

# Efficient Conditional Diffusion Model with Probability Flow Sampling for Image Super-resolution

Yutao Yuan, Chun Yuan

Tsinghua University  
yuanyt21@mails.tsinghua.edu.cn, yuanc@sz.tsinghua.edu.cn

## Abstract

Image super-resolution is a fundamentally ill-posed problem because multiple valid high-resolution images exist for one low-resolution image. Super-resolution methods based on diffusion probabilistic models can deal with the ill-posed nature by learning the distribution of high-resolution images conditioned on low-resolution images, avoiding the problem of blurry images in PSNR-oriented methods. However, existing diffusion-based super-resolution methods have high time consumption with the use of iterative sampling, while the quality and consistency of generated images are less than ideal due to problems like color shifting. In this paper, we propose Efficient Conditional Diffusion Model with Probability Flow Sampling (ECDP) for image super-resolution. To reduce the time consumption, we design a continuous-time conditional diffusion model for image super-resolution, which enables the use of probability flow sampling for efficient generation. Additionally, to improve the consistency of generated images, we propose a hybrid parametrization for the denoiser network, which interpolates between the data-predicting parametrization and the noise-predicting parametrization for different noise scales. Moreover, we design an image quality loss as a complement to the score matching loss of diffusion models, further improving the consistency and quality of super-resolution. Extensive experiments on DIV2K, ImageNet, and CelebA demonstrate that our method achieves higher super-resolution quality than existing diffusion-based image super-resolution methods while having lower time consumption. Our code is available at <https://github.com/Yuan-Yutao/ECDP>.

## Introduction

Image super-resolution, the task of recovering high-resolution (HR) images from low-resolution (LR) images, is fundamentally an ill-posed problem. Given an LR image, there are more than one HR images consistent with the input. Existing PSNR-oriented super-resolution methods (Dong et al. 2014; Lim et al. 2017; Zhang et al. 2018b) that learn deterministic mappings from LR images to HR images using pixel losses are effectively predicting the mean of all plausible HR images, and tend to generate blurry HR images with unsatisfactory visual quality. Super-resolution methods based on generative models deal with the ill-posed

nature by learning the distribution of HR images conditioned on LR images, allowing for the generation of multiple diverse results from a single input image and avoiding the problem of blurry images.

Recently, the use of diffusion probabilistic models (Ho, Jain, and Abbeel 2020; Song et al. 2021), a trending class of generative models, have grown popular in image super-resolution. SR3 (Saharia et al. 2021) and SRDiff (Li et al. 2021) adapts Diffusion Denoising Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020) for image super-resolution. They define a Markovian forward process that gradually adds Gaussian noise into image data, and use denoiser neural networks conditioned on LR images to learn its reverse process and generate new images from noise. They are able to generate diverse and realistic HR images with fine details.

However, there are still challenging aspects that remain to be improved for diffusion-based super-resolution. Diffusion models typically generate new images iteratively using a Markov chain, which necessitates many neural network evaluations and makes the super-resolution process **time-consuming**. Additionally, they are prone to problems like color shifting, making the **quality and consistency** of generated images less than ideal and reducing their performance on super-resolution.

To tackle the challenges, we propose Efficient Conditional Diffusion Model with Probability Flow Sampling (ECDP) for image super-resolution. It gradually corrupts HR images using stochastic differential equations (SDEs), and learns to restore the original images with a denoiser network conditioned on LR images. We generate super-resolution images using probability flow sampling, which can be performed with low time consumption using ordinary differential equation (ODE) solvers. Additionally, to improve the consistency of generated images with LR input, we use a hybrid parametrization in the denoiser network. It uses the  $x_0$ -parametrization that predicts the clean data directly in addition to the commonly used  $\epsilon$ -parametrization, and smoothly interpolates between them for different noise scales. Moreover, we introduce an image quality loss as a complement to the score matching loss of diffusion models. It measures the feature-space distance of generated HR images and ground-truth images in the dataset, further improving the consistency and quality of super-resolution. Extensive experiments

on multiple datasets encompassing face super-resolution and general super-resolution demonstrate the effectiveness of our approach.

Our main contributions are summarized as follows:

- We propose Efficient Conditional Diffusion Model with Probability Flow Sampling (ECDP) for image super-resolution, which generates realistic super-resolution images with low time costs.
- With a continuous-time conditional diffusion model based on SDEs designed for super-resolution, we can generate super-resolution images using probability flow sampling, which reduces the time consumption of super-resolution.
- We propose score matching with hybrid-parametrization and design an image quality loss for diffusion-based image super-resolution, improving the consistency and quality of generated images.
- Extensive experiments on DIV2K, ImageNet, and CelebA demonstrate that our method achieves higher super-resolution quality than existing diffusion-based image super-resolution methods while having lower time consumption.

## Related Work

**Diffusion probabilistic models** Diffusion probabilistic models are a family of deep generative models with great success in image generation. DDPM (Ho, Jain, and Abbeel 2020) defines a Markovian diffusion process on image data, gradually adding noise into the image, and learns to reproduce the original image with a sequence of denoisers. It achieves impressive high-quality image generation results. Various improvements to the model have been proposed, including improved architectures (Dhariwal and Nichol 2021) loss reweighting (Nichol and Dhariwal 2021), and fast sampling (Song, Meng, and Ermon 2021). DDPMs are shown to be equivalent to denoising score matching over multiple noise levels (Song and Ermon 2019), which are unified and generalized to the continuous case with SDEs (Song et al. 2021). We build our super-resolution method on top of the SDE formulation by extending it and conditioning on LR images.

