

# Stimulating Diffusion Model for Image Denoising via Adaptive Embedding and Ensembling

Tong Li, Hansen Feng, Lizhi Wang, *Member, IEEE*, Lin Zhu, *Member, IEEE*, Zhiwei Xiong, *Member, IEEE* and Hua Huang, *Senior Member, IEEE*

**Abstract**—Image denoising is a fundamental problem in computational photography, where achieving high perception with low distortion is highly demanding. Current methods either struggle with perceptual quality or suffer from significant distortion. Recently, the emerging diffusion model has achieved state-of-the-art performance in various tasks and demonstrates great potential for image denoising. However, stimulating diffusion models for image denoising is not straightforward and requires solving several critical problems. For one thing, the input inconsistency hinders the connection between diffusion models and image denoising. For another, the content inconsistency between the generated image and the desired denoised image introduces distortion. To tackle these problems, we present a novel strategy called the Diffusion Model for Image Denoising (DMID) by understanding and rethinking the diffusion model from a denoising perspective. Our DMID strategy includes an adaptive embedding method that embeds the noisy image into a pre-trained unconditional diffusion model and an adaptive ensembling method that reduces distortion in the denoised image. Our DMID strategy achieves state-of-the-art performance on both distortion-based and perception-based metrics, for both Gaussian and real-world image denoising. The code is available at <https://github.com/Li-Tong-621/DMID>.

**Index Terms**—Computational Photography, Image Denoising, Diffusion Model, Self-Supervised, Distortion-Perception.

## 1 INTRODUCTION

As smartphones have become ubiquitous, the pursuit of capturing high-quality images has become notably more demanding. Yet, when capturing images under challenging conditions, such as low-light environments, substantial information would be lost due to imaging noise. Consequently, the field of image denoising, with a focus on achieving low distortion and high perception, has been as a vital and thriving area of research.

Traditional methods [1], [2], [3] rely on image priors to guide the denoising process, but their effectiveness is limited under extreme conditions. With the development of deep learning, discriminative methods have become the mainstream method, which are usually trained by pixel-level losses [4], [5], [6], [7], [8]. Actually, the pixel-level losses tend to predict the median (or average) of all possible values rather than the realistic images [9], [10]. Thus, these discriminative methods often struggle with perceptual quality, particularly under extreme noise levels.

To improve perceptual quality, other image restoration tasks usually employ generative methods. Current state-of-the-art solutions usually rely on Generative Adversarial Networks (GANs) [10], [11], [12]. However, little work has been done to improve the perceptual quality of image denoising, even though generative approaches have been in-

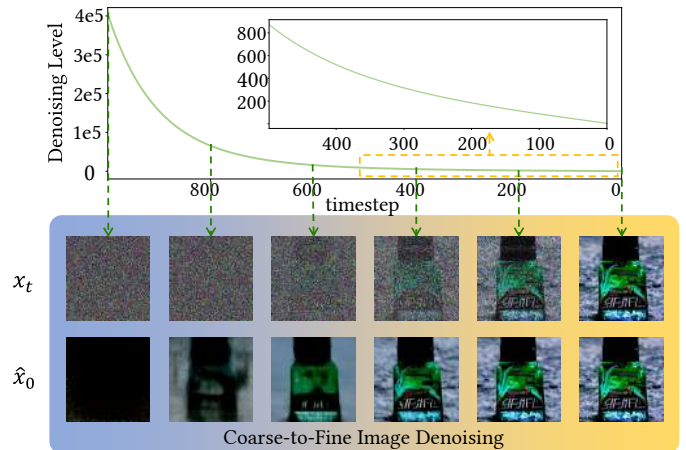


Fig. 1. From the denoising perspective, the reverse process of a diffusion model can be viewed as a coarse-to-fine iterative denoising process. The noise level corresponds to the standard deviation of noise on an 8-bit image, where the maximum signal for a clean image is 255.

troduced to address real noise modeling [13], [14]. Moreover, GAN-based methods are widely blamed for artifacts and inconsistency, leading to significant distortion, especially for extreme degradation.

Recently, the diffusion model [15], [16], [17] has achieved state-of-the-art (SOTA) performance in various tasks, such as image super-resolution [18], [19], image inpainting [20], [21], image deblurring [22]. The methods [18], [20], [21], [22] that utilize diffusion models exhibit higher perceptual quality and fewer artifacts. Therefore, we believe that the diffusion model holds significant potential for image denoising. However, stimulating the diffusion model for image

- Tong Li, Hansen Feng, Lizhi Wang and Lin Zhu are with the School of Computer Science and Technology, Beijing Institute of Technology, Beijing, 100081, China. Email: {litong, fenghansen, wanglizhi, zhulin}@bit.edu.cn
- Zhiwei Xiong is with the Department of Electronic Engineering and Information Science, University of Science and Technology of China, Hefei, 230027, China. Email: zwxiong@ustc.edu.cn
- Hua Huang is with the School of Artificial Intelligence, Beijing Normal University, Beijing, 100875, China. Email: hualiang@bnu.edu.cn
- Corresponding author: Lizhi Wang
- Tong Li and Hansen Feng contributed equally to this work

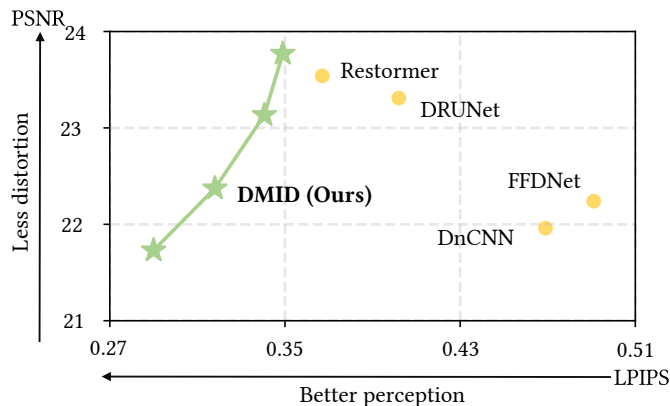


Fig. 2. Perception-distortion trade-off of different methods. Our method traverses through the perception-distortion curve and achieves SOTA performance.

denoising still remains several critical problems. Specifically, the diffusion model is designed to receive standard Gaussian noise as input, which is different from the noisy image input required for image denoising. Furthermore, the content difference between the generated image and the desired denoised image introduces distortion. The input inconsistency problem and the content inconsistency problem are crucial to be tackled to stimulate the diffusion model for image denoising.

In this paper, we propose a novel strategy to stimulate the Diffusion Model for Image Denoising (DMID), consisting of an adaptive embedding method and an adaptive ensembling method. Our insight comes from the denoising perspective towards the diffusion model. From the denoising perspective, the diffusion model and the iterative denoising methods share the same structure and similar functionalities. As illustrated in Figure 1, the diffusion model can be viewed as an iterative denoising method. This new perspective provides us with huge potential to employ a pre-trained unconditional diffusion model to tackle the aforementioned problems, which also eliminates the resource-intensive and time-consuming demands associated with training a diffusion model [23].

Firstly, the problem of input inconsistency is essentially an embedding problem. Here the noisy image and iterative denoising process correspond to the intermediate state and the reverse process subsequence of the pre-trained unconditional diffusion model, respectively. Therefore, we propose an adaptive embedding method that embeds the noisy image into an intermediate state of the pre-trained unconditional diffusion model. Our embedding method enables to perform Gaussian and real-world image denoising.

Secondly, the problem of content inconsistency fundamentally constitutes a stochasticity constraint problem. Since the distortion of the diffusion model mainly comes from the stochasticity inherent in the iterative process. To tackle this problem, we propose an adaptive ensembling method that can adjust distortion and perception without additional training. Our ensembling method enables to traverse through the perception-distortion curve [24] as shown in Figure 2, which achieves SOTA performance on both distortion-based and perception-based metrics.

In summary, our contributions are as follows:

- From the denoising perspective, we contribute a novel understanding and rethinking of the diffusion model and stimulate the diffusion model for image denoising.
- We propose an adaptive embedding method that embeds the noisy image into a pre-trained unconditional diffusion model, which enables us to perform Gaussian and real-world image denoising.
- We propose an adaptive ensembling method that constrains the distortion of the diffusion model, which enables us to traverse through the perception-distortion curve.
- Our method achieves SOTA performance on both distortion-based and perception-based metrics, and our advantages increase as the noise level becomes larger.

The organization of this paper is as follows. Firstly, we provide a brief introduction to the development of image denoising and the diffusion model in Section 2. Next, we offer a rethinking and understanding of the diffusion model from the denoising perspective in Section 3. Then, we propose our DMID strategy to tackle the existing problems in Section 4. After that, we demonstrate the performance and evaluate the effectiveness of our strategy in Section 5. Finally, we discuss the advantages and disadvantages in Section 6 and point out future directions in Section 7.

## 2 RELATED WORKS

In this section, we briefly introduce the development of image denoising and the diffusion model.

### 2.1 Image Denoising

Image denoising is an essential task in computer vision. It aims at restoring a clean image from its noisy counterpart. However, achieving high perception quality and low distortion in image denoising remains a challenging task. To tackle this problem, image denoising methods can be broadly divided into two categories: distortion-based methods and perception-based methods.

Distortion-based methods consist of traditional denoising methods and discriminative methods. Traditional image denoising methods focus on the usage of image priors such as sparsity [1], [25], low rank [2], self-similarity [26], [27], and smoothness [3], [28]. However, the ability of these priors in extreme conditions is ultimately limited, making it difficult to adequately denoise the corrupted images. With the development of deep learning, discriminative models [4], [5], [6], [7], [8], [29], [30], [31], [32], have become the mainstream methods. The main idea of discriminative methods is to use neural networks to learn mappings from noisy images to clean images through paired noisy-clean data or just noisy data [33], [34], [35], [36], [37].  $L_1$  loss and  $L_2$  loss are most widely adopted during training. However, it is well known that these pixel-level losses tend to predict the median (or average) of all possible values rather than the realistic images [9], [10], resulting in over-smooth images with poor perceptual quality.

In addition, some distortion-based methods [1], [38], [39], [40] employ an iterative framework to optimize denoising performance. The iterative denoising methods usu-

ally predict a rough image and refine it through multiple iterations. The iterative denoising framework serves as a regularization and constraint. In addition, the iterative denoising framework decomposes image denoising into a series of sub-problems, making denoising easier. However, these methods still struggle with perceptual quality.

In contrast to distortion-based methods, perception-based methods typically exhibit superior perceptual quality but have significant distortion. Perception-based methods usually refer to generative methods. Current perception-based approaches with competitive performance for other image restoration usually rely on Generative Adversarial Networks (GAN) [10], [11], [12]. However, GAN-based methods often introduce artifacts and inconsistent details that are not present in the original clean image. The artifacts usually lead to significant distortion, especially for extreme degradation. As for image denoising, many methods employ GAN and flow [41] to solve real noise modeling problems [13], [14], [42], [43], [44], [45], [46] to get paired data easily. It is really a shame that little work has attempted to improve the perceptual quality of image denoising.

## 2.2 Diffusion Model

The diffusion model is a burgeoning likelihood-based generative model and has demonstrated remarkable success over other models in various tasks [15], [23], [47].

The diffusion model is originally proposed separately from diffusion-based [16] and score-matching-based [48] perspectives. Later Ho *et al.* [15] demonstrate the enormous potential of unconditional diffusion model for image synthesis. Recently Dhariwal *et al.* propose a kind of conditional diffusion model and bring diffusion models back to the public view. After that, conditional diffusion models adapt to various tasks and have achieved brilliant performances in a series of image restoration tasks, such as image super-resolution [18], [19], image inpainting [20], [21], image deblurring [22]. Some methods also exhibit the capability to address multiple tasks. For instance, the novel method [49] proposes a framework for training a diffusion model on a single image and is applicable in various tasks, including style transfer and harmonization. Another advanced method [21] exhibits proficiency in tasks such as super-resolution, deblurring, and inpainting.

Recently, some methods [21], [50], [51] have exclusively employed pre-trained unconditional diffusion models, originally trained for image synthesis, to address linear image restoration problems. These methods typically assume that the observed image  $y$  is degraded as  $y = Hx + n$ , where  $H$  represents the degradation matrix,  $n$  denotes noise, and  $x$  signifies the desired clean image. The existing methods usually replace some content with the original degraded image and preserve the others during each sampling step of the diffusion model. These methods achieve this by decomposing the degradation matrix  $H$  to identify the boundary of the reserved information area. The popular techniques are Singular Value Decomposition and Range-Null Space Decomposition.

