

AsyncDSB: Schedule-Asynchronous Diffusion Schrödinger Bridge for Image Inpainting

Zihao Han¹, Baoquan Zhang^{1*}, Lisai Zhang⁴, Shanshan Feng³, Kenghong Lin¹, Guotao Liang¹, Yunming Ye¹, Joeq², Kolaye²

¹ Harbin Institute of Technology, Shenzhen;

² ShenZhen SiFar Co., Ltd.;

³ Centre for Frontier AI Research;

⁴ Shengshu AI.

{zihaoan, linkenghong}@stu.hit.edu.cn, {baoquanzhang, victor_fengss, yeyunming}@hit.edu.cn, LisaiZhang@foxmail.com, lianggt@pcl.ac.cn, {joeqxc1974, kolaygm}@gmail.com

Abstract

Image inpainting is an important image generation task, which aims to restore corrupted image from partial visible area. Recently, diffusion Schrödinger bridge methods effectively tackle this task by modeling the translation between corrupted and target images as a diffusion Schrödinger bridge process along a noising schedule path. Although these methods have shown superior performance, in this paper, we find that 1) existing methods suffer from a schedule-restoration mismatching issue, *i.e.*, the theoretical schedule and practical restoration processes usually exist a large discrepancy, which theoretically results in the schedule not fully leveraged for restoring images; and 2) the key reason causing such issue is that the restoration process of all pixels are actually asynchronous but existing methods set a synchronous noise schedule to them, *i.e.*, all pixels shares the same noise schedule. To this end, we propose a schedule-Asynchronous Diffusion Schrödinger Bridge (AsyncDSB) for image inpainting. Our insight is preferentially scheduling pixels with high frequency (*i.e.*, large gradients) and then low frequency (*i.e.*, small gradients). Based on this insight, given a corrupted image, we first train a network to predict its gradient map in corrupted area. Then, we regard the predicted image gradient as prior and design a simple yet effective pixel-asynchronous noise schedule strategy to enhance the diffusion Schrödinger bridge. Thanks to the asynchronous schedule at pixels, the temporal interdependence of restoration process between pixels can be fully characterized for high-quality image inpainting. Experiments on real-world datasets show that our AsyncDSB achieves superior performance, especially on FID with around 3% ~ 14% improvement over state-of-the-art baseline methods.

Introduction

Image inpainting, also known as image completion, is a well-established low-level computer vision task. Its objective is to restore a corrupted image by leveraging information from the visible portion. Early studies primarily focus on non-learning inpainting methods to tackle this challenge, such as diffusion (Ballester et al. 2001; Bertalmio et al. 2000; Zomet 2003) and PatchMatch (Barnes et al. 2009; Bertalmio et al. 2003; Criminisi, Pérez, and Toyama 2003; Criminisi, Pérez, and

Toyama 2004). These methods infer the missing details in corrupted area by using statistical information from visible area, and achieve reasonable structure and texture for small region or the ungraded parts with well-defined geometry. However, these methods often struggle to recover large region effectively due to the lack of semantic inference ability.

To address the above challenge, deep learning inpainting methods are proposed and have received wide attention recently. These methods can be roughly divided into three categories, *i.e.*, generative adversarial networks (GAN) -based (Dong, Cao, and Fu 2022; Li et al. 2022b), diffusion models-based (Liu et al. 2024; Lugmayr et al. 2022; Xia et al. 2023), and diffusion bridge models-based (Li et al. 2023; Liu et al. 2023) approaches. Among them, the diffusion bridge models-based approaches (Li et al. 2023; Liu et al. 2023) have shown remarkable success due to its realistic and robust generation ability, especially diffusion Schrödinger bridge methods (a typical method is I²SB (Liu et al. 2023), which models the translation between corrupted and target images as a diffusion Schrödinger bridge and then learns it in a diffusion training manner.) Although I²SB has shown superior performance on image inpainting, we notice that their noise schedule β_t is only roughly set as a symmetric form (see Figure 1(a) with experience setting for all images and pixels. However, there is few works to analyze the working mechanism of such noise schedule and argue the rationality of such experience setting.