Besides image generation, diffusion models have been applied to a large range of tasks in computer vision. A popular strategy among them is to formulate the task as a conditional generation problem, using diffusion models to predict the distribution of outputs conditioned on the inputs. This strategy has achieved success in text-to-image generation (Saharia et al. 2022b), image super-resolution (Saharia et al. 2021; Li et al. 2021), image inpainting (Saharia et al. 2022a), and image colorization (Saharia et al. 2022a), among others. A different line of research focuses on using existing diffusion models in a zero-shot manner. By taking an unconditional diffusion model and enforcing consistency with reference images during sampling, it is possible to perform image editing (Meng et al. 2022; Choi et al. 2021) and image inpainting (Lugmayr et al. 2022) without task-specific training. More recently, this approach has been generalized for

a family of linear and non-linear inverse problems (Kawar et al. 2022; Wang, Yu, and Zhang 2023; Chung et al. 2023).

**Image super-resolution** A lot of super-resolution methods based on deep learning have been proposed in recent years. Most of the early work takes a regression-based approach (Dong et al. 2014; Lim et al. 2017; Zhang et al. 2018b), learning a deterministic one-to-one mapping from LR images to HR images with L2 or L1 losses. Since the posterior distribution of HR images is highly multimodal, these methods tend to generate blurry images, effectively predicting the mean of the distribution. To improve the visual quality of generated images, GAN-based approaches (Ledig et al. 2017; Wang et al. 2018) are proposed for super-resolution. They are able to generate HR images with high quality, but tend to suffer from mode collapse, difficult optimization and low consistency with LR images. Normalizing flows have also been used for image super-resolution (Lugmayr et al. 2020; Liang et al. 2021). They are able to estimate the distribution of HR images conditioned on LR images, allowing for diverse and realistic image super-resolution. However, normalizing flows require invertible architectures, limiting the expressiveness of the models.

Recently, several super-resolution methods using diffusion models have been proposed. SR3 (Saharia et al. 2021) and SRDiff (Li et al. 2021) adapts DDPM (Ho, Jain, and Abbeel 2020) for super-resolution, making the model conditional on LR images. They are able to generate realistic, high quality HR images. SR3 concatenates upscaled LR images to noisy HR images as the input to the denoiser network, making the model conditional on LR images. SRDiff uses an LR encoder to extract features from LR input, and further uses residual prediction to improve the convergence speed and performance of the model. Some methods (Chung, Sim, and Ye 2022; Chung et al. 2023; Fei et al. 2023; Zhu et al. 2023) use a different zero-shot approach as opposed to aforementioned supervised methods, taking an unconditional diffusion model trained on image generation and modifying its sampling process with guidance. However, their performance is often less than ideal compared to supervised methods due to lack of dedicated training.

## Preliminaries

Given a dataset  $\mathcal{X} = \{x_i\}$  that follows an unknown distribution  $p(x)$ , continuous-time diffusion probabilistic models define a forward process that gradually injects noise into data  $x$ . This process can be described as the solution to a SDE running from  $t = 0$  to  $t = T$ , starting with i.i.d. samples  $x(0)$  from the dataset (Song et al. 2021):

$$dx = f(x, t)dt + g(t)dw \quad (1)$$

where  $f(x, t)$  and  $g(t)$  are predefined functions, and  $w$  is Brownian motion. Denote by  $x(t)$  the solution at time  $t$  and  $p_t(x)$  its probability distribution. The parameters of the forward process are chosen so  $p_T$  ends up as a prior distribution with tractable sampling.

It has been shown (Song et al. 2021) that the data distribution  $p_0$  can be recovered from  $p_T$  using another reverse SDE,