However, it should be emphasized that all the existing methods that employ pre-trained unconditional diffusion models to address restoration problems have been centered around image super-resolution or deblurring tasks,

overlooking the specific task of image denoising. Applying the existing diffusion-based image restoration methods to image denoising still encounters several issues. Firstly, existing methods fall short in the input inconsistency. Existing methods perform restoration by decomposing the degradation matrix  $H$  to identify the boundary of the reserved information area. Therefore these methods circumvent the input inconsistency for other restoration tasks. However, when it comes to image denoising, the degradation matrix  $H$  is simply an identity matrix, which renders matrix decomposition ineffective in providing additional valuable information to identify the boundary of the reserved information area. Therefore such decomposition-based methods fall short in the input inconsistency, making it ill-suited for image denoising. Secondly, existing methods overlook the content inconsistency. These methods often blindly pursue high perceptual quality, leading to a divergence between the denoised image and the desired clean image. Moreover, existing methods neglect the factors that impact distortion-based and perception-based performance, thus failing to meet the high requirements for both distortion-based and perception-based quality in image denoising.

More recently, a method [52] attempts to customize different diffusion models for each noise type. However, to customize distinct diffusion models, the method has to simplify noise types to standard Gaussian or standard Poisson noise, making it less effective for real-world image denoising.

## 3 RETHINKING AND UNDERSTANDING

In this section, we first introduce the pipeline of the diffusion model in Section 3.1. Next, we present the iterative denoising framework in Section 3.2. Finally, we integrate the diffusion model with the iterative denoising framework in Section 3.3, providing a new understanding of the diffusion model from the denoising perspective.

### 3.1 Diffusion Model Pipeline

In this section, we give a rough introduction to the diffusion model.

Diffusion models are composed of a T-timestep forward process and a T-timestep backward process. The forward process gradually adds Gaussian noise to the initial data  $x_0$  and converges to the standard Gaussian distribution. The reverse process starts with standard Gaussian noise  $x_T \sim \mathcal{N}(0, 1)$  and eventually obtains the expected high-quality image  $x_0$ .

The forward process is changed from the previous state  $x_{t-1}$  follows the Markov chain to generate current state  $x_t$ :

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}), \quad (1)$$

where  $t$  is an intermediate timestep,  $x_t$  represents the data (such as a noisy image) of the state at timestep  $t$ ,  $\beta_t$  can be predefined constants as hyperparameters or learned by reparameterization [53], and  $\mathcal{N}$  represents Gaussian distribution. Using the reparameterization technique and following Eq. (1), current state  $x_t$  can also be expressed as:

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (2)$$

---

**Algorithm 1** Iterative Denoising Framework.

---

```

1:  $x_N = y$ 
2: for  $t = N, \dots, 1$  do
3:    $\hat{x}_0 = \text{Denoiser}(x_t)$ 
4:    $x_{t-1} = \hat{x}_0 + \gamma_t(y - \hat{x}_0)$ 
5: end for
6: return  $\hat{x}_0$ 

```

---

where  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{i=1}^t (1 - \beta_i)$ .

The reverse process generates an image by a series of sampling processes, and we refer to the number of sampling times as  $S_t$ . For instance, from the initial state  $x_{1000}$  to get the intermediate state  $x_{500}$  and from the intermediate state  $x_{500}$  to obtain the final state  $x_0$ , the sampling times are  $S_t = 2$ . The sampling process from  $x_t$  to  $x_{t-1}$  is expressed as:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_\theta^2), \quad (3)$$

where  $\mu_\theta(x_t, t)$  is modeled by the neural network. The reverse process gradually transforms the Gaussian distribution  $x_T$  into the expected data distribution  $x_0$ .

Different sampling strategies [15], [54] of Eq. (3) can be summarized as a predicted item and an additional noise item:

$$x_{t-1} = \underbrace{\sqrt{\bar{\alpha}_{t-1}} \left( \frac{x_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(x_t, t)}{\sqrt{\bar{\alpha}_t}} \right)}_{\text{predicted } \hat{x}_0} + \underbrace{\sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \epsilon_\theta(x_t, t) + \sigma_t \epsilon_t}_{\text{additional noise}}, \quad (4)$$

where  $\sigma_t$  is an arbitrary constant and  $\epsilon_\theta(x_t, t)$  is modeled by neural network. In addition, the neural network is trained by optimizing the variational bound on negative data log likelihood  $\mathbb{E}_q[-\log p_\theta(x_0)] \leq \mathcal{L}$  to remove standard Gaussian noise. In each sampling step, the current image  $x_t$  is subtracted from the network-estimated noise  $\epsilon_\theta(x_t, t)$  to obtain the predicted clean image  $\hat{x}_0$ . Subsequently, based on this predicted clean image  $\hat{x}_0$ , some network-estimated noise  $\epsilon_\theta(x_t, t)$  and random Gaussian noise are added to produce the next image  $x_{t-1}$ .

### 3.2 Iterative Denoising Framework

In this section, we point out the iterative framework for image denoising.

Some distortion-based methods [1], [38], [39], [40] employ an iterative framework to optimize denoising performance. These methods usually predict a rough image and refine it through many iterations. Such an iterative framework breaks down the problem into a series of sub-problems, leading to a coarse-to-fine denoising process. The iterative denoising framework can be developed as Algorithm 1.

Specifically, the inference of this framework includes several iterations. In each iteration, a rough denoising result  $\hat{x}_0$  is firstly estimated from the current image  $x_t$ . After that, a weighted original noisy image  $y$  is introduced to obtain a

---

**Algorithm 2** Our Denoising Strategy.

---

```

1:  $x_N = \sqrt{\bar{\alpha}_N} \text{Transform}(y)$ 
2:  $k = \text{floor}(N/S_t)$ 
3: for  $i = 1, \dots, R_t$  do
4:   for  $t$  in reversed(range(0, N, k)) do
5:      $\epsilon \sim \mathcal{N}(0, 1)$ 
6:      $\hat{x}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(x_t, t))$ 
7:      $x_{t-k} = \sqrt{\bar{\alpha}_{t-k}} \hat{x}_0 + \sqrt{1 - \bar{\alpha}_{t-k} - \sigma_t^2} \epsilon_\theta(x_t, t) + \sigma_t \epsilon$ 
8:   end for
9:    $c_i = \hat{x}_0$ 
10: end for
11:  $\hat{x} = \sum_{i=1}^{R_t} c_i p(y|c_i) / \sum_{i=1}^{R_t} p(y|c_i)$ 
12: Inverse Transform  $\hat{x}$ 
13: return  $\hat{x}$ 

```

---

corrected version  $x_{t-1}$  with lower noise than the previous image  $x_t$ , as:

$$x_{t-1} = \hat{x}_0 + \gamma_t(y - \hat{x}_0). \quad (5)$$

This process is repeated iteratively to refine the final result. The iterative denoising framework transforms denoising into an alternating solution process of prior and recovery terms, which is a special case of plug-and-play methods [6], [40].

### 3.3 Discussion

**Denoising Perspective Understanding.** Form follows function [55], which is a well-known law, indicating that similar forms can support similar functions. The high formal similarity between the diffusion model and the iterative denoising framework inspires us to explore their relevance.

The diffusion model and the iterative denoising framework are similar from structures to details. In terms of structures, both techniques build a coarse-to-fine iterative process to generate high-quality images. Each step produces a better iteration result by the weighted fusion of coarse denoised image  $\hat{x}_0$  and noisy image  $y$ . In terms of details, both the sampling process Eq. (4) and iteration Eq. (5) can be decoupled into a prediction item  $\hat{x}_0$  and an additional noise item. This additional noise is mainly a residual noise as  $\epsilon_\theta(x_t, t)$  in Eq. (4) and  $y - \hat{x}_0$  in Eq. (5). Therefore, we declare that the diffusion model and the iterative denoising framework are indeed similar in form, where one sampling step of the former corresponds to one iteration of the latter.

Based on the above observations, we can rewrite the sampling process of the diffusion model from a denoising perspective. Suppose that we start sampling from the intermediate state  $x_N$  at timestep  $N$ . The generated image  $x_0$  after  $N$  steps of reverse process can be obtained by accumulating Eq. (4):

$$x_0 = \text{Denoiser}\left(y, \frac{\sqrt{1 - \bar{\alpha}_N}}{\sqrt{\bar{\alpha}_N}}\right) + \sum_{t=1}^{N-1} \frac{\sigma_{t+1}}{\sqrt{\bar{\alpha}_t}} (\epsilon_{t+1} - \epsilon_\theta(x_t, t)) + \sum_{t=1}^{N-1} \frac{\sqrt{1 - \bar{\alpha}_t - \sigma_{t+1}^2}}{\sqrt{\bar{\alpha}_t}} (\epsilon_\theta(x_{t+1}, t+1) - \epsilon_\theta(x_t, t)), \quad (6)$$

where  $y = \frac{x_N}{\sqrt{\bar{\alpha}_N}} = x + \frac{\sqrt{1 - \bar{\alpha}_N}}{\sqrt{\bar{\alpha}_N}} \epsilon$  and  $\epsilon \sim \mathcal{N}(0, 1)$ . Here  $x$  represents the desired clean image, and  $y$  denotes its

noisy observation. The derivation of Eq. (6) is given in the supplementary material.

The difference between a simple denoiser and Eq. (6) lies in the cumulative terms. Here the cumulative terms represent the additional noise added and removed during the reverse process. During the reverse process, the network-estimated noise  $\epsilon_\theta(x_t, t)$  is subtracted to obtain the predicted clean image  $\hat{x}_0$ , and is added back to produce the next image  $x_{t-1}$ . Theoretically, the cumulative terms should be zero and not affect the level of denoising. In fact, the iterative denoising framework also has a similar cumulative term.

Overall, we introduce a novel denoising perspective for understanding the diffusion model. The diffusion model can be classified as a variant of the iterative denoising framework. The neural network of the diffusion model can be regarded as a denoiser, where timestep  $t$  as a variable corresponds to the preset noise level  $\sqrt{1 - \bar{\alpha}_t}/\sqrt{\bar{\alpha}_t}$ .

Figure 1 illustrates the evolution of the intermediate state  $x_t$ , the estimated image  $\hat{x}_0$ , and the corresponding noise level during the reverse process. In the early stages of the reverse process, the noise level is high and the denoising effect is limited. So only low-frequency information can be determined. In the later stages of the reverse process, the remaining noise is minimal, and then high-frequency details are generated with iterations.

Denoising perspective understanding will intuitively and clearly guide us to stimulate the diffusion model for image denoising. To the best of our knowledge, this is the first time that the diffusion model is comprehended in such a concise denoising perspective.

**Denoising Perspective Rethinking.** The denoising perspective of the diffusion model is the starting point for us to rethink and stimulate the diffusion model for image denoising. The input of diffusion models is standard Gaussian noise while the input of image denoising is a noisy image, which constitutes input inconsistency. The content difference between the generated image and the desired denoised image introduces distortion, which constitutes content inconsistency. Guided by the denoising perspective, the approach to addressing input inconsistency and content inconsistency naturally emerges.

For input inconsistency, the solution lies in the comparison between the input of the diffusion model and the iterative denoising framework. From a denoising perspective, the diffusion model is a conditional AWGN denoiser, iteratively denoising according to a schedule where the expected noise level of the intermediate state  $x_t$  gradually decreases. Therefore, by proposing a method to embed the input into the corresponding intermediate state of the pre-trained diffusion model based on the noise level, we naturally address input inconsistency.

For content inconsistency, the solution lies in the comparison between the process of the diffusion model and the iterative denoising framework. The diffusion model excels in preserving high perceptual quality, while the classical iterative denoising framework primarily aims to reduce distortion. Based on our observations, the difference is mainly attributed to the introduction of random noise  $\epsilon$  besides the original noisy image  $y$  during the sampling process of the diffusion model. The excessive stochasticity of random noise  $\epsilon$  allows the diffusion model to potentially generate content

beyond the original clean image  $x_0$ . From the denoising perspective, this phenomenon corresponds to the local content flickering caused by the stochasticity of i.i.d. noise in low-light video denoising [56]. Taking cues from the late fusion in video processing, we introduce an ensembling method to effectively adjust distortion and perception in the diffusion model, thereby addressing content inconsistency.

It is worth highlighting that our method is insightful compared to existing diffusion-based image denoising methods [50], [51]. Existing methods are not specifically designed for image denoising, thus they typically start from pure Gaussian noise and recommend extremely high sampling times for one inference, which is time-consuming and inefficient. The denoising perspective helps us rethink the diffusion model and addresses the problems of the input inconsistency and the content inconsistency in a concise and efficient manner.

## 4 METHOD

In this section, we first give an overview of our DMID strategy in Section 4.1, which includes an embedding method and an ensembling method. Then we describe the procedure for the embedding method and ensembling method in Section 4.2 and Section 4.3, respectively.

### 4.1 Overview

In this section, we propose a novel strategy to stimulate the diffusion model for image denoising (DMID), which mainly consists of an adaptive embedding method and an adaptive ensembling method. In summary, we first perform the embedding method to construct the initial sampling state. Subsequently, we generate multiple denoised images using the pre-trained diffusion model through multiple inferences. Finally, we employ the ensembling method to reduce distortion.