To this end, in this paper, we conduct an in-depth theoretical and qualitative analysis on the noise schedule β_t and find that: **1**) the noise schedule β_t actually controls the restoration speed of corrupted region (see theoretical analysis of Section). As shown in Figure 1, during the noise schedule, the main contents of an image are restored in few schedule steps between $t = 0.6T \rightarrow 0.4T$ since they have a large β_t value. Instead, the other steps close to $0T$ and $1.0T$ are only responsible for cold start and detail refinement. It is such three-stage rough-fine restoration schedule that makes existing I²SB method (Liu et al. 2023) achieve superior performance on image inpainting; **2**) However, the I²SB simply set the same schedule β_t for all pixels. The assumption behind such design is the three-stage restoration process of all pixels are pixel-synchronous. Unfortunately, the restoration process of pixels are not synchronous but actually asynchronous. For

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Connection between noise schedule (β_t) and the restoration process of I^2SB (Liu et al. 2023).

example, in general, the high-frequency structure pixels (e.g., face contour) are restored early than the low-frequency details pixels (e.g., face colors). Therefore, the synchronous noise schedule actually mismatch the asynchronous pixels restoration process; and **3**) Because of the above setting mismatch, the existing I^2SB method (Liu et al. 2023) suffers from a schedule-restoration mismatching issue, i.e., there is usually a large discrepancy between the theoretical denoise schedule and practical restoration process (see our insight of Section for more details). These facts mean that existing pixel-synchronous schedule used in existing I^2SB method (Liu et al. 2023) is actually not optimal, which limits its image inpainting performance.

To address this problem, we present a novel schedule-Asynchronous Schrödinger Diffusion Bridge (AsyncDSB) framework for image inpainting. The main idea is that introducing a pixel-asynchronous noise schedule for each pixel of corrupted images such that the sequential dependence of pixel restoration can be fully characterized. Intuitively, the pixels with high frequency information (i.e., large gradients) are restored early than low-frequency pixels (i.e., small gradients). Based on this insight, given a corrupted image, we first restore its gradient map (i.e., frequency component) through an adversarial learning manner. Then, we regard the image gradient as frequency priors and design a simple yet effective pixel-asynchronous noise schedule strategy for diffusion Schrödinger bridge. Thanks to the asynchronous schedule at pixels, the temporal interdependence between pixel restoration can be fully captured for high-quality image inpainting.

In summary, our contributions are three-fold as follows:

- We inspect the working mechanism of existing noise schedule β_t and find that applying the same noise schedule to all pixels results in a schedule-restoration mismatching issue, i.e., there is usually a large discrepancy between the theoretical noise schedule and practical restoration processes. This limits the performance of existing methods.
- To address the above issue, we propose a novel schedule-Asynchronous Schrödinger Diffusion Bridge (AsyncDSB) framework. It effectively aligns the theoretical schedule and practical restoration process such that better inpainting performance can be achieved.
- Extensive experiments conducted on two datasets (i.e., CelebA-HQ and Places2) demonstrate the superior performance of the proposed AsyncDSB method in comparison to various state-of-the-art image inpainting methods.

Related Work

Image Inpainting

Image inpainting is a challenging task, which aims restore a corrupted image with partial visible area. Recently, deep learning-based methods are proposed to tackle this challenge, which mainly contains three groups: (1) *GAN-based approaches* (Bendel, Ahmad, and Schniter 2023; Chu et al. 2023; Feng et al. 2022; Li et al. 2022a; Nazeri et al. 2019; Sargsyan et al. 2023; Wang et al. 2022; Zeng et al. 2021). This type of methods focuses on the encoder-decoder network to restore corrupted image and learn it through adversarial learning manner, such as MISF (Li et al. 2022b) and ZITS (Dong, Cao, and Fu 2022). (2) *Diffusion models-based approaches* (Liu et al. 2024; Lugmayr et al. 2022; Xia et al. 2023; Zhang et al. 2023). This group of methods aims to design a denoising diffusion probabilistic model with input-conditional UNet to transform image inpainting as a conditional generation task, such as StrDiffusion (Liu et al. 2024), Repaint (Lugmayr et al. 2022), DiffIR (Xia et al. 2023). (3) *Diffusion bridge models-based approaches* (De Bortoli et al. 2021; Li et al. 2023; Liu et al. 2023; Liu and Wu 2023; Peluchetti 2023; Shi et al. 2023; Su et al. 2022). This kind of approach is a new model of conditional diffusion, which directly models the translation between corrupted image and target image in a diffusion bridge manner (Li et al. 2023; Liu et al. 2023). For example, in (Liu et al. 2023), Liu et al. propose a nonlinear diffusion bridge, i.e., diffusion Schrödinger bridge (called I^2SB), to model the translation between corrupted and target images. The non-linear characteristic of diffusion Schrödinger bridge significantly improves image inpainting performance.