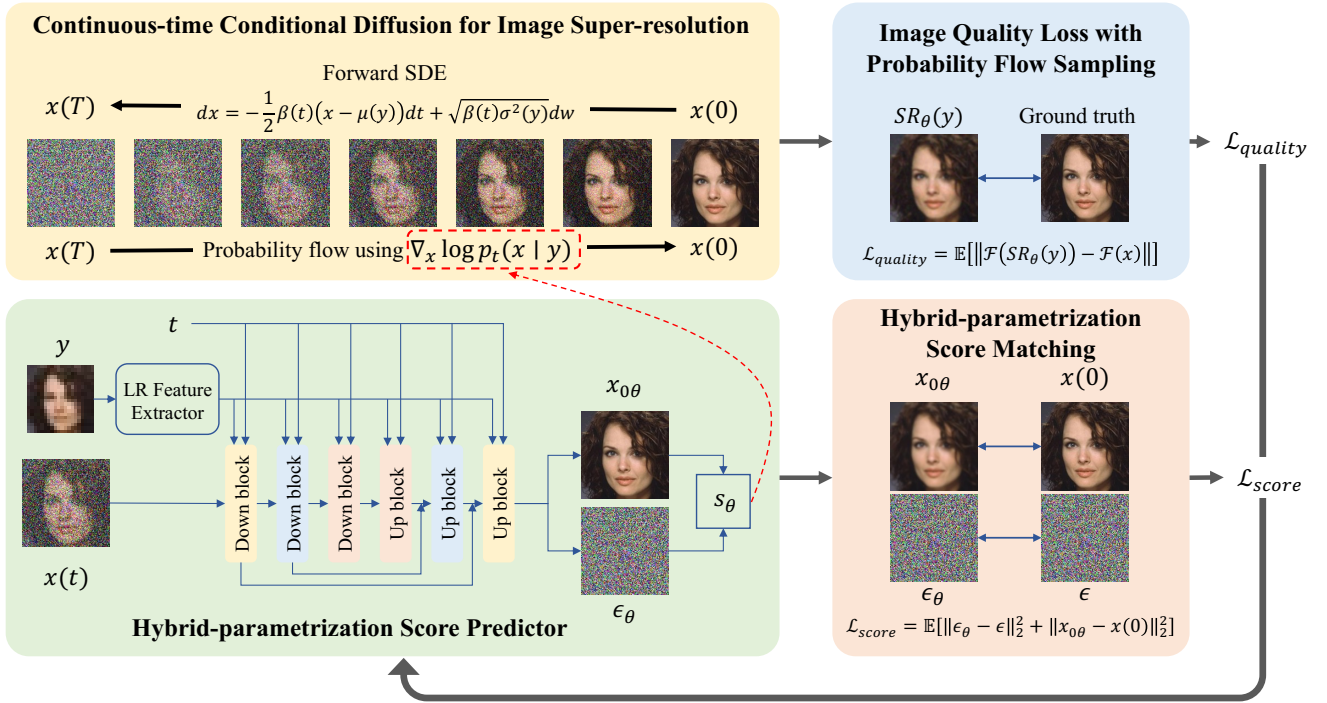


Figure 1: Overview of ECDP. *Top left*: Continuous-time conditional diffusion uses a forward SDE to transform images into noise, and generate new images from noise using probability flow. *Bottom*: The conditional score in the probability flow is approximated with a hybrid-parametrization score predictor  $s_\theta$ , which is trained using score matching. *Top right*: An additional image quality loss that compares the generated HR images with the ground truth is computed using probability flow sampling, improving the quality of super-resolution results.

running backwards from  $t = T$  to  $t = 0$ :

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t) d\mathbf{w} \quad (2)$$

The term  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$  is the score of the distribution  $p_t(\mathbf{x})$ , which is intractable because the data distribution is unknown. Diffusion models learn the data distribution by approximating  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$  with a score prediction network  $s_\theta(\mathbf{x}, t)$ . Due to the intractability of the marginal distribution  $p_t(\mathbf{x})$ , score matching techniques are deployed in training, giving rise to the score matching loss:

$$\mathcal{L} = \mathbb{E}_{\mathbf{x}, t} \mathbb{E}_{\mathbf{x}(t)} \left[ \left\| s_\theta(\mathbf{x}(t), t) - \nabla_{\mathbf{x}(t)} \log p_{0t}(\mathbf{x}(t) | \mathbf{x}) \right\|_2^2 \right] \quad (3)$$

where  $p_{0t}$  is the transition probability of the forward process and can often be computed analytically. In practice the loss is often reweighted, where terms associated with different  $t$  are assigned different weights, to ensure better convergence in training.

Besides the reverse SDE, it is also possible to sample from the learned distribution using probability flow, which takes the form of an ODE:

$$d\mathbf{x} = \left[ \mathbf{f}(\mathbf{x}, t) - \frac{1}{2} g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt \quad (4)$$

It can be solved using off-the-shelf ODE solvers, allowing

efficient sampling using much less time than the reverse SDE.

## Proposed Method

In this section, we present Efficient Conditional Diffusion Model with Probability Flow Sampling (ECDP) for image super-resolution. From a dataset of paired HR and LR images  $\mathcal{D} = \{(x_i, y_i)\}$ , our model learns the conditional distribution  $p(\mathbf{x} | \mathbf{y})$ . Given LR images, super-resolution images are generated by sampling from this conditional distribution. Our method is illustrated in Figure 1.

### Continuous-Time Conditional Diffusion for Image Super-Resolution

In our method, we design a conditional diffusion model for image super-resolution with a forward SDE that transforms HR images  $\mathbf{x}$  into noise while conditioned on LR images  $\mathbf{y}$ :

$$d\mathbf{x} = -\frac{1}{2} \beta(t) (\mathbf{x} - \boldsymbol{\mu}(\mathbf{y})) dt + \sqrt{\beta(t) \sigma^2(\mathbf{y})} d\mathbf{w} \quad (5)$$

where  $\boldsymbol{\mu}(\mathbf{y})$  and  $\sigma^2(\mathbf{y})$  are the per-pixel mean and variance of  $p(\mathbf{x} | \mathbf{y})$  respectively, and  $\beta(t)$  is a hyperparameter that controls how fast noise is injected into data.

Since it is impossible to compute  $\boldsymbol{\mu}(\mathbf{y})$  and  $\sigma^2(\mathbf{y})$  without direct access to the true data distribution, we approximate  $\boldsymbol{\mu}(\mathbf{y})$  by upscaling  $\mathbf{y}$  using bicubic interpolation, and

set  $\sigma^2(\mathbf{y})$  to a predefined constant determined using the empirical variance of  $\mathbf{x} - \boldsymbol{\mu}(\mathbf{y})$  over the training dataset. It is worth noting that the use of  $\boldsymbol{\mu}(\mathbf{y})$  in the forward process is similar to residual prediction (Li et al. 2021), which subtracts upscaled LR images from HR images before diffusion. We differ from residual prediction in that we integrate the LR images to the forward process of diffusion directly, and we additionally considers the variance of HR images.

With the use of  $\boldsymbol{\mu}(\mathbf{y})$  and  $\sigma^2(\mathbf{y})$ , the forward process has the mean-preserving and variance-preserving properties, as formalized below:

**Proposition 1.** *The forward process given by (5) keeps the mean and variance of  $\mathbf{x}(t)$  conditioned on  $\mathbf{y}$  unchanged during the transform from  $t = 0$  to  $t = T$ . More specifically:*

$$\mathbb{E}[\mathbf{x}(t) | \mathbf{y}] = \boldsymbol{\mu}(\mathbf{y}) \tag{6}$$

$$\text{Var}[\mathbf{x}(t) | \mathbf{y}] = \sigma^2(\mathbf{y}) \tag{7}$$

By maintaining the mean and variance of data during the forward process, the amount of change in the data distribution is minimized, making model training easier and image generation faster.