The embedding method first transforms the input noisy image  $y$  into a latent space where the Gaussian noise assumption is valid. Subsequently, the embedding method normalizes the image range to match the data range expected by diffusion models. Ultimately, the embedding method converts the latent image into the intermediate  $x_N$  state of the diffusion model. The embedding method can be summarised as:

$$x_N = \sqrt{\bar{\alpha}_N} \text{Transform}(y), \quad (7)$$

where  $0 \leq N < T$  is an intermediate timestep. The function  $\text{Transform}(y)$  serves as both a noise transformation technique and a data normalizer. Following the  $\text{Transform}(y)$  function, we multiply the latent noisy image by  $\sqrt{\bar{\alpha}_N}$  to obtain the intermediate  $x_N$  state of the diffusion model. The value of  $N$  should be chosen carefully according to the noise level of the latent noisy image to achieve optimal denoising. After that, we conclude the embedding method.

Furthermore, we then generate a desirable denoised image by directly sampling from timestep  $N$ . The sampling strategy can be any existing approach, such as DDIM [54], DDRM [50], or DDNM [51] among others. In our experiments, we simply employ DDIM. The denoising process,

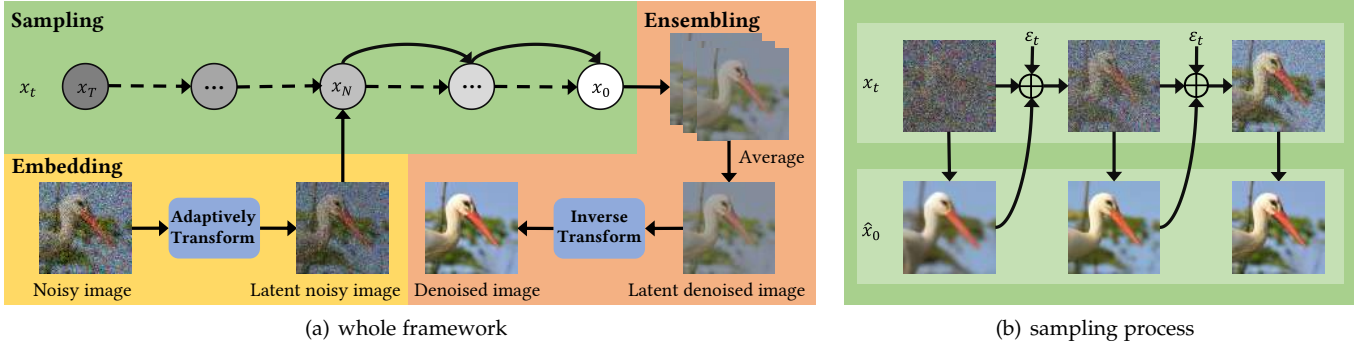


Fig. 3. Our Diffusion Model for Image Denoising (DMID) strategy. Our adaptive embedding method connects diffusion model and image denoising, while our adaptive ensemble method reduces distortion in the final denoised image.

achieved through direct sampling from timestep  $N$ , can be mathematically described as:

$$p(x_0|y) = \prod_{t=1}^N p_\theta(x_{t-1}|x_t), \quad (8)$$

where  $x_0$  is the denoised output image. DMID performs denoising within the latent space and, thus needs to inverse transform the latent denoised image, which is originally within the data range of the diffusion model, back to the original data range.

In addition, we employ our ensembling method to further reduce distortion. Specifically, we adjust the sampling times (referred to  $S_t$ ) in one inference process and repeat the inference process several times (referred to  $R_t$ ). Each inference on the same image corresponds to a Monte Carlo sampling, yielding  $R_t$  denoised images. To mitigate stochasticity and distortion, we conduct a Minimum Mean Square Error (MMSE) averaging of all denoised images, ultimately obtaining the final result.

The whole procedure for the inference of our method is described in Figure 3 and Algorithm 2. We will delve into the procedures of the embedding method and ensembling method in the following sections.

## 4.2 Adaptive Embedding

In this section, we will describe the procedure for the embedding method. As discussed before, the diffusion model is reframed as a Gaussian denoiser. The intermediate states within the diffusion model are adept at handling Gaussian noise with distinct levels. Consequently, our embedding method is tasked with transforming the noise within a noisy image into Gaussian noise and converting the image into a suitable intermediate state based on the noise level. The procedure of the embedding method is shown in Figure 4.

The noise transformation technique involves converting different types of noise into Gaussian noise. The primary goal of our approach is to address general image denoising, allowing it to handle various forms of image noise. Although we have successfully treated the diffusion model as a Gaussian denoiser, it is crucial to account for other noise types, such as real-world noise, which does not adhere to a Gaussian distribution. To surmount this challenge, we improve a noise transformation technique NN [57] to convert diverse noise types into Gaussian noise. In essence, noise transformation entails finding a latent image  $z$  that

exhibits correlation with the input noisy image  $y$ , while the noise in  $z$  conforms to the assumption of additive white Gaussian noise (AWGN). NN [57] achieves this by training an encoder-decoder neural network using a single input noisy image for autoregressive modeling. The latent image  $z$  is derived from the encoder and is expected to adhere to the AWGN assumption. We designate the encoder network of the VAE with parameters  $\theta_1$  as  $G_{\theta_1}$ , the decoder network parameterized by  $\theta_2$  as  $F_{\theta_2}$ , and the latent image as  $z = G_{\theta_1}(y) \sim \mathcal{N}(x, \sigma^2 I)$ . The loss function of the noise transformation technique is presented as:

$$\mathcal{L}(x, \theta_1, \theta_2) = \frac{1}{2} E_\epsilon \|F_{\theta_2}(G_{\theta_1}(y) + \epsilon) - y\|^2 + \frac{1}{2\sigma^2} \|G_{\theta_1}(y) - x\|^2 + \lambda R(x), \quad (9)$$

where  $x$  represents the clean image,  $R(x)$  is a regularization function, and  $\epsilon \sim \mathcal{N}(0, 1)$ . The derivation of Eq. (9) is given in the supplementary material. In the specific approach, three terms in the Eq. (9) work in synergy. The first and second terms ensure data fidelity, while the final term serves as a regularization component. Together, these three terms prevent zero mapping and identity mapping. To facilitate unsupervised training using only the input noisy image  $y$  without the clean image  $x$ , Eq. (9) is optimized through ADMM [58]:

$$\mathcal{L}_\rho(\bar{x}, \theta_1, \theta_2, p, q) = \mathcal{L}(\bar{x}, \theta_1, \theta_2) - \lambda R(\bar{x}) + R(p) + \frac{\rho}{2} \|\bar{x} - p + \frac{q}{\rho}\|^2 - \frac{\rho}{2} \|\frac{q}{\rho}\|^2, \quad (10)$$

$$\bar{x}^{k+1}, \theta_1^{k+1}, \theta_2^{k+1} = \arg \min_{\bar{x}, \theta_1, \theta_2} \mathcal{L}_\rho(\bar{x}, \theta_1, \theta_2, p^k, q^k), \quad (11)$$

$$p^{k+1} = \arg \min_p \mathcal{L}_\rho(\bar{x}^{k+1}, \theta_1^{k+1}, \theta_2^{k+1}, p, q^k), \quad (12)$$

$$q^{k+1} = q^k + \rho(\bar{x}^{k+1} - p^{k+1}), \quad (13)$$

where  $\bar{x}$  is an estimation of the clean image  $x$ ,  $p$  is an auxiliary variable,  $q$  is the dual variable and  $\rho > 0$  is a chosen constant. The subproblems are solved by alternating minimization since the clean image  $x$  is not accessible during unsupervised training. Specifically Eq. (11) is solved by updating the clean image estimation  $\bar{x}$  and the network parameters  $\theta_1$  and  $\theta_2$  alternatively. In addition, Eq. (12) is solved through the plug-and-play idea:

$$p^{k+1} = D(\bar{x}^{k+1} + q^k/\rho), \quad (14)$$




Fig. 4. The procedure for the embedding method. The embedding method first transforms the noise into Gaussian noise and subsequently normalizes the latent noisy image. After that, we search all the timestep  $t$  to find a timestep  $N$  to guarantee  $\frac{\sqrt{1-\bar{\alpha}_N}}{\sqrt{\bar{\alpha}_N}}$  closest to  $\sigma$ . Ultimately, we multiply the noisy image by  $\sqrt{\bar{\alpha}_N}$  to convert the image to the intermediate state  $x_N$ .

where  $D$  is any existing Gaussian denoiser. We choose BM3D [27]. Instead of directly using the latent image  $z = G_{\theta_1}(y)$ , NN employs a linear combination of the noisy image  $y$  and the latent image  $z$  as the final result after transformation.

However, determining when to stop the optimization is a critical challenge. The original NN method required clean images to calculate PSNR and determine the stopping time. We improve this process by using SURE (Stein’s Unbiased Risk Estimator) [59], [60], which eliminates the need for clean images to determine when to stop. SURE provides an unbiased estimate of MSE and represents the quality of the image:

$$SURE(z) = \frac{\|z - D(z)\|^2}{K} - \sigma^2 + \frac{2\sigma^2}{K} \sum_{i=1}^K \frac{\partial D_i(z)}{\partial z_i}, \quad (15)$$

where  $K$  is the image size and  $z_i$  is the  $i$ th element of  $z$ . Additionally, we employ the Monte-Carlo (MC) approximation [61] of the divergence term in Eq. (15) as follows:

$$\sum_{i=1}^K \frac{\partial D_i(z)}{\partial z_i} \approx \frac{1}{\mu} \epsilon^T (D(z + \mu\epsilon) - D(z)), \quad (16)$$

where  $T$  is the transpose operator and  $\mu$  is a small positive value. Every 500 iterations, we calculate the SURE of the latent image  $z$ . Smaller SURE values indicate higher image quality, and we stop the iterations when SURE begins to increase. It merits mentioning that our improved transformation technique requires no pre-training on external datasets or the use of the clean image  $x$ . The noise transformation process solely relies on the single input noisy image  $y$ .

The value of  $N$  should be chosen carefully to achieve optimal denoising performance based on the noise level of the latent noisy image. For the latent noisy image with noise level  $\sigma$ , we search all the timestep  $t$  to find a timestep  $N$  to guarantee  $\frac{\sqrt{1-\bar{\alpha}_N}}{\sqrt{\bar{\alpha}_N}}$  closest to  $\sigma$ . After that, we multiply the noisy image by  $\sqrt{\bar{\alpha}_N}$  to convert the image to the intermediate state  $x_N$ . The reason is that we should align the denoising ability of the diffusion model at timestep  $N$  with the noise level  $\sigma$  of the latent noisy image to attain optimal denoising outcomes. And we consider the diffusion model as a denoiser, whose denoising ability at timestep  $t$  is  $\sqrt{1-\bar{\alpha}_t}/\sqrt{\bar{\alpha}_t}$ , as discussed in Section 3.3. In practice, we calculate the corresponding level  $\frac{\sqrt{1-\bar{\alpha}_t}}{\sqrt{\bar{\alpha}_t}}$  for each timestep  $t$  in advance. For a latent noisy image with noise level  $\sigma$ , we readily identify an  $N$  that corresponds to the level  $\frac{\sqrt{1-\bar{\alpha}_N}}{\sqrt{\bar{\alpha}_N}}$  closest to  $\sigma$ . For instance, we choose  $N = 115$  for noise level  $\sigma = \frac{2*50}{255} \approx 0.3922$ , as  $\sqrt{1-\bar{\alpha}_t}/\sqrt{\bar{\alpha}_t}$  closest to  $\sigma$  at  $t = 115$  among all timestep  $t$ , as shown in Figure 4. For tasks such as Gaussian denoising, the noise level of the

latent noisy image is known. In cases of unknown noise levels, established noise estimation techniques [62], [63] can be applied to address the issue.

### 4.3 Adaptive Ensembling

In this section, we will describe the procedure for the ensembling method. Our embedding method is tasked with adjusting distortion and perception based on specific requirements. The primary factor influencing distortion and perception is the stochasticity brought by the Gaussian noise in the additional noise item. In practice, excessive stochasticity can lead to significant distortion in the final denoised image. In theory, moderate stochasticity can facilitate convergence to better results [48]. Thus our ensembling strategy involves adjusting the sampling times to control stochasticity in one inference process and averaging the images obtained from multiple inferences to converge to higher quality.