In this paper, we focus on I^2SB due to its superiority but find that 1) it suffers from a schedule-restoration mismatching issue (see Section); and 2) its reason is pixel-synchronous noise schedule setting. To this end, we present a novel pixel-asynchronous schedule strategy, which effectively aligns the theoretical schedule and practical restoration process for better inpainting performance.

Noise Schedule in Diffusion Models

Noise schedule β_t is an important hyper-parameter in diffusion models, which determines how much weight we assign to each denoising step. As existing noise schedules (e.g., linear-based (Ho, Jain, and Abbeel 2020) and cosine-based (Nichol and Dhariwal 2021) schedule) are only empirical settings, some studies raised a questionable point of view on

noise schedule recently and propose some novel noise schedules for diffusion models, such as re-scaling schedule (Lin et al. 2024), WSNR-equivalent schedule (Guo et al. 2023), compound schedule (Chen 2023), and optimization-based schedule (Strasman et al. 2024). However, to our best knowledge, there is few works to explore it for diffusion bridges, especially I²SB. In this paper, we take an in-depth analysis on noise schedule and propose a novel pixel-asynchronous schedule, which effectively improves the inpainting performance of I²SB.

Methodology

Problem Definition

Formally, given a target image x_g , we assume that the target image x_g is corrupted to a corrupted image x_c along an image mask x_m where $x_m^{i,j} = 1$ if pixel $x_c^{i,j}$ at spatial position (i, j) is corrupted region otherwise $x_m^{i,j} = 0$, i.e., $x_c = (1 - x_m) \times x_g$. The goal of image inpainting is to restore the missing area (i.e., $x_m \times x_g$) within the corrupted image x_c . In particular, the restored areas should blend seamlessly with the surrounding visible region and be logically consistent.

Preliminaries: Diffusion Schrödinger Bridge for Image Inpainting

Recently, diffusion Schrödinger bridge models (DSB), as a new conditional diffusion models, has achieve success on image inpainting. Different from previous input-conditional diffusion models, DSB model directly learns the transformation processes between target image distributions p_{x_g} and corrupted image distribution p_{x_c} via a nonlinear Schrödinger bridge. Its key challenge is how learn the Schrödinger bridge on image inpainting. Next, we briefly review the Schrödinger bridge and its adaptation on image inpainting, i.e., Diffusion Schrödinger Bridge (Liu et al. 2023).

Schrödinger Bridge (SB) is an entropy-regularized optimal transport problem, which has been widely applied to quantum mechanics (Pavon 2002) and optimal transport and control (Chen, Georgiou, and Pavon 2016). Given two arbitrary distributions $p(x_0)$ and $p(x_1)$, SB aims to seek an optimal path pair between $p(x_0)$ and $p(x_1)$ with the following forward and backward stochastic differential equations (SDEs):

$$\text{Forward : } dx_t = [f_t + \beta_t \nabla \log \Psi(x_t, t)]dt + \sqrt{\beta_t}dW_t, \quad t \in [0, 1], \quad (1)$$

$$\text{Backward : } dx_t = [f_t - \beta_t \nabla \log \hat{\Psi}(x_t, t)]dt + \sqrt{\beta_t}d\hat{W}_t, \quad t \in [0, 1],$$

where the forward and backward SDEs denote the optimal transformation from $t = 0$ to $t = 1$ and its reverse from $t = 1$ to $t = 0$, respectively; $x_0 \sim p_{x_0}$ and $x_1 \sim p_{x_1}$ denote samples from two arbitrary distributions; W_t and \hat{W}_t denote the standard Brownian motion and its reversed counterpart, respectively; f_t is a drift term; and $\Psi(\cdot)$ and $\hat{\Psi}(\cdot)$ denotes the time-varying energy functions.

Diffusion Schrödinger Bridge In fact, the image inpainting process can also be regarded as a Schrödinger bridge when we set the target image distributions p_{x_g} and its corrupted image distributions p_{x_c} as p_{x_0} and p_{x_1} , respectively, i.e., $p_{x_0} = p_{x_g}$ and $p_{x_1} = p_{x_c}$. However, due to the complexity of SB, the key challenge of adapting SB to image

inpainting is how to solve the SDEs defined in Eq (1). To address this challenge, Liu et al. (Liu et al. 2023) propose a tractable diffusion-based Schrödinger bridge (called I²SB (Liu et al. 2023)) recently, which simplifies the SDEs of Eq (1) by dropping $f := 0$ and deduce an analytic solution for forward SDEs with the following Gaussian posterior:

$$x_t \sim q(x_t|x_0, x_1) = N(x_t; \mu_t, \Sigma_t),$$

$$\mu_t = \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \sigma_t^2}x_0 + \frac{\sigma_t^2}{\bar{\sigma}_t^2 + \sigma_t^2}x_1, \Sigma_t = \frac{\sigma_t^2 \bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \sigma_t^2} \cdot I, \quad (2)$$

where σ_t and $\bar{\sigma}_t$ are two variances accumulated from either sides along a noise schedule β_t , i.e.,

$$\sigma_t^2 = \int_0^t \beta_\tau d\tau, \quad \bar{\sigma}_t^2 = \int_t^1 \beta_\tau d\tau, \quad (3)$$

where the noise schedule β_τ is experimentally set to a simple symmetric form (see Figure 1(a)). Based on the above analytic solution, the score function for $\nabla \log \hat{\Psi}(x_t, t)$ can be directly calculated, i.e., $\nabla \log \hat{\Psi}(x_t, t) = \frac{x_t - x_0}{\sigma_t^2}$, which is further used for learning a score function $s_\theta(x_t, t)$. Following standard diffusion (i.e., DDPM (Ho, Jain, and Abbeel 2020)), the overall score-matching objective L can be defined as:

$$L = \|s_\theta(x_t, t) - \frac{x_t - x_0}{\sigma_t}\|. \quad (4)$$

After training (i.e., during inference), given a corrupted image, i.e., $x_1 = x_c$, we can directly traverse the restoration process induced by Eq (1) by the following recursive posterior sampling in DDPM:

$$q(x_{t-\frac{1}{T}}|x_t) = \mathcal{N}\left(x_t - \int_{t-\frac{1}{T}}^t [-\beta_t \nabla \log \hat{\Psi}(x_t, t)] dt, \int_{t-\frac{1}{T}}^t \beta_t dt\right), t = 1, 1 - \frac{1}{T}, \dots, 0, \quad (5)$$

where $\nabla \log \hat{\Psi}(x_t, t) = \frac{s_\theta(x_t, t)}{\sigma_t}$ and T denotes the total time step of discretized SDE. Following (Liu et al. 2023), $T = 1000$ is used. Finally, a restored image, i.e., $\hat{x}_g = \hat{x}_0$, can be achieved at time step $t = 0$.

A Closer Look at Existing Noise Schedule β_t

Although existing I²SB (Liu et al. 2023) has shown superior performance, we find that 1) existing method focuses on experience manner to set the noise schedule β_t defined in Eq (3); and 2) as shown in Figure 1(a), the noise schedule β_t is usually set as a symmetric form for all images and pixels. However, few works analyze the working mechanism of such design and argue the rationality of such experience setting. In this paper, we conduct an in-depth theoretical and qualitative analysis to clarify this point.

Theoretical Analysis on Noise Schedule Intuitively, the noise schedule β_t may control the restoration process (i.e., speed) of image inpainting along the diffusion Schrödinger bridge. To verify this point, in this part, we conduct a theoretical analysis on the transformation path of diffusion Schrödinger bridge. Formally, given a pair corrupted and target samples (x_c, x_g) . Then, we set the pair sample as the two endpoints of SB, i.e., $x_0 = x_g$ and $x_1 = x_c$. After that,

following the Eq (2), we can represent each restored image x_t at time step t on the transport path as:

$$x_t = \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \sigma_t^2} x_0 + \frac{\sigma_t^2}{\bar{\sigma}_t^2 + \sigma_t^2} x_1 + \frac{\bar{\sigma}_t \sigma_t}{\sqrt{\bar{\sigma}_t^2 + \sigma_t^2}} z, \quad (6)$$

where z denotes a Gaussian noise, *i.e.*, $z \sim \mathcal{N}(0, 1)$. Then, the restoration speed of image inpainting process at time t can be characterized as a derivative of x_t about the time t . That is,