To enable conditional generation of HR images, we learn to approximate the conditional score  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x} | \mathbf{y})$  with a conditional score prediction network  $\mathbf{s}_{\theta}(\mathbf{x}, \mathbf{y}, t)$ . The score prediction network is trained with denoising score matching (Vincent 2011) using the following loss:

$$\mathcal{L}_{\text{score}} = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \mathbb{E}_t \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[ \left\| \mathbf{s}_{\theta} \left( \hat{\boldsymbol{\mu}}(\mathbf{x}, \mathbf{y}, t) + \sqrt{\hat{\sigma}^2(\mathbf{y}, t)} \boldsymbol{\epsilon}, \mathbf{y}, t \right) + \frac{\boldsymbol{\epsilon}}{\sqrt{\hat{\sigma}^2(\mathbf{y}, t)}} \right\|_2^2 \right] \tag{8}$$

where

$$\hat{\boldsymbol{\mu}}(\mathbf{x}, \mathbf{y}, t) = \sqrt{\alpha(t)}(\mathbf{x} - \boldsymbol{\mu}(\mathbf{y})) + \boldsymbol{\mu}(\mathbf{y}) \tag{9}$$

$$\hat{\sigma}^2(\mathbf{y}, t) = (1 - \alpha(t))\sigma^2(\mathbf{y}) \tag{10}$$

$$\alpha(t) = \exp \left( - \int_0^t \beta(s) ds \right) \tag{11}$$

For efficient generation of super-resolution images, we sample HR images with probability flow, which can be much faster than SDE-based sampling due to its deterministic nature. Given an LR image  $\mathbf{y}$ , new HR images can be sampled from the learned conditional distribution by solving the following ODE from  $t = T$  to  $t = 0$ , starting with  $\mathbf{x}(T)$  sampled from the limit distribution of the forward SDE:

$$\mathbf{x}(T) \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{y}), \sigma^2(\mathbf{y})\mathbf{I}) \tag{12}$$

$$d\mathbf{x} = \left[ -\frac{1}{2}\beta(t)(\mathbf{x} - \boldsymbol{\mu}(\mathbf{y})) - \frac{1}{2}\beta(t)\sigma^2(\mathbf{y})\mathbf{s}_{\theta}(\mathbf{x}, \mathbf{y}, t) \right] dt \tag{13}$$

### Hybrid-Parametrization Score Matching

Existing diffusion probabilistic models (Ho, Jain, and Abbeel 2020; Song et al. 2021) typically parametrize the denoiser network in such a way that its output is matched




Figure 2: The denoised images produced by the  $\epsilon$ -parametrization, the  $x_0$ -parametrization and the hybrid-parametrization.

against the normalized noise component  $\boldsymbol{\epsilon}$  in noisy data  $\mathbf{x}(t)$  during training, so the network output has identical variance across different  $t$ . We denote it the  $\epsilon$ -parametrization.  $\epsilon$ -parametrization has been proved effective in image generation (Ho, Jain, and Abbeel 2020), improving the quality of denoising predictions when the amount of noise is low. However,  $\epsilon$ -parametrization alone does not work well for image super-resolution. In image super-resolution, the space of plausible HR images  $\mathbf{x}$  is strongly constrained by the consistency with paired LR images  $\mathbf{y}$  and concentrated around a single point.  $\epsilon$ -parametrization has difficulty learning this consistency constraint, because it can only recover the clean HR images indirectly by subtracting the predicted noise from noisy images, requiring wildly varying prediction values for different noisy data. It tends to produce very inconsistent denoising results when the amount of noise is large. This effect can be seen in Figure 2a, where  $\epsilon$ -parametrization does not produce satisfactory denoising predictions, leaving artifacts in the denoised image. As a result, super-resolution methods using the  $\epsilon$ -parametrization tends to produce HR images inconsistent with LR input.

A natural alternative to  $\epsilon$ -parametrization is to predict the clean data component in noisy data as opposed to the noise component. We denote it the  $x_0$ -parametrization.  $x_0$ -parametrization has been investigated in the context of unconditional image generation (Benny and Wolf 2022; Salimans and Ho 2022), where it was found to have better performance than  $\epsilon$ -parametrization in certain cases. In image super-resolution,  $x_0$ -parametrization produces clean images with no artifacts even when the amount of noise is large, because the denoiser network can use the LR inputs to recover a good estimate of clean HR images directly.

For the conditional image diffusion in our method,  $\epsilon$ -parametrization and  $x_0$ -parametrization are defined by expressing the score prediction values in terms of neural networks  $\boldsymbol{\epsilon}_{\theta}$  and  $\mathbf{x}_{0\theta}$  respectively:

$$\mathbf{s}_{\theta, \epsilon}(\mathbf{x}, \mathbf{y}, t) = -\frac{1}{\sqrt{\hat{\sigma}^2(\mathbf{y}, t)}} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}, \mathbf{y}, t) \tag{14}$$

$$\mathbf{s}_{\theta, x_0}(\mathbf{x}, \mathbf{y}, t) = -\frac{1}{\hat{\sigma}^2(\mathbf{y})} (\mathbf{x} - \hat{\boldsymbol{\mu}}(\mathbf{x}_{0\theta}(\mathbf{x}, \mathbf{y}, t), \mathbf{y}, t)) \tag{15}$$