Firstly we constrain the stochasticity in one inference. The level of stochasticity is primarily determined by the number of sampling times since the additional noise item appears in each sampling process. A higher number of sampling times  $S_t$  allows for more refinement and introduces more stochasticity, resulting in a more detailed image with improved perceptual quality. However, this also leads to a smaller weighted noisy image  $y$  at each iteration, which reduces its controllability at the same time. Therefore, larger sampling times generally yield better perceptual quality but also result in larger distortion. To adjust distortion and perception, we recommend adjusting the number of sampling times according to the desired outcome. To reduce distortion, setting sampling times to be less than 10 is typically sufficient.

Furthermore, we enhance the denoised images through the stochasticity in multiple inferences. The Gaussian noise in the additional noise item introduces stochasticity into the sampling process. This enables generating multiple clean images for the same noisy image when the sampling times are greater than one. Thus, we propose an ensembling approach based on the Monte Carlo method to utilize the stochasticity of multiple inferences. Specifically, we repeat the inference process multiple times to generate multiple candidate denoised images for the same noisy input. We refer to the repetition times of inference as  $R_t$ , and the candidate denoised image as  $c_i$ . Additionally, we incorporate the Minimum Mean Square Error (MMSE) averaging [73] to further enhance the ensembling method. The final denoised image is expressed as

$$x^{\text{MMSE}} = \frac{\sum_{i=1}^{R_t} c_i p(y|c_i)}{\sum_{i=1}^{R_t} p(y|c_i)}, \quad (17)$$

TABLE 1

Classical Gaussian image denoising. The top row is training a separate model for specific noise level, the bottom row is the models that designed to deal with various noise levels. The best and second-best methods are in red and blue.

Method	CBSD68 [64]			Kodak24 [65]			McMaster [66]			Urban100 [67]		
	$\sigma=15$	$\sigma=25$	$\sigma=50$	$\sigma=15$	$\sigma=25$	$\sigma=50$	$\sigma=15$	$\sigma=25$	$\sigma=50$	$\sigma=15$	$\sigma=25$	$\sigma=50$
BRDNet [68]	34.10	31.43	28.16	34.88	32.41	29.22	35.08	32.75	29.52	34.42	31.99	28.56
RNAN [69]	-	-	28.27	-	-	29.58	-	-	29.72	-	-	29.08
RDN [70]	-	-	28.31	-	-	29.66	-	-	-	-	-	29.38
IPT [30]	-	-	28.39	-	-	29.64	-	-	29.98	-	-	29.71
SwinIR [4]	<b>34.42</b>	31.78	28.56	35.34	32.89	29.79	<b>35.61</b>	33.20	30.22	<b>35.13</b>	32.90	29.82
Restormer [5]	34.40	<b>31.79</b>	<b>28.60</b>	<b>35.47</b>	<b>33.04</b>	<b>30.01</b>	<b>35.61</b>	<b>33.34</b>	<b>30.30</b>	<b>35.13</b>	<b>32.96</b>	<b>30.02</b>
CODE [71]	34.33	31.69	28.47	35.32	32.88	29.82	35.38	33.11	30.03	-	-	-
CBM3D [27]	33.52	30.71	27.38	34.28	32.15	28.46	34.06	31.66	28.51	32.35	29.70	25.95
DnCNN [7]	33.90	31.24	27.95	34.60	32.14	28.95	33.45	31.52	28.62	32.98	30.81	27.59
FFDNet [8]	33.87	31.21	27.96	34.63	32.13	28.98	34.66	32.35	29.18	33.83	31.40	28.05
DSNet [72]	33.91	31.28	28.05	34.63	32.16	29.05	34.67	32.40	29.28	-	-	-
DRUNet [6]	34.30	31.69	28.51	35.31	32.89	29.86	35.40	33.14	30.08	34.81	32.60	29.61
Restormer [5]	34.39	31.78	28.59	35.44	33.02	30.00	35.55	33.31	30.29	35.06	32.91	<b>30.02</b>
<b>DMID-d (Ours)</b>	<b>34.45</b>	<b>31.86</b>	<b>28.72</b>	<b>35.51</b>	<b>33.12</b>	<b>30.14</b>	<b>35.72</b>	<b>33.49</b>	<b>30.50</b>	<b>35.26</b>	<b>33.11</b>	<b>30.28</b>

where  $p(y|c_i) \sim \mathcal{N}(c_i, \sigma^2\mathbf{I})$  is calculated by the Gaussian noise model,  $\sigma$  is the noise level of the noisy image, and  $\mathcal{N}$  represents the Gaussian distribution. The derivation of Eq. (17) is given in the supplementary material. Inspired by our denoising prespective understanding, we reframe the diffusion model as a Gaussian denoiser and transform all noise to Gaussian noise. Thus we can leverage existing noisy images  $y$  and the Gaussian noise model  $p(y|x)$  to compute the observation likelihood  $p(y|c_i)$  for each candidate restored image  $c_i$ . This likelihood indicates how well the candidate restored image  $c_i$  aligns with the denoised outcome of the noisy image  $y$ . Therefore, by weighting our predicted denoised results with the corresponding observation likelihood, we can further minimize stochasticity. According to the Law of Large Numbers, a higher number of repetition times  $R_t$  yield better distortion-based quality.

It is imperative to highlight that our contribution is pioneering in multiple facets. Firstly, we reframe the diffusion model as a Gaussian denoiser. This allows us to denoise in a single step or potentially optimize through multiple iterations. Secondly, we elucidate the embedding method and the noise-correlated sampling starting point for image denoising. This not only reduces the required sampling times in one inference but also enhances the quality of the results. Finally, we present the elucidation of how sampling times  $S_t$  and repetition times  $R_t$  impact distortion-based and perception-based quality. This helps reduce distortion and lays the foundation for further research in image denoising area.

## 5 EXPERIMENTS

Our experiments are generally divided into three parts. First of all, we demonstrate the performance of our method on Gaussian noise in Section 5.2 and real-world noise in Section 5.3. In addition, we conduct detailed ablation studies to fully evaluate our embedding and ensembling method in Section 5.4. Finally, we engage in a comparative analysis and extension to other diffusion-based methods in Section 5.5.

### 5.1 Implementation Details

We employ a pre-trained model from [23] which is trained on  $256 \times 256$  images from ImageNet [74], with full timesteps  $T = 1000$ . Following [23],  $\beta_t$  is defined increasing linearly from 0.0001 to 0.02 for  $t$  from 1 to 1000,  $\sigma_t$  is set to be  $\gamma \sqrt{\frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}} \sqrt{1 - \frac{\bar{\alpha}_t}{\bar{\alpha}_{t-1}}}$ , and  $\gamma = 0.85$ .

Since our method can traverse through the perception-distortion curve [24], [75], we present two variants of our method. The first variant, named “DMID-d”, achieves the least distortion and satisfactory perceptual quality. The second variant, named “DMID-p”, achieves the best perceptual quality and tolerable distortion. To achieve this, we need to determine the number of sampling times (referred to  $S_t$ ) in an inference process and how many times will we repeat and get an average (referred to  $R_t$ ). For the first variant, we set the full sampling times  $S_t * R_t$  to be 1000 following [15], [20] to ensure fairness. For the second variant, we set  $R_t = 1$  and the value of  $S_t$  varies from 2 to 200 for different datasets and various noise levels. As described in Section 4.3, decreasing the number of sampling times  $S_t$  and increasing the number of repetition times  $R_t$  will reduce distortion while conversely improving perceptual quality. This is why we introduced the two variations. Further explanations can be found in Section 5.4.

### 5.2 Gaussian Image Denoising

In this section, we conduct Gaussian image denoising experiments on synthetic benchmark datasets.

Image denoising is commonly evaluated using the distortion-based PSNR metric with noise levels  $\sigma = 15, 25, 50$ . However, our method is robust to much higher noise levels, excelling in both distortion-based and perception-based metrics. Hence, we perform classical Gaussian denoising experiments as well as robust Gaussian denoising experiments.

For classical Gaussian denoising, we evaluate our method using the distortion-based PSNR metric with noise levels  $\sigma = 15, 25, 50$ , consistent with previous classical





Fig. 5. Visual results on classical Gaussian image denoising. The images restored by our model exhibit more details and realism.

comparisons [5], [71], [76]. Table 1 presents the PSNR scores achieved by various SOTA approaches on the synthetic benchmark datasets (CBSD68 [64], Kodak24 [65], McMaster [66], Urban100 [67]). All the results are reported from [5], [76]. The top row in Table 1 is training a separate model for a specific noise level, the bottom row is the models that are designed to deal with various noise levels. The methods (such as Restormer [5] and SwinIR [4]) that can only handle one noise level still perform worse than us. As shown in Figure 5, these two images presented significant challenges for previous methods, with their denoised results showing highly unrealistic facial features. In contrast, our method successfully denoises these images to a natural and realistic quality, highlighting the effectiveness of our approach.

For the robust Gaussian denoising, we evaluate our method using both distortion-based and perception-based metrics with larger noise levels. We compare our method with DnCNN [7], DRUNet [6], Restormer [5] and ART [76]. Specifically, we evaluate different methods on ImageNet 1K (crop the center  $256 \times 256$  image as input) [74], CBSD68 [64], Kodak24 [65], and McMaster [66] datasets. In addition, we evaluate different methods on two representative distortion-based metrics (PSNR and SSIM [77]) and perception-based metric (LPIPS [78]). Remarkably, our method is effective even when the standard deviation greatly exceeds 255, a previously unexplored capability that sets our method apart from other models. To ensure a fair comparison with other models, we evaluate different methods across a series of representative noise levels, spanning from 0 to 255.

Subsequently, we will provide a detailed explanation of our comparative methods. DnCNN [7] and DRUNet [6] are considered classical methods, while Restormer [5] and ART [76] represent the current state-of-the-art approaches. For DnCNN [7] and DRUNet [6], we retrain them following their specified training details. However, it is worth noting

that retraining Restormer [5] and ART [76] is a resource-intensive process. Their publicly available pre-trained models are not designed to handle noise levels with standard deviations exceeding 50. Therefore, following the approach of [6], we multiply by a constant to ensure that the standard deviation is 50.

As illustrated in Table 2, our method demonstrates robustness across various noise levels and achieves SOTA performance on all metrics. Under extreme conditions of the ImageNet [74] dataset, we outperform Restormer by more than 0.5dB in PSNR. Furthermore, we achieve substantial improvements in perception-based metrics in all datasets. This notable enhancement directly translates into improved perceptual quality, as vividly depicted in Figure 6. Whether dealing with irregular animal fur, intricate grass patterns, or structured circular designs, our method consistently exhibits superior performance when compared to alternative approaches. This experiment substantiates the robustness of our method across a spectrum of noise levels and evaluation metrics.

### 5.3 Real-world Image Denoising

In this section, we conduct real-world image denoising experiments on real-world benchmark datasets.

We compare our method with both supervised and unsupervised methods. As our method is essentially an unsupervised solution for real-world denoising, which does not require training on pairs of noisy-clean real-world images. Specifically, we evaluate different methods on three datasets: CC [79], PolyU [80], and FMDD [81], following [37], [57], [87], [88].

Subsequently, we will provide a detailed explanation of the comparative methods. For supervised methods, DANet+ [82] is the only and the latest generative method that can be directly employed for denoising. In addition,

TABLE 2

Robust Gaussian image denoising. We report the results of method labeled "Ours-d" with least distortion, and a second method labeled "Ours-p" with greatest perceptual quality. Our method achieves SOTA performance on all metrics (PSNR↑ / SSIM↑ / LPIPS↓) and on all noise levels.