$$\begin{aligned} \frac{\partial x_t}{\partial t} &= \frac{\partial \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \sigma_t^2}}{\partial t} x_0 + \frac{\partial \frac{\sigma_t^2}{\bar{\sigma}_t^2 + \sigma_t^2}}{\partial t} x_1 + \underbrace{\frac{\partial \frac{\bar{\sigma}_t \sigma_t}{\sqrt{\bar{\sigma}_t^2 + \sigma_t^2}}}{\partial t}}_{Term3} z \\ &= \frac{\sigma_t^2 \frac{\partial \bar{\sigma}_t^2}{\partial t} - \frac{\partial \sigma_t^2}{\partial t} \bar{\sigma}_t^2}{(\bar{\sigma}_t^2 + \sigma_t^2)^2} x_0 + \frac{\bar{\sigma}_t^2 \frac{\partial \sigma_t^2}{\partial t} - \frac{\partial \bar{\sigma}_t^2}{\partial t} \sigma_t^2}{(\bar{\sigma}_t^2 + \sigma_t^2)^2} x_1 \\ &= \frac{(x_1 - x_0)(\beta_t \bar{\sigma}_t^2 + \sigma_t^2 \beta_t)}{(\bar{\sigma}_t^2 + \sigma_t^2)^2} = \beta_t \frac{(x_1 - x_0)}{\bar{\sigma}_t^2 + \sigma_t^2}, \end{aligned} \quad (7)$$

where *Term3* can be overlooked since its expectation is zero. In particular, the corrupted image x_1 satisfies $x_1 = x_0 \times (1 - x_m)$, which suggests that the derivative $\frac{\partial x_t}{\partial t}$ can be further expressed as:

$$\frac{\partial x_t}{\partial t} = \beta_t \frac{x_0 \times (1 - x_m) - x_0}{\bar{\sigma}_t^2 + \sigma_t^2} = -\beta_t \frac{x_0 x_m}{\bar{\sigma}_t^2 + \sigma_t^2}. \quad (8)$$

Note that 1) x_m denotes the mask of corrupted region, *i.e.*, $x_m^{i,j} = 1$ if pixel $x_1^{i,j}$ at spatial position (i, j) is corrupted region otherwise $x_m = 0$, and 2) $\bar{\sigma}_t^2 + \sigma_t^2 = \int_0^1 \beta_\tau d\tau$ is actually a constant and can be overlooked. Let's take a closer look at the Eq (8), we can see that the restoration speed $\frac{\partial x_t}{\partial t}$ of image inpainting process is only relevant with the noise schedule β_t and the value $x_0^{i,j}$ of each pixel (i, j) at the corrupted region. In particular, the noise schedule β_t controls the overall restoration tread and the pixel value $x_0^{i,j}$ controls the amplitude of restoration at corrupted region.

Based on this fact, let's revisit the the noise schedule β_t shown in Figure 1(a), we can see that β_t is symmetric and two bound points are small but middle point is big. This means that the restoration process at two bound points are very slow (a reasonable explain is that the early slow speed is to achieve a good cold start, however the later is to achieve detail refinement) and only few intermediate schedule steps are used for quickly restoring images. It is such nonlinear noise schedule with well cold start and detail refinement that I²SB achieves superior performance on image inpainting.

Our Insight: Schedule-restoration Mismatching Issue

The Eq (8) is only theoretical analysis result, a nature question is *Is the practical restoration process really like theoretical schedule?*. To answer this question, we evaluate the SSIM performance of CelebA-HQ test set at each denoising time step t from $t = 1.0T$ to $t = 0T$ ($T=1000$) and calculate its derivative. For clarity, we normalize its amplitude for removing the impact of pixel values.

The result is shown in Figure 2. From Figure 2, we can see that 1) the shape of restoration process is very similar to the




Figure 2: Qualitative Analysis.

theoretical schedule; but 2) the overall tread has been delayed by a large margin. This means that existing I²SB methods suffers from a schedule-restoration mismatching issue, *i.e.*, there is a large discrepancy between theoretical schedule (see blue line) and practical restoration process (see blue dotted line). Such mismatching issue will results in that the theoretical nonlinear schedule is not fully leveraged on overall schedule time steps such that the inpainting performance of existing I²SB method is limited.




Figure 3: Comparison between theoretical schedule and practical process on different frequency.

We guess the reason of causing such schedule-restoration mismatching issue is the existing I²SB method sets all pixels to the same noise schedule (*i.e.*, its assumption is the restoration process of all pixels are synchronous). However, intuitively, the generation process are not synchronous but exist a potential temporal interdependence where some high-frequency pixels (*e.g.*, outline or texture) is usually restored first and then some lower frequency pixels is repaired in later. To verify our guess, in Figure 3, we evaluate the restoration process (*i.e.*, the SSIM derivative) of CelebA-HQ test set at three type of pixels with different frequency (*i.e.*, high-frequency pixels, middle-frequency pixels, and lower-frequency pixels), respectively. From Figure 3, we can see that 1) these pixels with different frequency are indeed not restored at the same time but asynchronous; and 2) the restoration process of high-frequency pixels is very consist to the theoretical schedule, but middle- and lower-frequency pixels brings a large discrepancy which results in the overall restoration process to be delayed (*i.e.*, causing the overall mismatching between the theoretical schedule and the practical restoration process). This fact verify our guess, *i.e.*, the schedule-restoration mismatching issue is indeed caused by that the restoration process of all pixels are actually asynchronous but existing I²SB method uses a pixel-synchronous noise schedule β_t .