where  $\boldsymbol{\epsilon}_{\theta}$  and  $\mathbf{x}_{0\theta}$  are trained using the following reweighted

version of the score matching loss (8):

$$\mathcal{L}_{\text{score}} = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \mathbb{E}_t \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[ \|\epsilon_{\theta}(\mathbf{x}_t, \mathbf{y}, t) - \epsilon\|_2^2 + \|\mathbf{x}_{0\theta}(\mathbf{x}_t, \mathbf{y}, t) - \mathbf{x}\|_2^2 \right] \quad (16)$$

where

$$\mathbf{x}_t = \hat{\boldsymbol{\mu}}(\mathbf{x}, \mathbf{y}, t) + \sqrt{\hat{\sigma}^2(\mathbf{y}, t)}\epsilon \quad (17)$$

As the  $\epsilon$ -parametrization has low estimation errors in the low-noise region, while  $x_0$ -parametrization has low estimation errors in the high-noise region, to take advantage of both parametrizations, we use a hybrid approach to smoothly interpolate between these two parametrizations:

$$s_{\theta}(\mathbf{x}, \mathbf{y}, t) = \lambda(t)s_{\theta, \epsilon}(\mathbf{x}, \mathbf{y}, t) + (1 - \lambda(t))s_{\theta, x_0}(\mathbf{x}, \mathbf{y}, t) \quad (18)$$

where  $\lambda(t)$  is the interpolation coefficient. In our method, we choose  $\lambda(t) = \alpha(t)^c$ , where  $c$  is a constant in the range  $[0.5, 1.5]$ , selected for each dataset individually to minimize the prediction error for the hybrid parametrization. The hybrid-parametrization is able to achieve low estimation errors in all noise scales.

### Image Quality Loss with Probability Flow Sampling

Diffusion probabilistic models do not directly optimize for the quality of generated images during training. Instead, the quality is only optimized indirectly in the learning of data distribution with score matching loss. In image super-resolution, each LR image has only one paired HR image in the dataset, and it is difficult for diffusion models to estimate the conditional distribution of HR images accurately from this single data point. It is therefore important for diffusion-based super-resolution methods to add additional image quality losses into training.

In CNN-based super-resolution methods, losses like the pixel loss and the perceptual loss directly measure the distance between super-resolution images and the ground truth. However, these losses have not received usage in diffusion-based super-resolution methods, as they require generating new HR images during training, which is computationally expensive for diffusion models using stochastic sampling.

With the use of probability flow sampling in our method for efficient super-resolution, it becomes possible to introduce such losses into the training of the diffusion model. Therefore, we propose an image quality loss for diffusion-based image super-resolution, defined as the feature-space distance between generated HR images and the ground truth:

$$\mathcal{L}_{\text{quality}} = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [\|\mathcal{F}(\text{SR}_{\theta}(\mathbf{y})) - \mathcal{F}(\mathbf{x})\|] \quad (19)$$

where  $\text{SR}_{\theta}(\mathbf{y})$  is an HR image sample generated using the probability flow.  $\mathcal{F}$  is chosen as the feature maps of a VGG network pretrained on image classification, making  $\mathcal{L}_{\text{quality}}$  equivalent to the perceptual loss in CNN-based methods, but in principle it can be any function that converts images to feature vectors.

To compute gradients of the image quality loss with regards to the network parameters, it is necessary to backpropagate through  $\text{SR}_{\theta}(\mathbf{y})$ , the solution of the probability




Figure 3: Visualization of images generated by our method trained without and with  $\mathcal{L}_{\text{quality}}$ . The image generated by the model with  $\mathcal{L}_{\text{quality}}$  has more visible structure (the lines on the pillar) and less background noise.

flow ODE. This can be achieved efficiently using the adjoint method (Chen et al. 2018), which expresses the gradient of ODE solutions with regards to model parameters in terms of another augmented ODE, making it possible to compute gradients without depending on the intermediate values of the original ODE. Compared with direct backpropagation through the ODE solver, the memory consumption of computing the image quality loss is reduced from  $O(s)$  to  $O(1)$ , where  $s$  is the number of sampling steps.

As seen in Figure 3, the image quality loss can significantly improve the quality of super-resolution images for diffusion-based methods. The model trained with  $\mathcal{L}_{\text{quality}}$  is able to generate cleaner images with more visible structure and less background noise compared to the model trained without  $\mathcal{L}_{\text{quality}}$ .

## Experiments

In this section, we conduct experiments on multiple datasets encompassing both face images and general images to demonstrate the performance of our method.

### Experimental Setup

**Model architecture and hyperparameters** In our experiments, we set  $\beta(t)$  in our forward process (5) to a linear function increasing from  $\beta(0) = 0.1$  and  $\beta(T) = 20$ , matching the settings in VP-SDE (Song et al. 2021). Score prediction from  $\epsilon$ -parametrization and  $x_0$ -parametrization are produced using a single network with two output heads, using an U-Net architecture with BigGAN (Brock, Donahue, and Simonyan 2019) residual blocks. To make the model conditional on LR images, we use an LR feature extractor with the RRDB (Wang et al. 2018) architecture and add the features to each layer of the score prediction network. To improve the training efficiency, we use only the standard score matching loss in early stages of training, and then add our image quality loss later on. To generate HR images, we perform probability flow sampling using a standard Runge-Kutta ODE solver with absolute and relative error tolerance of  $10^{-4}$ .