Dataset	Noise Level	DnCNN [7]	DRUNet [6]	Restormer [5]	ART [76]	DMID-d (Ours)	DMID-p (Ours)
		PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS
ImageNet	$\sigma = 50$	28.21 / 0.8806 / 0.179	29.46 / 0.9081 / 0.145	29.61 / 0.9110 / 0.136	29.62 / 0.9105 / 0.132	29.90 / 0.9157 / 0.114	27.59 / 0.8722 / 0.087
	$\sigma = 100$	25.34 / 0.8025 / 0.293	26.48 / 0.8482 / 0.252	26.61 / 0.8532 / 0.234	26.57 / 0.8501 / 0.231	27.00 / 0.8626 / 0.201	24.61 / 0.7987 / 0.156
	$\sigma = 150$	23.69 / 0.7421 / 0.370	24.81 / 0.8032 / 0.331	24.93 / 0.8103 / 0.306	24.88 / 0.8043 / 0.307	25.39 / 0.8236 / 0.263	22.94 / 0.7465 / 0.210
	$\sigma = 200$	22.52 / 0.6910 / 0.432	23.66 / 0.7673 / 0.393	23.78 / 0.7762 / 0.363	23.72 / 0.7674 / 0.368	24.25 / 0.7915 / 0.295	21.56 / 0.6932 / 0.259
	$\sigma = 250$	21.62 / 0.6448 / 0.478	22.80 / 0.7374 / 0.445	22.92 / 0.7479 / 0.409	22.84 / 0.7363 / 0.417	23.44 / 0.7675 / 0.346	20.87 / 0.6701 / 0.289
CBSD68	$\sigma = 50$	27.84 / 0.8844 / 0.226	28.51 / 0.8991 / 0.183	28.59 / 0.9011 / 0.177	28.63 / 0.9015 / 0.173	28.69 / 0.9029 / 0.162	26.63 / 0.8605 / 0.122
	$\sigma = 100$	25.06 / 0.8090 / 0.365	25.76 / 0.8357 / 0.308	25.84 / 0.8389 / 0.291	25.86 / 0.8376 / 0.295	25.96 / 0.8413 / 0.283	23.94 / 0.7840 / 0.208
	$\sigma = 150$	23.58 / 0.7555 / 0.449	24.32 / 0.7932 / 0.395	24.41 / 0.7980 / 0.367	24.41 / 0.7945 / 0.380	24.54 / 0.8001 / 0.361	22.47 / 0.7314 / 0.264
	$\sigma = 200$	22.56 / 0.7131 / 0.514	23.36 / 0.7616 / 0.464	23.47 / 0.7682 / 0.426	23.44 / 0.7623 / 0.447	23.57 / 0.7686 / 0.392	21.37 / 0.6842 / 0.312
	$\sigma = 250$	21.80 / 0.6773 / 0.558	22.64 / 0.7368 / 0.519	22.77 / 0.7451 / 0.475	22.72 / 0.7364 / 0.500	22.88 / 0.7462 / 0.455	20.77 / 0.6610 / 0.352
Kodak24	$\sigma = 50$	28.84 / 0.8921 / 0.247	29.86 / 0.9132 / 0.188	30.00 / 0.9153 / 0.185	30.02 / 0.9152 / 0.181	30.13 / 0.9174 / 0.172	27.90 / 0.8770 / 0.131
	$\sigma = 100$	26.00 / 0.8206 / 0.387	27.16 / 0.8609 / 0.297	27.30 / 0.8642 / 0.287	27.27 / 0.8616 / 0.289	27.50 / 0.8682 / 0.274	25.31 / 0.8107 / 0.211
	$\sigma = 150$	24.41 / 0.7670 / 0.475	25.69 / 0.8233 / 0.381	25.85 / 0.8281 / 0.363	25.78 / 0.8219 / 0.372	26.08 / 0.8336 / 0.348	24.06 / 0.7707 / 0.275
	$\sigma = 200$	23.33 / 0.7228 / 0.543	24.65 / 0.7930 / 0.450	24.84 / 0.7999 / 0.422	24.74 / 0.7901 / 0.439	25.08 / 0.8054 / 0.382	22.99 / 0.7311 / 0.318
	$\sigma = 250$	22.51 / 0.6830 / 0.589	23.89 / 0.7686 / 0.504	24.09 / 0.7772 / 0.469	23.94 / 0.7631 / 0.491	24.40 / 0.7853 / 0.446	22.44 / 0.7133 / 0.356
McMaster	$\sigma = 50$	28.35 / 0.9078 / 0.180	30.04 / 0.9350 / 0.140	30.29 / 0.9378 / 0.134	30.31 / 0.9378 / 0.132	30.51 / 0.9396 / 0.124	28.34 / 0.9112 / 0.092
	$\sigma = 100$	25.59 / 0.8499 / 0.283	27.05 / 0.8914 / 0.233	27.25 / 0.8962 / 0.220	27.22 / 0.8945 / 0.218	27.57 / 0.8990 / 0.206	25.37 / 0.8560 / 0.158
	$\sigma = 150$	23.98 / 0.8041 / 0.360	25.36 / 0.8570 / 0.300	25.57 / 0.8645 / 0.279	25.51 / 0.8604 / 0.281	25.89 / 0.8672 / 0.267	23.72 / 0.8148 / 0.209
	$\sigma = 200$	22.85 / 0.7646 / 0.423	24.20 / 0.8279 / 0.354	24.41 / 0.8384 / 0.327	24.33 / 0.8321 / 0.334	24.68 / 0.8392 / 0.303	22.51 / 0.7816 / 0.252
	$\sigma = 250$	21.96 / 0.7279 / 0.469	23.31 / 0.8031 / 0.402	23.54 / 0.8161 / 0.367	23.43 / 0.8077 / 0.380	23.81 / 0.8174 / 0.359	21.73 / 0.7553 / 0.290

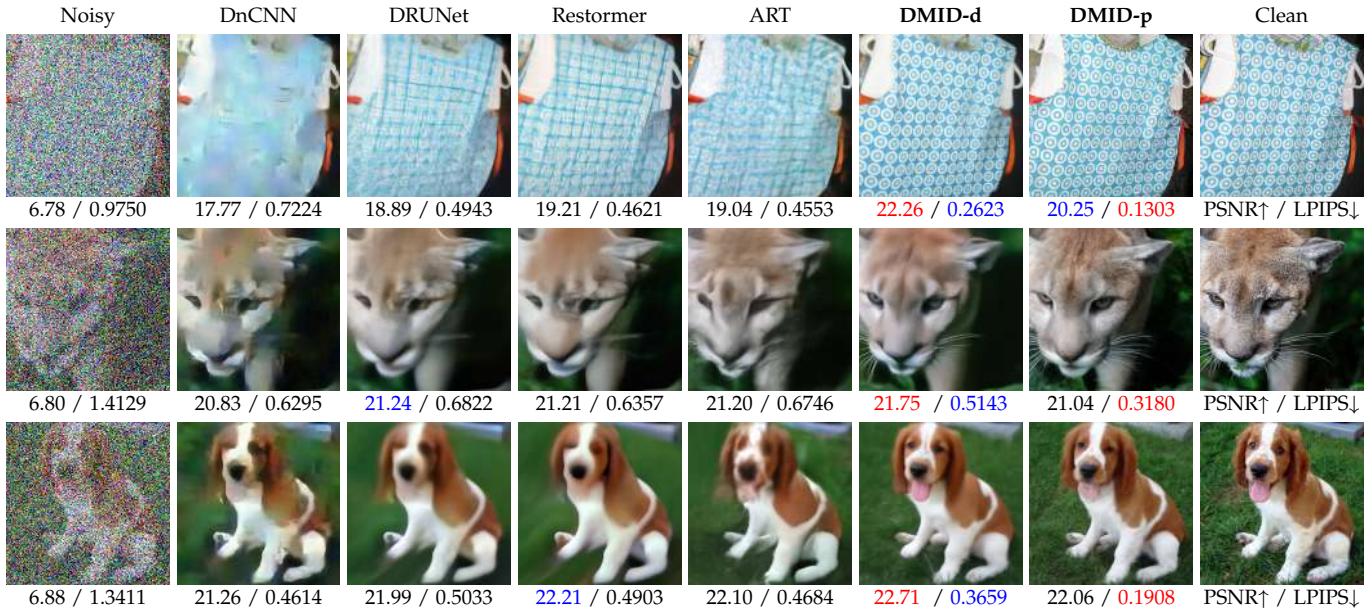


Fig. 6. Visual results on robust Gaussian image denoising. Our method can generate detailed texture, while other models even have severe chromatic aberration and blur.

all the supervised methods are pre-trained on real-world dataset SIDD [89] and borrowed from their officially released versions. For unsupervised methods, the results of AP-BSN [35], LG-BPN [36], ZS-N2N [87], R2R [37] are reproduced and evaluated by ourselves, since these methods are not evaluated on all the datasets and metrics we use in their original paper. The results of N2V [33], N2S [34] and S2S [86] are reported from R2R [37] and ScoreDVI [88].

Table 3 and Figure 7 demonstrate our significant advantages. The Table 3 segregates supervised methods in the top row and unsupervised methods in the bottom row. In

contrast to supervised methods, our unsupervised approach yields significantly higher PSNR and SSIM scores. Supervised methods, trained on the SIDD dataset [89], suffer from poor generalization, rendering them highly vulnerable to out-of-distribution data. As a result, their outcomes consistently exhibit significant residual noise. In contrast to unsupervised methods, our results demonstrate a substantial improvement. AP-BSN [35] and LG-BPN [36] often yield overly blurred outputs, accompanied by texture distortion, whereas our results maintain remarkably clear textures. R2R [37] and ZS-N2N [87] appear to grapple with complete

TABLE 3

Real-world image denoising. Our method achieves excellent performance across all metrics (PSNR $\uparrow$  / SSIM $\uparrow$  / LPIPS $\downarrow$ ) when compared to both supervised (the top row) and unsupervised methods (the bottom row). The best and second-best methods are in red and blue.

Method	CC [79]			PolyU [80]			FMDD [81]		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
DANet $_+$ [82]	35.91	0.9816	0.073	37.23	0.9796	0.088	31.59	0.7962	0.238
MIRNet [83]	36.04	0.9797	0.088	37.45	0.9783	0.101	31.83	0.8106	0.227
MPRNet [32]	36.20	0.9769	0.094	37.43	0.9759	0.109	31.21	0.7915	0.292
DeamNet [84]	35.64	0.9652	0.089	32.46	0.8615	0.130	32.28	0.8116	0.261
NAFNet [85]	34.39	0.9784	0.073	36.04	0.9607	0.107	26.27	0.6519	0.340
Uformer [29]	36.31	0.9795	0.085	37.31	0.9782	0.096	31.81	0.8118	0.207
Restormer [5]	36.27	0.9810	0.077	37.51	0.9776	0.102	31.81	0.8046	0.211
N2V [33]	32.27	0.862-	-	33.83	0.873-	-	-	-	-
N2S [34]	33.38	0.846-	-	35.04	0.902-	-	-	-	-
S2S [86]	37.52	0.951-	-	38.37	0.962-	-	30.76	0.695-	-
AP-BSN [35]	34.86	0.9744	0.131	36.45	0.9750	0.099	32.40	0.8461	0.335
R2R [37]	33.43	0.9564	0.227	36.23	0.9655	0.151	27.17	0.5250	0.448
LG-BPN [36]	34.58	0.9755	0.135	36.59	0.9763	0.102	33.12	0.8668	0.283
ZS-N2N [87]	33.51	0.9571	0.224	35.99	0.9587	0.197	31.65	0.7674	0.222
ScoreDVI [88]	37.09	0.945-	-	37.77	0.959-	-	33.10	0.865-	-
<b>DMID-d (Ours)</b>	<b>37.99</b>	<b>0.9880</b>	0.078	<b>38.62</b>	<b>0.9853</b>	<b>0.069</b>	<b>33.40</b>	<b>0.8747</b>	0.266
<b>DMID-p (Ours)</b>	37.09	<b>0.9854</b>	<b>0.072</b>	<b>38.46</b>	<b>0.9843</b>	<b>0.067</b>	33.09	0.8616	0.232

TABLE 4

Results of the ablation study on the Embedding method with whether perform noise transformation.

Noise Transformation	CC			FMDD		
	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS
✗	35.97 / 0.9769 / 0.095	29.47 / 0.6806 / 0.334				
✓	37.09 / 0.9854 / 0.072	31.97 / 0.7830 / 0.314				

noise removal, leaving residual noise in the final results. These experiments unequivocally underscore the adeptness of our method in handling denoising tasks, achieving high perceptual quality while introducing minimal distortion.

The abnormal situation observed in the FMDD [81] dataset concerning the LPIPS metric is worth discussing. It is notable that our method and DANet $_+$  [82], both generative methods, perform well on LPIPS for CC [79] and PolyU [80] but fail for FMDD [81]. Additionally, our method and LG-BPN [36] perform well on PSNR and SSIM for FMDD [81] but not on LPIPS. This situation appears abnormal. We hypothesize that it may be due to the perception-distortion curve phenomenon [24]. The perception-distortion curve implies a contradiction between distortion-based metrics and perception-based metrics. Therefore, outstanding performance on distortion-based metrics can not align with outstanding perception-based metrics. This explanation is consistent with our results on CC [79], where our DMID-d does not achieve the second-best results on LPIPS. Despite the presence of such an abnormal situation, our method still stands out among unsupervised methods. Disregarding this abnormal situation, our method shines brightly even when compared to various supervised methods.

### 5.4 Ablation Study

In this section, we conduct ablation studies to evaluate our embedding method and ensembling method.

#### 5.4.1 For embedding method

Our embedding method first performs noise transformation and then converts the image to the intermediate state  $x_N$ . Therefore, for the embedding method, we first assess the impact of noise transformation. Then, we evaluate the impact of the choice of  $N$ , which is crucial for our method, as denoising cannot be performed without it for our DMID method.

(a). **Influence of noise transformation.** We conduct experiments on CC [79] and FMDD [81] with a fixed number of sampling times  $S_t = 1$  and repetition times  $R_t = 1$ .

Without converting the image into the intermediate state  $x_N$  of the diffusion model, our method cannot function. Therefore, we consistently convert the image into the intermediate state  $x_N$ , whether performing noise transformation or not. The results are presented in Table 4, and it is evident that noise transformation brings about significant improvement. This experiment demonstrates the importance of noise transformation for our method.