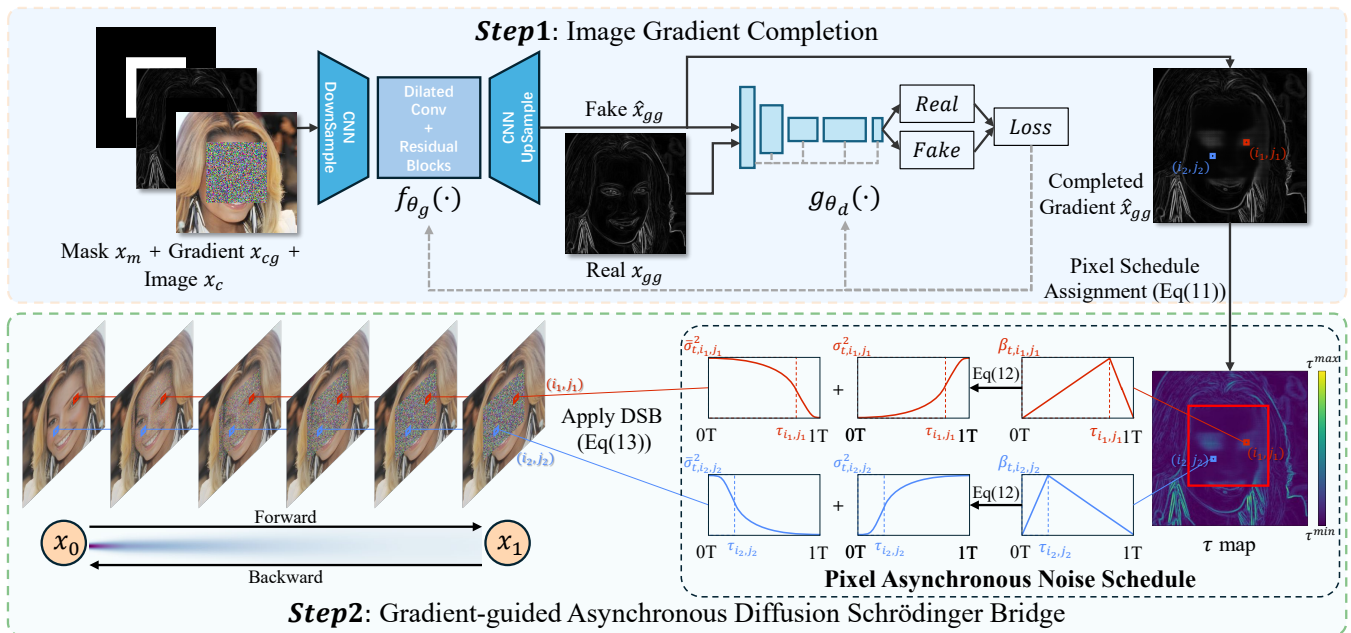


Figure 4: The overall framework of our AsyncDSB, which consists of two steps, an image gradient completion step, and a gradient-guided asynchronous Schrödinger diffusion bridge step. Here, given a corrupted image, the *Step 1* aims to predict the image gradient of the missing area. Then, the *Step 2* accounts for assigning a pixel-asynchronous noise schedule, which is employed to make a asynchronous Schrödinger diffusion bridge for completing the corrupted region in origin pixel space according to the image gradient information obtained by the *Step 1*.

AsyncDSB Framework

As introduced in Section , existing noise schedule methods suffers from a schedule-restoration mismatching issue, which is because that the generation process of different pixels are asynchronous but existing I²SB method employs a synchronous noise schedule β_t . According to the above mentioned analysis, an intuitive idea is employing a non-symmetric noise schedule to match the asynchronous generation according to pixel’s frequency information, *i.e.*, the pixels with high-frequency information (*e.g.*, contour) is preferentially scheduled and then the pixels with low-frequency information. To achieve this idea, we propose a schedule-Asynchronous Diffusion Schrödinger Bridge (AsyncDSB) framework. As shown in Figure 4, our AsyncDSB framework consists of two steps, *i.e.*, image gradient completion and gradient-guided asynchronous Schrödinger diffusion bridge.