**Datasets** For general image super-resolution ( $4\times$ ), we evaluate the performance of various methods on two datasets, DIV2K (Agustsson and Timofte 2017) and ImageNet (Russakovsky et al. 2015). For DIV2K, models are trained using random HR image crops of size  $160 \times 160$  and evaluated with full-size images. The training dataset is augmented with images from Flickr2K following the practice in earlier methods (Lugmayr et al. 2020; Li et al. 2021). For ImageNet, images are center cropped and resized to  $256 \times 256$  for HR following SR3 (Saharia et al. 2021), and further downsampled to  $64 \times 64$  with bicubic interpolation for LR.

For face image super-resolution ( $8\times$ ), we train and evaluate on CelebA. Images are cropped and resized to  $160 \times 160$  following the procedures in (Lugmayr et al. 2020), and then downsampled as LR using bicubic interpolation.

**Baselines** We compare our image super-resolution method with the following baselines: PSNR-oriented method RRDB (Wang et al. 2018), GAN-based method ESRGAN (Wang et al. 2018), normalizing flow-based methods SRFlow (Lugmayr et al. 2020) and HCFlow (Liang et al. 2021), as well as several diffusion-based methods, SR3 (Saharia et al. 2021), SRDiff (Li et al. 2021), IR-SDE (Luo et al. 2023), GDP (Fei et al. 2023), and DiffPIR (Zhu et al. 2023). Among diffusion-based methods, GDP and DiffPIR uses pretrained unconditional diffusion models, while other methods including ours train conditional models for super-resolution. We use official results for baselines where available, and train the models from scratch otherwise.

**Evaluation metrics** In addition to PSNR and SSIM, two standard metrics for image super-resolution, we use LPIPS (Zhang et al. 2018a) to measure the visual quality of super-resolution results. It is a perceptual metric that is known to correlate better with human perception than traditional metrics like PSNR and SSIM, and is therefore considered the main metric in our experiments.

## Results of Image Super-Resolution

**General image super-resolution** The quantitative results on DIV2K and ImageNet are shown in Table 1 and 2 respectively, and the visual results are shown in Figure 4. On both datasets, our method outperforms all baseline methods in LPIPS and achieves the best overall super-resolution quality. RRDB generates blurry images and has high LPIPS distance despite achieving high PSNR and SSIM scores, further confirming that PSNR and SSIM does not correlate well with perceptual quality. SR3 has good PSNR and SSIM metrics on ImageNet, but its performance is poor in the main metric LPIPS. GDP and DiffPIR, the two diffusion-based methods that use pretrained unconditional models as opposed to training dedicated conditional models, have inferior performance in the three metrics compared to other diffusion-based methods, demonstrating the importance of designing diffusion models specifically for image super-resolution.

In addition, the sampling speed of diffusion-based super-resolution methods in our experiments is listed in Table 3. Our method is able to perform super-resolution with the least

Method	LPIPS↓	PSNR↑	SSIM↑	#Params
RRDB	0.253	29.44	0.84	16.7M
ESRGAN	0.124	26.22	0.75	16.7M
SRFlow	0.120	27.09	0.76	39.5M
HCFlow	0.110	26.61	0.74	23.2M
SR3	0.175	25.90	0.75	97.8M
SRDiff	0.136	27.41	<b>0.79</b>	37.6M
IR-SDE	0.231	25.90	0.66	137.1M
Ours	<b>0.108</b>	<b>28.03</b>	<b>0.79</b>	42.6M

Table 1: Results of  $4\times$  image super-resolution on DIV2K. Diffusion-based methods and non-diffusion-based methods are grouped together respectively. Best results among diffusion-based methods are highlighted in bold.

Method	LPIPS↓	PSNR↑	SSIM↑	#Params
RRDB	0.245	27.23	0.78	16.7M
ESRGAN	0.123	24.18	0.67	16.7M
SRFlow	0.142	24.09	0.67	39.5M
HCFlow	0.129	25.07	0.70	23.2M
SR3	0.191	<b>26.40</b>	<b>0.76</b>	625M
SRDiff	0.154	24.04	0.59	37.6M
GDP	0.240	24.42	0.68	—
DiffPIR	0.219	25.19	0.70	—
Ours	<b>0.110</b>	25.81	0.74	37.6M

Table 2: Results of  $64 \times 64 \rightarrow 256 \times 256$  image super-resolution on ImageNet. Best results among diffusion-based methods are highlighted in bold. #Params for GDP and DiffPIR are omitted because they use pretrained unconditional diffusion models and do not train new models on their own.

Method	SR3	SRDiff	IR-SDE
Inference Time	83.1s	2.4s	6.2s
Method	GDP	DiffPIR	Ours
Inference Time	94.6s	13.0s	<b>2.0s</b>

Table 3: Time required to generate a single  $256 \times 256$  HR image for diffusion-based methods.

Method	LPIPS↓	PSNR↑	SSIM↑	#Params
RRDB	0.230	26.59	0.77	16.7M
ESRGAN	0.120	22.88	0.63	16.7M
SRFlow	0.110	25.24	0.71	40.0M
HCFlow	0.090	24.83	0.69	27.0M
SR3	0.109	24.26	0.68	97.8M
SRDiff	0.106	25.38	<b>0.74</b>	12.0M
Ours	<b>0.097</b>	<b>25.48</b>	0.73	40.0M

Table 4: Results of face image super-resolution on CelebA. Best results among diffusion-based methods are highlighted in bold.