(b). **Influence of  $N$ .** We conduct experiments on McMaster with noise level  $\sigma = 25$  using different choices of  $N$  with a fixed number of sampling times  $S_t = 1$  and repetition times  $R_t = 1$ .

As explained in Section 4.2, the most optimal denoising outcomes are realized when the denoising ability of the diffusion model at timestep  $N$  matches the noise level of the noisy image  $x_N$ . In simpler terms, we define the ratio of the denoising ability and the noise level as  $\eta$ , and when  $\eta \approx 1$  at timestep  $N$ , here we can get optimal denoising outcomes and the timestep  $N$  corresponds to such a noise level. This concept is intuitively straightforward, and the denoising outcomes fluctuate based on alterations in the correlation with  $\eta$  as illustrated in Figure 8. The distinction between appropriate and inappropriate values of timestep  $N$  can result in a significant difference in PSNR and LPIPS.

Since LPIPS is more tolerant to noise than to blurring [78], the highest LPIPS score is frequently achieved when

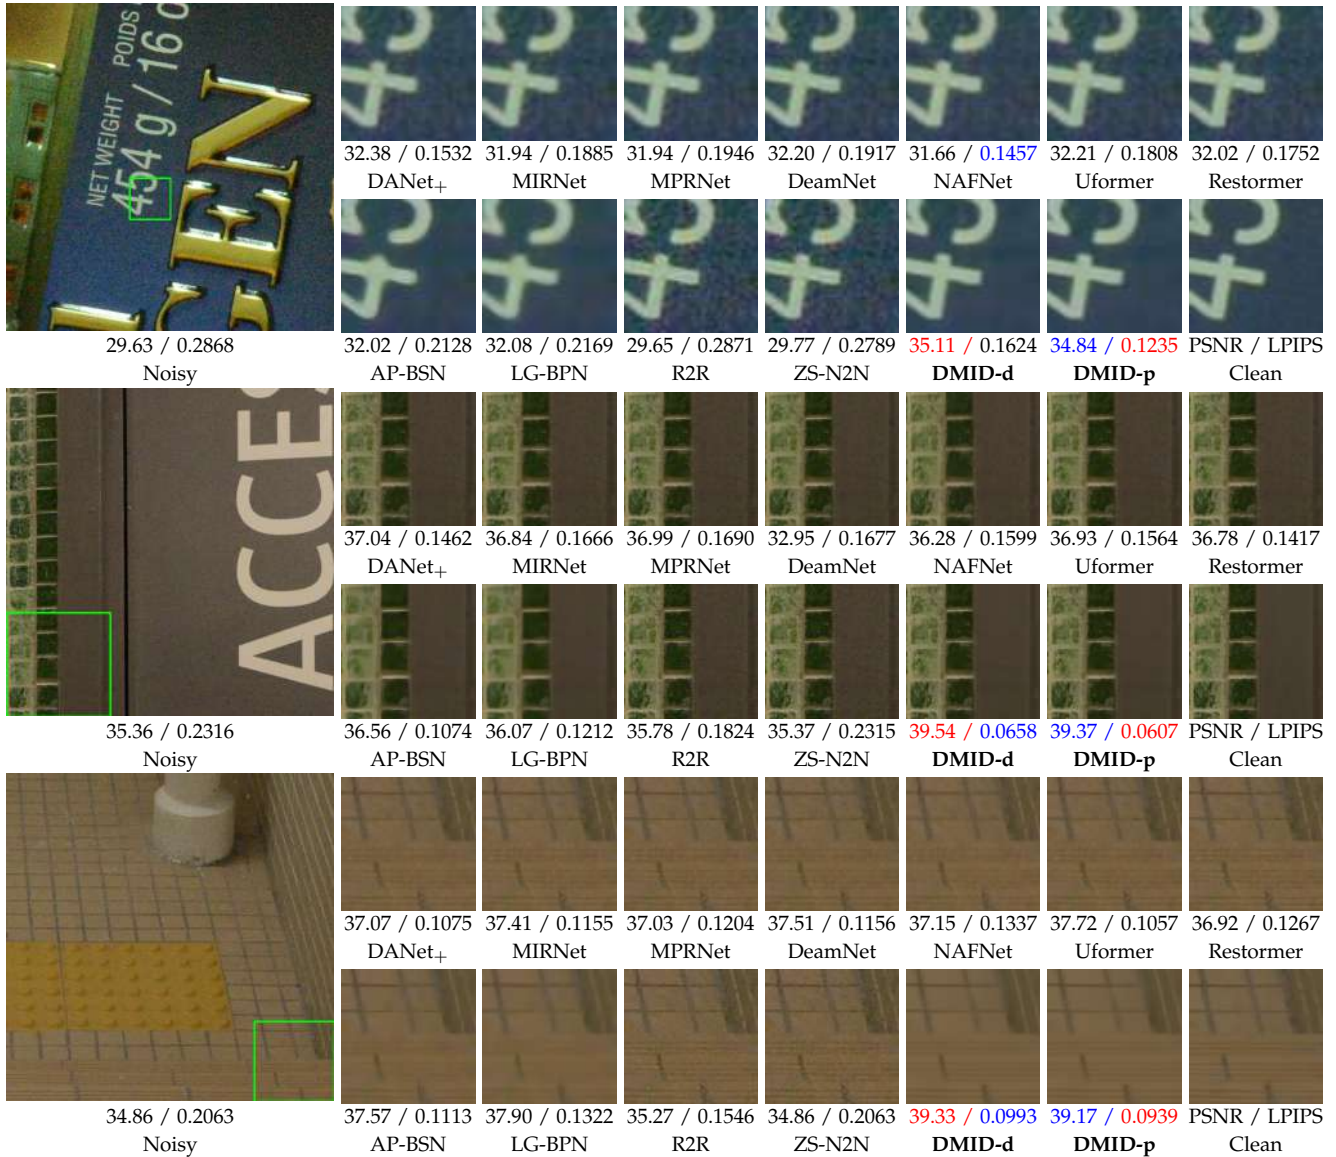


Fig. 7. Visual results on real-world image denoising. Our method achieves better denoising results and pictures restored by our method are more detailed and realistic.

denoising is slightly incomplete. When  $\eta \approx 1$  at timesetp  $N$ , this optimal timesetp  $N$  is denoted as  $N_o$ . As for our method, the optimal LPIPS is generally obtained within the range of  $N_o - 1$  to  $N_o - 5$  as shown in Figure 9. In our experiments, we just set  $N$  to be  $N_o$ , and further tuning could potentially yield even better LPIPS results. This phenomenon, in which the optimal LPIPS score is not attained at  $N_o$  while PSNR is achieved at  $N_o$ , further validates the credibility of our approach. It underscores the accuracy of our  $N$  calculation. Furthermore, our method demonstrates robustness, showing favorable denoising outcomes when  $\eta$  varies within the range of  $1 \pm 10\%$  as shown in Figure 9.

This experiment demonstrates the necessity of embedding into intermediate states and the appropriate choice of  $N$  for our method. Without embedding into intermediate states, our method cannot perform denoising. With the appropriate choice of  $N$ , our denoising performance reaches its maximum potential.

#### 5.4.2 For ensembling method

Our ensembling method adjusts distortion and perception based on requirements by adjusting the sampling times and repeating the inference process. Therefore, for the ensembling method, we separately evaluate the impact of different values for sampling times  $S_t$  and repetition times  $R_t$ .

(a). **Influence of  $S_t$ .** We conduct experiments on McMaster with noise level  $\sigma = 250$  using different sampling times  $S_t$  with a fixed number of repetition times  $R_t = 1$  for Gaussian denoising. In addition, we conduct experiments on CC [79], PolyU [80], and FMDD [81] for real-world denoising.

Increasing sample timesteps  $S_t$  can lead to more corruption and reconstruction times, and weaken the controllability of the noisy image  $y$  introduced in each iteration. While this can improve perceptual quality, it may also increase the uncontrollability of picture details resulting in less similarity in PSNR. The experimental results are as shown in Figure 10

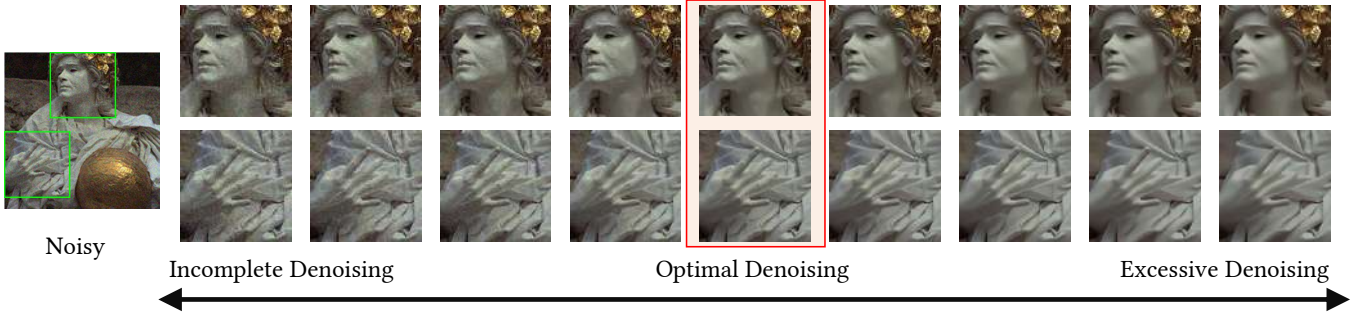


Fig. 8. Visual results of ablation study on the embedding method. The most optimal denoising outcomes are realized when the denoising ability of the diffusion model at timestep  $N$  matches the noise level of the noisy image  $x_N$ .




Fig. 9. Results of the ablation study on the embedding method with varying timestep  $N$ . The most optimal denoising outcomes are realized when  $\eta \approx 1$ .




Fig. 10. Results of the ablation study on the ensembling method with varying sampling times  $S_t$ . Increased sampling times  $S_t$  result in higher distortion, particularly when  $S_t > 1$  for Gaussian denoising.

for Gaussian denoising and Table 5 for real-world denoising.

The overall impact trend of sample timesteps  $S_t$  is similar for both Gaussian denoising and real-world denoising. However, there are some noteworthy differences to consider. For Gaussian denoising, when the repetition times are set to  $R_t = 1$ , the optimal distortion result (the highest PSNR and SSIM) is attained with a sampling time of  $S_t = 1$ . However, with an increase in repetition times, results obtained within the range of  $1 < S_t < 10$  can surpass the outcome obtained with a sampling time of  $S_t = 1$ . For real-world, denoising, the optimal distortion result is always attained with sampling times of  $S_t = 2$  or  $S_t = 3$ . The results are shown in Table 5. This is because real-world denoising is more challenging, and 2-3 iterations tend to yield better results with relatively less stochasticity. Increasing the number of iterations introduces excessive stochasticity, and the performance improvement is not sufficient to compensate for the

TABLE 5  
Results of the ablation study on the ensembling method with varying sampling timesteps  $S_t$ . Increased sampling timesteps  $S_t$  result in higher distortion, particularly when  $S_t > 3$  for real-world denoising.

$S_t$	$R_t$	CC [79]	PolyU [80]	FMDD [81]
		PSNR / SSIM	PSNR / SSIM	PSNR / SSIM
$S_t = 1$	$R_t = 1$	37.09 / 0.9854	38.46 / 0.9843	31.97 / 0.7830
$S_t = 2$	$R_t = 1$	<b>37.84 / 0.9877</b>	<b>38.56 / 0.9849</b>	33.09 / 0.8616
$S_t = 3$	$R_t = 1$	<b>37.82 / 0.9876</b>	<b>38.56 / 0.9849</b>	<b>33.22 / 0.8693</b>
$S_t = 4$	$R_t = 1$	37.73 / 0.9873	<b>38.52 / 0.9847</b>	<b>33.15 / 0.8671</b>
$S_t = 5$	$R_t = 1$	37.67 / 0.9871	38.51 / 0.9847	33.11 / 0.8668
$S_t = N$	$R_t = 1$	37.49 / 0.9866	38.50 / 0.9847	32.85 / 0.8578




Fig. 11. Results of the ablation study on the ensembling method with varying repetition times  $R_t$ . Increased repetition times  $R_t$  result in less distortion, when  $R_t > 1$  for both Gaussian denoising and real-world denoising.

loss caused by increased stochasticity.

This experiment illustrates the capacity of our method to adjust for distortion and perception.

**(b). Influence of  $R_t$ .** We perform experiments on Mc-Master with noise level  $\sigma = 250$  using different repetition times  $R_t$  and the same sampling times  $S_t = 10$ .