It is worth noting that the AsyncSDB framework is not limited to the image frequency guided schedule strategy. The proposed idea of asynchronous schedule actually enables any $\beta_{t,i,j}$ strategy that can match the pixel-asynchronous image generation process, which is highly potential. In this paper, we use the image frequency as guidance mainly to verify the correctness of the AsyncSDB framework.

Step 1: Image Gradient Completion. Intuitively, a pixel with a larger gradient has a larger frequency. To achieve the frequency prior of a corrupted image x_c , we first take the corrupted image x_c and its mask x_m and gradient map x_{cg} as inputs and then design an image gradient completion module $f_{\theta_g}(\cdot)$ with parameter θ_g . During inference, given a corrupted pixels x_c and its mask x_m and gradient map x_{cg} , the module

$f_{\theta_g}(\cdot)$ is directly used to predict its target gradient map \hat{x}_{gg} . That is:

$$\hat{x}_{gg} = f_{\theta_g}(x_c, x_m, x_{cg}), \quad (9)$$

and then the gradient could be used as guidance of the asynchronous $\beta_{t,i,j}$ for each pixel.

During training, following edge prediction (Nazeri et al. 2019), we adopt the adversarial learning manner to learn the image gradient completion module $f_{\theta_g}(\cdot)$. The adversarial training loss is formulated as:

$$\min_{\theta_g} \max_{\theta_d} L_G = (\lambda_{adv} \max_{\theta_d} (L_{adv}) + \lambda_{FM} L_{FM}), \quad (10)$$


where λ_{adv} and λ_{FM} are hyper-parameters; L_{adv} and L_{FM} denotes the adversarial and feature matching loss(Nazeri et al. 2019), respectively; and θ_d denotes the parameters of discriminator $g_{\theta_d}(\cdot)$.

Step 2: Gradient-guided Asynchronous Diffusion Schrödinger Bridge.

As analyzed in Section, high-frequency pixels should restore earlier than pixels with low frequency. To this end, we introduce a new variable $\tau_{i,j} \in [0T, 1.0T]$ and transform the traditional noise schedule β_t into an asymmetric and pixel-wise noise schedule form, *i.e.*, $\beta_{t,i,j}$ (see bottom right part of Figure 4 for its details). The advantage of such design is that the schedule sequence between pixels can be freely adjusted by assigning different variable $\tau_{i,j}$. Intuitively, the pixel with high frequency should assign larger $\tau_{i,j}$ than pixel with low frequency since the pixel with high frequency are priority restored. To achieve this idea, we design a simple yet

Method	CelebA-HQ								Places2							
	FID↓				IS↑				FID↓				IS↑			
	center	half	wide	narrow	center	half	wide	narrow	center	half	wide	narrow	center	half	wide	narrow
MISF (Li et al. 2022b)	3.0	11.9	3.7	2.3	4.23	6.21	8.69	11.12	14.2	26.4	7.7	4.5	61.3	37.3	65.8	77.3
ZITS (Dong, Cao, and Fu 2022)	4.6	10.9	5.4	3.2	5.28	7.24	8.63	10.32	<u>8.3</u>	11.7	<u>5.4</u>	3.2	<u>64.9</u>	51.6	69.8	79.6
StrDiffusion (Liu et al. 2024)	6.7	17.3	8.3	3.7	5.46	7.08	8.83	11.10	16.0	29.6	8.5	3.7	57.1	33.9	63.9	78.3
RePaint (Lugmayr et al. 2022)	9.0	12.1	7.1	7.3	6.40	8.25	<u>9.61</u>	<u>11.31</u>	9.0	13.0	5.5	3.4	64.2	49.12	71.1	79.6
DiffIR _{S2} (Xia et al. 2023)	2.2	9.7	3.4	3.7	5.86	8.34	9.60	11.27	8.8	11.1	5.7	3.7	64.9	<u>51.6</u>	<u>70.2</u>	<u>80.0</u>
BBDM (Li et al. 2023)	6.9	8.5	7.4	7.5	6.00	6.43	6.46	6.42	28.5	32.6	30.6	27.1	31.8	26.1	26.8	33.9
I ² SB (baseline) (Liu et al. 2023)	<u>2.2</u>	<u>5.4</u>	<u>2.9</u>	<u>1.7</u>	<u>6.71</u>	<u>8.35</u>	9.50	11.30	8.6	<u>10.8</u>	5.6	<u>3.1</u>	64.6	43.3	69.8	79.7
AsyncDSB	1.9	5.2	2.8	1.6	6.74	8.38	9.62	11.41	8.3	10.5	5.2	3.0	65.1	52.3	69.9	80.0
	▲14%	▲4%	▲3%	▲6%	▲4%	▲4%	▲12%	▲9%	▲4%	▲3%	▲7%	▲3%	▲8%	▲14%	▲1%	▲1%

Table 1: Quantitative results on CelebA-HQ and Places2. ↓: Lower is better for FID; ↑: Higher is better for IS. The best and second-best results are marked with **boldface** and underline, respectively.