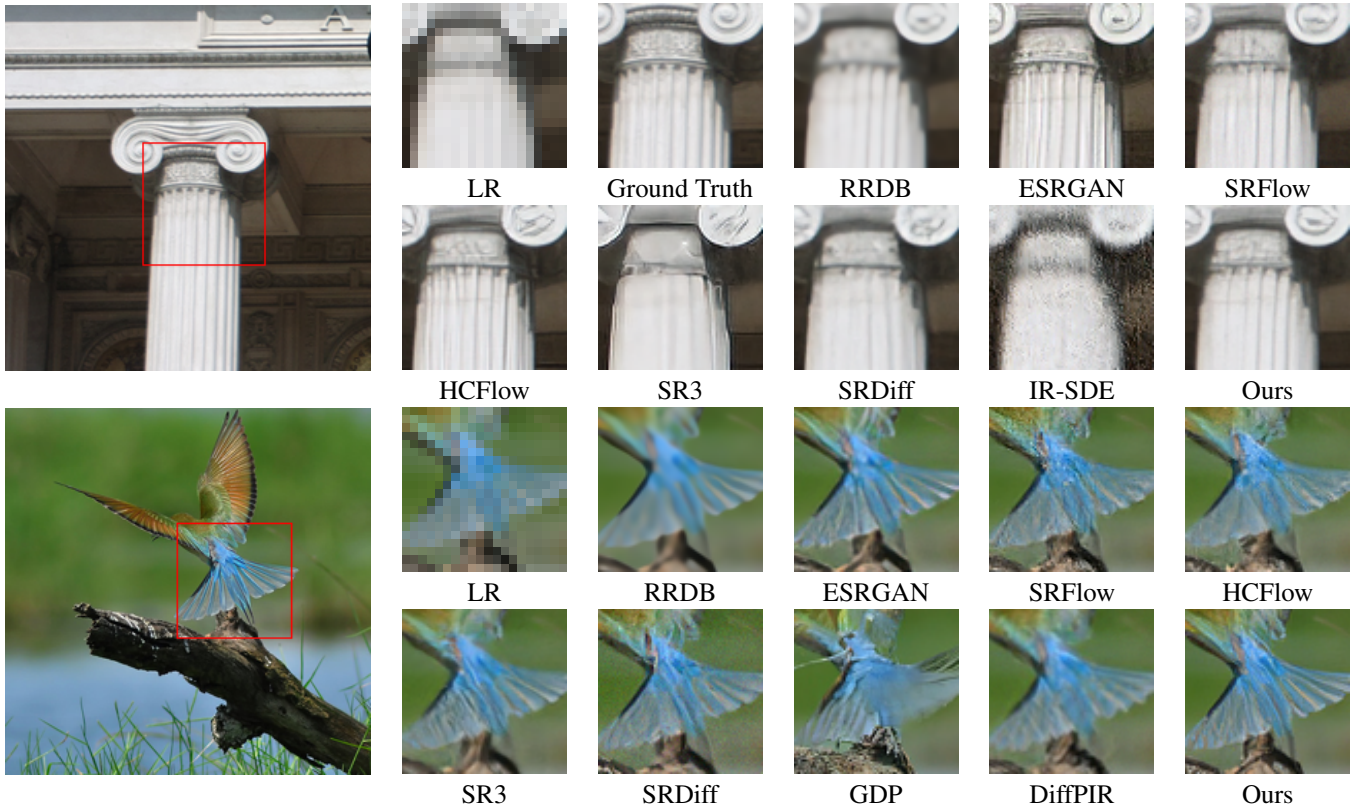


Figure 4: Visual results of general image super-resolution. The first two rows are results on DIV2K validation set. The last two rows are results on ImageNet dev set.

Parametrization	$\mathcal{L}_{\text{quality}}$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$
hybrid, $c = 1.0$	✓	0.108	28.03	0.79
hybrid, $c = 0.5$	✓	0.109	28.07	0.79
hybrid, $c = 1.5$	✓	0.107	28.01	0.79
$\epsilon$ -parametrization	✓	0.128	28.02	0.79
$x_0$ -parametrization	✓	0.112	27.72	0.77
hybrid, $c = 1.0$	×	0.120	27.26	0.76

Table 5: Results of the ablation study. All results are measured on the DIV2K 4 $\times$  task.

amount of time among all diffusion-based methods, confirming the efficiency of probability flow sampling.

**Face image super-resolution** The results on CelebA are shown in Table 4. Similar to the general image super-resolution case, our method reaches state-of-the-art performance and has the best LPIPS score among diffusion-based methods. Our method is able to produce faces with realistic details without generating unnecessary noise and distorting the images.

### Ablation Study

To study the influence of difference choices of parametrizations in the denoiser network as well as the use of im-

age quality loss, we conduct ablation studies as illustrated in Table 5. Our hybrid-parametrization achieves the best super-resolution results in all metrics among three choices of parametrizations since it has the advantages of both  $\epsilon$  and  $x_0$ -parametrization.  $\epsilon$ -parametrization has the lowest LR-PSNR score among all, confirming our analysis that it generates images with low consistency. The performance of hybrid-parametrization is mostly insensitive to the choice of  $c$ . For the image quality loss, it can be seen that the model trained with  $\mathcal{L}_{\text{quality}}$  achieves significantly better metrics than the model without  $\mathcal{L}_{\text{quality}}$ , confirming its importance for diffusion-based super-resolution.

## Conclusion

In this paper, we proposed ECDP, an image super-resolution framework with a continuous-time conditional diffusion model. It deploys a hybrid-parametrization denoiser network to learn the conditional score function, and generates super-resolution images efficiently using probability flow sampling. An additional image quality loss for diffusion-based super-resolution is introduced, which is computed efficiently and improves the quality of super-resolution results. Experiments demonstrate that our method achieves higher super-resolution quality than existing diffusion-based image super-resolution methods while having lower time consumption.

