zhang. Generalizable synthetic image detection via language-guided contrastive learning. 2023. [1](#), [3](#)
- [42] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *CVPR*, pages 2955–2966, 2023. [3](#)
- [43] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. 2022. [3](#)
- [44] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *WIFS*, pages 1–6. IEEE, 2019. [3](#)
- [45] Xu Zhang, Svebor Karaman, and Shih-Fu Chang. Detecting and simulating artifacts in gan fake images. In *WIFS*, pages 1–6. IEEE, 2019. [6](#)
- [46] Wenliang Zhao, Yongming Rao, Zuyan Liu, Benlin Liu, Jie Zhou, and Jiwen Lu. Unleashing text-to-image diffusion models for visual perception. 2023. [3](#)
- [47] Mingjian Zhu, Hanting Chen, Qiangyu Yan, Xudong Huang, Guanyu Lin, Wei Li, Zhijun Tu, Hailin Hu, Jie Hu, and Yunhe Wang. Genimage: A million-scale benchmark for detecting ai-generated image. 2023. [2](#), [5](#)