More repetition times  $R_t$  can result in denoised images that are more likely to match probability, resulting in less distortion. The increase brought about by the ensembling strategy has a severe marginal utility for distortion. The increase after 10 times is small, but the impact on perception lasts longer as shown in Figure 11. While each candidate image obtained exhibits slight variations, averaging them could potentially lead to the loss of minor details. It is important to emphasize that critical details are, nonetheless, preserved. For instance, in Figure 12, we can observe a grad-

TABLE 6  
Compared with other diffusion-based methods on real-world denoising.

Method	Extension with our method		CC [79]			PolyU [80]			FMDD [81]		
	Embedding	Ensembling	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
DDRM	✗	✗	33.52	0.9575	0.221	36.00	0.9590	0.196	27.23	0.5451	0.450
DDRM+Clean	✗	✗	36.03	0.9788	0.081	38.07	0.9828	0.057	29.61	0.7156	0.291
DDNM	✗	✗	33.52	0.9575	0.222	36.01	0.9590	0.196	27.22	0.5442	0.450
DDNM+Clean	✗	✗	35.27	0.9723	0.126	36.94	0.9700	0.133	29.18	0.6790	0.328
DM+Ours	✓	✗	37.20	0.9855	0.072	38.45	0.9844	0.067	30.99	0.7721	0.308
DM+Ours	✓	✓	37.97	0.9878	0.079	38.57	0.9851	0.070	32.64	0.8788	0.254
DMID-d (Ours)	✓	✓	37.99	0.9880	0.078	38.62	0.9853	0.069	33.40	0.8747	0.266
DMID-p (Ours)	✓	✓	37.09	0.9854	0.072	38.46	0.9843	0.067	33.09	0.8616	0.232

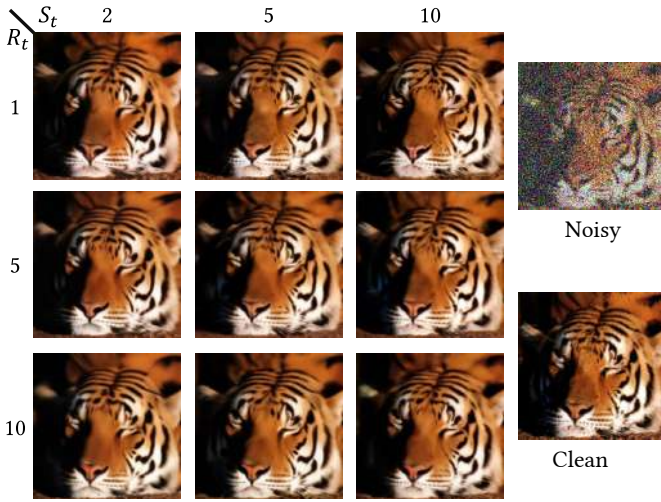


Fig. 12. Visual results of ablation study on the ensembling method. (Best viewed with zoom-in)

ual enrichment of details as the number of sampling times  $S_t$  increases. Furthermore, with an increase in repetition times  $R_t$ , we notice that most details tend to converge closer to the clean image, including fine features such as whiskers, which remain faithfully intact.

This experiment validates the robustness of our analysis and further illustrates the capacity of our method to reduce distortion.

### 5.5 Comparative Analysis and Extension to Diffusion-Based Methods

In this section, we embark on a comprehensive analysis that involves comparing our proposed method with other diffusion-based image restoration methods. Additionally, we explore the adaptability of our method to address specific challenges and limitations encountered by existing diffusion-based methods.

Firstly, we conduct experiments compared with other diffusion-based image restoration methods. DDRM [50] and DDNM [51] are the representative diffusion-based image restoration methods. Tables 7 and 6 showcase the results of Gaussian denoising on ImageNet and real-world denoising on CC [79], PolyU [80], and FMDD [81]. Our method consistently yields optimal results under the same runtime as shown in Figure 13. As DDRM [50] and DDNM [51] do

TABLE 7

Compared with other diffusion-based methods on Gaussian denoising.

Noise Level	DDRM [50]	DDNM [51]	DMID-d (Ours)
	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS	PSNR / SSIM / LPIPS
$\sigma = 50$	28.76 / 0.8950 / 0.163	29.20 / 0.9030 / 0.158	29.90 / 0.9157 / 0.114
$\sigma = 100$	26.12 / 0.8395 / 0.261	26.40 / 0.8484 / 0.249	27.00 / 0.8626 / 0.201
$\sigma = 150$	24.70 / 0.8036 / 0.316	24.77 / 0.8085 / 0.308	25.39 / 0.8236 / 0.263
$\sigma = 200$	23.66 / 0.7737 / 0.362	23.61 / 0.7763 / 0.352	24.25 / 0.7915 / 0.295
$\sigma = 250$	22.81 / 0.7465 / 0.405	22.71 / 0.7487 / 0.387	23.44 / 0.7675 / 0.346




Fig. 13. The comparison of runtime on the ImageNet dataset with noise level  $\sigma = 250$ . Notably, our method consistently yields optimal results under the same runtime.

not focus on image denoising, we do not compare them in the previous Section 5.2 and Section 5.3, and the results in the table are produced by ourselves. The hyperparameters of noise level used in DDRM [50] and DDNM [51] are estimated using the method [90]. Other hyperparameters are employed their recommended settings. Due to their poor performance on real-world denoising, we introduce the clean image as auxiliary information to calculate the noise level and upgrade DDRM and DDNM to “DDRM+Clean” and “DDNM+Clean”, respectively. However, the performance is still limited.

As depicted in Figure 14, both DDRM [50] and DDNM [51] struggle to handle image denoising tasks. That is because these methods are not well-suited for image denoising tasks. As discussed in Section 2.2, they can not circumvent the input inconsistency and overlook the content inconsistency.

Next, we will extend our method to DDRM [50] and DDNM [51]. The sampling strategy in our method can be arbitrary. However, practical applications require some adjustments when applying different sampling strategies. This is because various sampling methods have unique




Fig. 14. Visual results compared with other diffusion-based methods. Other methods can not deal with real-world noise.




Fig. 15. Visual results of the diffusion-based method and extension with our methods.

designs and cannot be directly integrated. Both DDRM [50] and DDNM [51] utilize the same sampling strategy for noise handling, with the primary difference being that DDRM [50] involves a specific initialization step. This initialization, however, conflicts with our embedding method and is therefore replaced. Thus, DDRM [50] and DDNM [51] are the same for our method, and we denote them as DM.

The results are presented in Table 6. For the embedding method, we apply the same noise transformation as in our method. In addition, we convert the image to the intermediate state  $x_N$  and start sampling from  $x_N$ . For the ensembling method, we set the sampling times to be  $S_t = 3$  for CC [79] and PolyU [80], and  $S_t = 11$  for FMDD [81].

The visual results are presented in Figure 15. The versions featuring our embedding method and ensembling

method (referred to "DM +Ours"), outperform their original counterparts DDRM [50] and DDNM [51] easily. Additionally, "DM +Ours" employs significantly fewer sampling times compared to DDRM [50] and DDNM [51], respectively. For example, the sampling times  $S_t$  are 3 for "DM +Ours", whereas DDRM [50] and DDNM [51] require 20 and 100 sampling times, respectively, on the PolyU [80] dataset. DDRM and DDNM require more runtime because DDRM [50] and DDNM [51] start sampling from standard Gaussian noise for image denoising. Moreover, when DDRM and DDNM produce the intermediate image  $x_N$  with a noise level comparable to that of the original input noisy image  $y$ , the content and information of  $x_N$  do not align with those of the original input noisy image  $y$ . Consequently, DDRM [50] and DDNM [51] not only exhibit poor performance but also demand a larger number of sampling times. This experiment further emphasizes the superiority, effectiveness, and scalability of our method.

## 6 DISCUSSION

In this paper, we introduce a method that utilizes the diffusion model pre-trained for image synthesis. At present, there is no established paradigm for effectively leveraging the diffusion model for image denoising. The comparison between trainable and pre-trained solutions raises an intriguing question that has yet to be thoroughly explored. In the subsequent paragraphs, we will analyze the merits and drawbacks of both approaches and aim to draw a conclusion.

The pre-trained solution employs the same pre-trained diffusion model for universal image denoising. Utilizing pre-trained diffusion models eliminates the need for training and provides extensive image priors, enhancing denoising capabilities. Despite these advantages, challenges

related to input inconsistency and content inconsistency are required to be corrected.

The trainable solution customizes diffusion models for various noise types and datasets [52], improving the handling of specific noise. However, complexities arise from the unknown distribution of real noise and the sensor-specific nature of noise [91], leading to improbability in customizing specific models for real noise and inefficiencies in training different models for various noise. In addition, diffusion models are delicate and only excel with Gaussian noise [92], [93]. Customizing diffusion models and altering the noise distribution leads to suboptimal performance [92], [93].

In summary, our solution which employs pre-trained diffusion models offers a flexible solution for handling various types of noise. However, some modifications are required to adapt the pre-trained diffusion model. On the other hand, trainable diffusion models provide a solution for customizing the training of different diffusion models based on distinct noise profiles. Nevertheless, challenges arise in customizing the training process and optimizing the denoising capabilities.

## 7 CONCLUSION

In this paper, we present a novel strategy to stimulate the diffusion model for image denoising. Specifically, we revisit diffusion models from the denoising perspective. Furthermore, we propose an adaptive embedding method to perform denoising and an adaptive ensembling method to reduce distortion. Our method achieves SOTA performance on both distortion-based and perception-based metrics, for both Gaussian and real-world image denoising. In future research endeavors, we intend to stimulate the diffusion model for multiple other restoration tasks.

## REFERENCES

- [1] M. Aharon, M. Elad, and A. Bruckstein, "K-svd: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, pp. 4311–4322, 2006.
- [2] S. Gu, Z. Lei, W. Zuo, and X. Feng, "Weighted nuclear norm minimization with application to image denoising," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2862–2869, 2014.
- [3] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of gaussians in the wavelet domain," *IEEE Transactions on Image processing (TIP)*, pp. 1338–1351, 2003.
- [4] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, "Swinir: Image restoration using swin transformer," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1833–1844, 2021.
- [5] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, "Restormer: Efficient transformer for high-resolution image restoration," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5728–5739, 2022.
- [6] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 6360–6376, 2021.
- [7] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE transactions on image processing (TIP)*, pp. 3142–3155, 2017.
- [8] K. Zhang, W. Zuo, and L. Zhang, "Ffdnet: Toward a fast and flexible solution for cnn-based image denoising," *IEEE Transactions on Image Processing (TIP)*, pp. 4608–4622, 2018.
- [9] M. Mathieu, C. Couprie, and Y. LeCun, "Deep multi-scale video prediction beyond mean square error," *International Conference on Learning Representations (ICLR)*, pp. 1–14, 2016.
- [10] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1–9, 2014.
- [12] O. Kupyn, V. Budzan, M. Mykhailych, D. Mishkin, and J. Matas, "Deblurgan: Blind motion deblurring using conditional adversarial networks," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8183–8192, 2018.
- [13] A. Maleky, S. Kousha, M. S. Brown, and M. A. Brubaker, "Noise2noiseflow: Realistic camera noise modeling without clean images," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17 632–17 641, 2022.
- [14] G. Jang, W. Lee, S. Son, and K. M. Lee, "C2n: Practical generative noise modeling for real-world denoising," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2350–2359, 2021.
- [15] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6840–6851, 2020.
- [16] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," *International Conference on Machine Learning (ICML)*, pp. 2256–2265, 2015.
- [17] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," *International Conference on Learning Representations (ICLR)*, pp. 1–12, 2021.
- [18] H. Li, Y. Yang, M. Chang, S. Chen, H. Feng, Z. Xu, Q. Li, and Y. Chen, "Srdiff: Single image super-resolution with diffusion probabilistic models," *Neurocomputing*, pp. 47–59, 2022.
- [19] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 1–12, 2022.
- [20] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "Repaint: Inpainting using denoising diffusion probabilistic models," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11 461–11 471, 2022.
- [21] Y. Zhu, K. Zhang, J. Liang, J. Cao, B. Wen, R. Timofte, and L. Van Gool, "Denoising diffusion models for plug-and-play image restoration," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1219–1229, 2023.
- [22] J. Whang, M. Delbracio, H. Talebi, C. Saharia, A. G. Dimakis, and P. Milanfar, "Deblurring via stochastic refinement," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16 293–16 303, 2022.
- [23] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 8780–8794, 2021.
- [24] Y. Blau and T. Michaeli, "The perception-distortion tradeoff," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6228–6237, 2018.
- [25] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Transactions on Image processing (TIP)*, pp. 3736–3745, 2006.
- [26] A. Buades, B. Coll, and J. M. Morel, "A non-local algorithm for image denoising," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 60–65, 2005.
- [27] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE Transactions on image processing (TIP)*, pp. 2080–2095, 2007.
- [28] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D Nonlinear Phenomena*, pp. 259–268, 1992.
- [29] Z. Wang, X. Cun, J. Bao, W. Zhou, J. Liu, and H. Li, "Uformer: A general u-shaped transformer for image restoration," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17 683–17 693, 2022.