## Acknowledgments

This work was supported by the National Key R&D Program of China (2022YFB4701400/4701402), SSTIC Grant (JCYJ20190809172201639, WZC20200820200655001), Shenzhen Key Laboratory (ZDSYS20210623092001004), and Beijing Key Lab of Networked Multimedia.

## References

- Agustsson, E.; and Timofte, R. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*, 1122–1131. IEEE Computer Society.
- Benny, Y.; and Wolf, L. 2022. Dynamic Dual-Output Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 11472–11481. IEEE.
- Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. 2018. Neural Ordinary Differential Equations. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 6572–6583.
- Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, 14347–14356. IEEE.
- Chung, H.; Kim, J.; McCann, M. T.; Klasky, M. L.; and Ye, J. C. 2023. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Chung, H.; Sim, B.; and Ye, J. C. 2022. Come-Closer-Diffuse-Faster: Accelerating Conditional Diffusion Models for Inverse Problems through Stochastic Contraction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 12403–12412. IEEE.
- Dhariwal, P.; and Nichol, A. Q. 2021. Diffusion Models Beat GANs on Image Synthesis. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 8780–8794.
- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2014. Learning a Deep Convolutional Network for Image Super-Resolution. In Fleet, D. J.; Pajdla, T.; Schiele, B.; and Tuytelaars, T., eds., *European Conference on Computer Vision, Lecture Notes in Computer Science*, 184–199. Springer.
- Fei, B.; Lyu, Z.; Pan, L.; Zhang, J.; Yang, W.; Luo, T.; Zhang, B.; and Dai, B. 2023. Generative Diffusion Prior for Unified Image Restoration and Enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 9935–9946. IEEE.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022. Denoising Diffusion Restoration Models. In *NeurIPS*.
- Ledig, C.; Theis, L.; Huszar, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A. P.; Tejani, A.; Totz, J.; Wang, Z.; and Shi, W. 2017. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. In *IEEE Conference on Computer Vision and Pattern Recognition*, 105–114. IEEE Computer Society.
- Li, H.; Yang, Y.; Chang, M.; Feng, H.; Xu, Z.; Li, Q.; and Chen, Y. 2021. SRDiff: Single Image Super-Resolution with Diffusion Probabilistic Models. arXiv:2104.14951.
- Liang, J.; Lugmayr, A.; Zhang, K.; Danelljan, M.; Gool, L. V.; and Timofte, R. 2021. Hierarchical Conditional Flow: A Unified Framework for Image Super-Resolution and Image Rescaling. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, 4056–4065. IEEE.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Lee, K. M. 2017. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 1132–1140. IEEE Computer Society.
- Lugmayr, A.; Danelljan, M.; Gool, L. V.; and Timofte, R. 2020. SRFlow: Learning the Super-Resolution Space with Normalizing Flow. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part V*, volume 12350 of *Lecture Notes in Computer Science*, 715–732. Springer.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Gool, L. V. 2022. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 11451–11461. IEEE.
- Luo, Z.; Gustafsson, F. K.; Zhao, Z.; Sjölund, J.; and Schön, T. B. 2023. Image Restoration with Mean-Reverting Stochastic Differential Equations. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 23045–23066. PMLR.

- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.; and Ermon, S. 2022. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved Denoising Diffusion Probabilistic Models. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 8162–8171. PMLR.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.*, 115(3): 211–252.
- Saharia, C.; Chan, W.; Chang, H.; Lee, C. A.; Ho, J.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022a. Palette: Image-to-Image Diffusion Models. In Nandigjavi, M.; Mitra, N. J.; and Hertzmann, A., eds., *SIGGRAPH '22: Special Interest Group on Computer Graphics and Interactive Techniques Conference, Vancouver, BC, Canada, August 7 - 11, 2022*, 15:1–15:10. ACM.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, S. K. S.; Lopes, R. G.; Ayan, B. K.; Salimans, T.; Ho, J.; Fleet, D. J.; and Norouzi, M. 2022b. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In *NeurIPS*.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2021. Image Super-Resolution via Iterative Refinement. arXiv:2104.07636.
- Salimans, T.; and Ho, J. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Song, Y.; and Ermon, S. 2019. Generative Modeling by Estimating Gradients of the Data Distribution. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 11895–11907.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Vincent, P. 2011. A Connection Between Score Matching and Denoising Autoencoders. *Neural Comput.*, 23(7): 1661–1674.
- Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Loy, C. C. 2018. ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks. In Leal-Taixé, L.; and Roth, S., eds., *Computer Vision - ECCV 2018 Workshops - Munich, Germany, September 8-14, 2018, Proceedings, Part V*, volume 11133 of *Lecture Notes in Computer Science*, 63–79. Springer.
- Wang, Y.; Yu, J.; and Zhang, J. 2023. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018a. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 586–595. Computer Vision Foundation / IEEE Computer Society.
- Zhang, Y.; Li, K.; Li, K.; Wang, L.; Zhong, B.; and Fu, Y. 2018b. Image Super-Resolution Using Very Deep Residual Channel Attention Networks. In Ferrari, V.; Hebert, M.; Sminchisescu, C.; and Weiss, Y., eds., *European Conference on Computer Vision*, Lecture Notes in Computer Science, 294–310. Springer.
- Zhu, Y.; Zhang, K.; Liang, J.; Cao, J.; Wen, B.; Timofte, R.; and Gool, L. V. 2023. Denoising Diffusion Models for Plug-and-Play Image Restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023 - Workshops, Vancouver, BC, Canada, June 17-24, 2023*, 1219–1229. IEEE.