- [30] H. Chen, Y. Wang, T. Guo, C. Xu, Y. Deng, Z. Liu, S. Ma, C. Xu, C. Xu, and W. Gao, "Pre-trained image processing transformer," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12 299–12 310, 2021.
- [31] S. Cheng, Y. Wang, H. Huang, D. Liu, H. Fan, and S. Liu, "Nbnnet: Noise basis learning for image denoising with subspace projection," *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 4896–4906, 2021.
- [32] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, "Multi-stage progressive image restoration," *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 14 821–14 831, 2021.
- [33] A. Krull, T.-O. Buchholz, and F. Jug, "Noise2void-learning denoising from single noisy images," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2129–2137, 2019.
- [34] J. Batson and L. Royer, "Noise2self: Blind denoising by self-supervision," *International Conference on Machine Learning (ICML)*, pp. 524–533, 2019.
- [35] W. Lee, S. Son, and K. M. Lee, "Ap-bsn: Self-supervised denoising for real-world images via asymmetric pd and blind-spot network," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17 725–17 734, 2022.
- [36] Z. Wang, Y. Fu, J. Liu, and Y. Zhang, "LG-BPN: Local and global blind-patch network for self-supervised real-world denoising," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18 156–18 165, 2023.
- [37] T. Pang, H. Zheng, Y. Quan, and H. Ji, "Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2043–2052, 2021.
- [38] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, "An iterative regularization method for total variation-based image restoration," *Multiscale Modeling & Simulation*, pp. 460–489, 2005.
- [39] W. Dong, X. Li, L. Zhang, and G. Shi, "Sparsity-based image denoising via dictionary learning and structural clustering," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 457–464, 2011.
- [40] W. Dong, P. Wang, W. Yin, G. Shi, F. Wu, and X. Lu, "Denoising prior driven deep neural network for image restoration," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 2305–2318, 2019.
- [41] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," *International Conference on Machine Learning (ICML)*, pp. 1530–1538, 2015.
- [42] I. Marras, G. G. Chrysos, I. Alexiou, G. Slabaugh, and S. Zafeiriou, "Reconstructing the noise manifold for image denoising," *arXiv preprint arXiv:2002.04147*, pp. 1–18, 2020.
- [43] A. Abdelhamed, M. A. Brubaker, and M. S. Brown, "Noise flow: Noise modeling with conditional normalizing flows," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3165–3173, 2019.
- [44] K.-C. Chang, R. Wang, H.-J. Lin, Y.-L. Liu, C.-P. Chen, Y.-L. Chang, and H.-T. Chen, "Learning camera-aware noise models," *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 343–358, 2020.
- [45] Y. Cai, X. Hu, H. Wang, Y. Zhang, H. Pfister, and D. Wei, "Learning to generate realistic noisy images via pixel-level noise-aware adversarial training," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3259–3270, 2021.
- [46] K. Lin, T. H. Li, S. Liu, and G. Li, "Real photographs denoising with noise domain adaptation and attentive generative adversarial network," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 3155–3164, 2019.
- [47] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10 684–10 695, 2022.
- [48] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 11 895–11 907, 2019.
- [49] V. Kulikov, S. Yadin, M. Kleiner, and T. Michaeli, "Sinddm: A single image denoising diffusion model," *International Conference on Machine Learning (ICLR)*, pp. 17 920–17 930, 2023.
- [50] B. Kawar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1–13, 2022.
- [51] Y. Wang, J. Yu, and J. Zhang, "Zero-shot image restoration using denoising diffusion null-space model," *International Conference on Learning Representations (ICLR)*, pp. 1–12, 2022.
- [52] Y. Xie, M. Yuan, B. Dong, and Q. Li, "Diffusion model for generative image denoising," *arXiv preprint arXiv:2302.02398*, 2023.
- [53] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [54] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502*, pp. 1–12, 2020.
- [55] L. H. Sullivan, *The tall office building artistically considered*. Lippincott's Magazine, 1896.
- [56] C. Chen, Q. Chen, M. N. Do, and V. Koltun, "Seeing motion in the dark," 2019, pp. 3185–3194.
- [57] D. Zheng, S. H. Tan, X. Zhang, Z. Shi, K. Ma, and C. Bao, "An unsupervised deep learning approach for real-world image denoising," *International Conference on Learning Representations (ICLR)*, pp. 1–23, 2021.
- [58] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, pp. 1–122, 2011.
- [59] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *The Annals of Statistics*, pp. 1135–1151, 1981.
- [60] S. Soltanayev and S. Y. Chun, "Training deep learning based denoisers without ground truth data," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3261–3271, 2018.
- [61] S. Ramani, T. Blu, and M. Unser, "Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Transactions on Image Processing*, pp. 1540–1554, 2008.
- [62] C. Liu, W. Freeman, R. Szeliski, and S. B. Kang, "Noise estimation from a single image," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 901–908, 2006.
- [63] A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian, "Practical poissonian-gaussian noise modeling and fitting for single-image raw-data," *IEEE Transactions on Image Processing (TIP)*, pp. 1737–1754, 2008.
- [64] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 416–423, 2001.
- [65] R. Franzen, "Kodak lossless true color image suite," *source: <http://r0k.us/graphics/kodak>*, pp. 1–1, 1999.
- [66] L. Zhang, X. Wu, A. Buades, and X. Li, "Color demosaicking by local directional interpolation and nonlocal adaptive thresholding," *Journal of Electronic Imaging*, pp. 023 016–023 016, 2011.
- [67] J. B. Huang, A. Singh, and N. Ahuja, "Single image super-resolution from transformed self-exemplars," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5197–5206, 2015.
- [68] C. Tian, Y. Xu, and W. Zuo, "Image denoising using deep cnn with batch renormalization," *Neural Networks*, pp. 461–473, 2020.
- [69] Y. Zhang, K. Li, K. Li, B. Zhong, and Y. Fu, "Residual non-local attention networks for image restoration," *arXiv preprint arXiv:1903.10082*, pp. 1–13, 2019.
- [70] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, "Residual dense network for image restoration," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 2480–2495, 2020.
- [71] H. Zhao, Y. Gou, B. Li, D. Peng, J. Lv, and X. Peng, "Comprehensive and delicate: An efficient transformer for image restoration," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14 122–14 132, 2023.
- [72] Y. Peng, L. Zhang, S. Liu, X. Wu, Y. Zhang, and X. Wang, "Dilated residual networks with symmetric skip connection for image denoising," *Neurocomputing*, pp. 67–76, 2019.
- [73] D. Guo, S. Shamai, and S. Verdú, "Mutual information and minimum mean-square error in gaussian channels," *IEEE transactions on information theory*, pp. 1261–1282, 2005.
- [74] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision (IJCV)*, pp. 211–252, 2015.

- [75] D. Liu, H. Zhang, and Z. Xiong, "On the classification-distortion-perception tradeoff," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1–8, 2019.
- [76] J. Zhang, Y. Zhang, J. Gu, Y. Zhang, L. Kong, and X. Yuan, "Accurate image restoration with attention retractable transformer," *International Conference on Learning Representations (ICLR)*, pp. 1–13, 2023.
- [77] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing (TIP)*, pp. 600–612, 2004.
- [78] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018.
- [79] S. Nam, Y. Hwang, Y. Matsushita, and S. J. Kim, "A holistic approach to cross-channel image noise modeling and its application to image denoising," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1683–1691, 2016.
- [80] J. Xu, H. Li, Z. Liang, D. Zhang, and L. Zhang, "Real-world noisy image denoising: A new benchmark," *arXiv preprint arXiv:1804.02603*, pp. 1–13, 2018.
- [81] Y. Zhang, Y. Zhu, E. Nichols, Q. Wang, S. Zhang, C. Smith, and S. Howard, "A poisson-gaussian denoising dataset with real fluorescence microscopy images," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11710–11718, 2019.
- [82] Z. Yue, Q. Zhao, L. Zhang, and D. Meng, "Dual adversarial network: Toward real-world noise removal and noise generation," *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 41–58, 2020.
- [83] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, "Learning enriched features for real image restoration and enhancement," *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 492–511, 2020.
- [84] C. Ren, X. He, C. Wang, and Z. Zhao, "Adaptive consistency prior based deep network for image denoising," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8596–8606, 2021.
- [85] L. Chen, X. Chu, X. Zhang, and J. Sun, "Simple baselines for image restoration," *Proceedings of the European conference on computer vision (ECCV)*, pp. 17–33, 2022.
- [86] Y. Quan, M. Chen, T. Pang, and H. Ji, "Self2self with dropout: Learning self-supervised denoising from single image," *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1890–1898, 2020.
- [87] Y. Mansour and R. Heckel, "Zero-shot noise2noise: Efficient image denoising without any data," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14018–14027, 2023.
- [88] J. Cheng, T. Liu, and S. Tan, "Score priors guided deep variational inference for unsupervised real-world single image denoising," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 1–12, 2023.
- [89] A. Abdelhamed, S. Lin, and M. S. Brown, "A high-quality denoising dataset for smartphone cameras," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1692–1700, 2018.
- [90] G. Chen, F. Zhu, and P. Ann Heng, "An efficient statistical method for image noise level estimation," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 477–485, 2015.
- [91] H. Feng, L. Wang, Y. Wang, H. Fan, and H. Huang, "Learnability enhancement for low-light raw image denoising: A data perspective," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 46, no. 1, pp. 370–387, 2024.
- [92] A. Bansal, E. Borgnia, H.-M. Chu, J. Li, H. Kazemi, F. Huang, M. Goldblum, J. Geiping, and T. Goldstein, "Cold diffusion: Inverting arbitrary image transforms without noise," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 41259–41282, 2023.
- [93] A. Jolicœur-Martineau, K. Fatras, K. Li, and T. Kachman, "Diffusion models with location-scale noise," *arXiv preprint arXiv:2304.05907*, pp. 1–5, 2023.



**Tong Li** received the BS degree from Beijing Institute of Technology, China, in 2023. He is currently a Master student with the School of Computer Science and Technology at Beijing Institute of Technology. His research interests include computational photography and image processing.



**Hansen Feng** received the BS degree from the University of Science and Technology Beijing, China, in 2020. He is currently a Ph.D. student with the School of Computer Science and Technology, Beijing Institute of Technology. His research interests include computational photography and image processing. He received the Best Paper Runner-Up Award of ACM MM 2022.



**Lizhi Wang** (Member, IEEE) received the B.S. and Ph.D. degrees from Xidian University, Xi'an, China, in 2011 and 2016, respectively. He is currently a professor with the School of Computer Science and Technology, Beijing Institute of Technology. His research interests include computational photography and image processing. He received the Best Paper Award of IEEE VCIP 2016.



**Lin Zhu** (Member, IEEE) received the B.S. degree in computer science from the Northwestern Polytechnical University, China, in 2014, the M.S. degree in computer science from the North Automatic Control Technology Institute, China, in 2018, and the Ph.D. degree with the School of Electronics Engineering and Computer Science, Peking University, China, in 2022. He is currently an assistant professor with the School of Computer Science, Beijing Institute of Technology, China. His current research interests include image processing, computer vision, neuromorphic computing, and spiking neural network.

age processing, computer vision, neuromorphic computing, and spiking neural network.



**Zhiwei Xiong** (Member, IEEE) received the BS and PhD degrees in electronic engineering from the University of Science and Technology of China (USTC), in 2006 and 2011, respectively. He is currently a professor with USTC. Before that, he was a researcher with Microsoft Research Asia (MSRA). His research interests include computational photography, 3D vision, and biomedical image analysis. He has authored or coauthored more than 100 papers in premium journals and conferences. He received the Best Paper Award of IEEE VCIP 2016 and MSRA Fellowship 2009. He and his students were winners of 8 technical challenges held in CVPR / ICCV / ECCV / MM / ICME / ISBI.

Paper Award of IEEE VCIP 2016 and MSRA Fellowship 2009. He and his students were winners of 8 technical challenges held in CVPR / ICCV / ECCV / MM / ICME / ISBI.



**Hua Huang** (Senior Member, IEEE) received the B.S. and Ph.D. degrees from Xi'an Jiaotong University, in 1996 and 2006, respectively. He is currently a professor in the School of Artificial Intelligence, Beijing Normal University. He is also an adjunct professor at Xi'an Jiaotong University and Beijing Institute of Technology. His main research interests include image and video processing, computational photography, and computer graphics. He received the Best Paper Award of ICML2020 / EURASIP2020 /

PRCV2019 / ChinaMM2017.