(a) Regular masks (i.e., center and half) (b) Irregular masks (i.e., wide and narrow)

Figure 5: Qualitative results on CelebA-HQ and Places2 under regular and irregular masks.

effective pixel-wise schedule assignment strategy. That is,

$$\tau_{i,j} = \frac{\text{gauss}(\hat{x}_{i,j}^{gg}) - \min(\text{gauss}(\hat{x}_{gg}))}{\max(\text{gauss}(\hat{x}_{gg}) - \min(\text{gauss}(\hat{x}_{gg}))} (\tau^{\max} - \tau^{\min}) + \tau^{\min} \quad (11)$$

where τ^{\max} and τ^{\min} are two hyper-parameters, which controls the asynchronous range of pixel schedule; and $\text{gauss}()$ is a Gaussian filter, which aims to achieve a flat noise schedule for all pixels.

After assigning noise schedule for each pixel, we follow Eq (3) to calculate two variances $\sigma_{t,i,j}$ and $\bar{\sigma}_{t,i,j}$ for each pixel by accumulating from either sides along a noise schedule $\beta_{t,i,j}$. That is,

$$\sigma_{t,i,j}^2 = \int_0^t \beta_{\tau,i,j} d\tau, \quad \bar{\sigma}_{t,i,j}^2 = \int_t^1 \beta_{\tau,i,j} d\tau. \quad (12)$$

Finally, during inference, we directly use the pixel form of Eq (5) to restore images. That is,

$$q(x_{t-\frac{1}{\sigma_t},i,j} | x_{t,i,j}) = \mathcal{N} \left(x_{t,i,j} - \int_{t-\frac{1}{\sigma_t}}^t [-\beta_{t,i,j} \nabla \log \hat{\Psi}(x_{t,i,j}, t)] dt, \int_{t-\frac{1}{\sigma_t}}^t \beta_{t,i,j} dt \right), \quad (13)$$

where $\nabla \log \hat{\Psi}(x_{t,i,j}, t) = \frac{s_{\theta}(x_{t,i,j}, t)}{\sigma_t}$. Please see Appendix for more training and inference details.

Experiments

Experimental Settings

Dataset. We validate our method and various baselines over two typical dataset: **1) CelebA-HQ** is a face dataset. We split this dataset by following standard dataset split (Liu et al. 2015). **2) Places2** is a scene image dataset, which split into 1.8 million training images and 36500 validation images,

from which we randomly select 6000 images as test. Following RePaint (Lugmayr et al. 2022), we use four types of masks (i.e., center, half, wide and narrow) to conduct experiments where all images are set to 256x256 resolution.

Evaluation Metrics. Following (Liu et al. 2023), we employ two widely used metrics, Fréchet Inception Distance (FID) and Inception Score (IS), to evaluate the quality and diversity of the generated images.

Results and Discussion

We select seven latest methods: MISF (Li et al. 2022b), ZITS (Dong, Cao, and Fu 2022), StrDiffusion (Liu et al. 2024), RePaint (Lugmayr et al. 2022), DiffIR_{S2} (Xia et al. 2023), BBDM (Li et al. 2023), and I²SB (Liu et al. 2023) as our baseline. The results are shown in Table 1 and Figure 5.

Quantitative Evaluation. In Table 1, we evaluate all methods on CelebA-HQ and Places2 under varied mask settings. We can see that 1) our AsyncDSB achieves the best performance, especially on FID, which verifies our superiority; 2) our AsyncDSB exceeds the key competitor (i.e., I²SB) around 3% ~ 14%. This means that employing a pixel-asynchronous noise schedule is very beneficial for I²SB, which verifies the effectiveness of the proposed AsyncDSB; 3) the performance improvement is more significant on center&half masks than wide&narrow, which may be because the missing of high-frequency information (e.g., outline) is even worse on center&half masks.

Qualitative Evaluation. In Figure. 5, we show some visual results of our AsyncDSB and some representative methods. It can be observed that our AsyncDSB generates images with finer details and more natural textures that are well blended with the context. This further verifies the effectiveness.

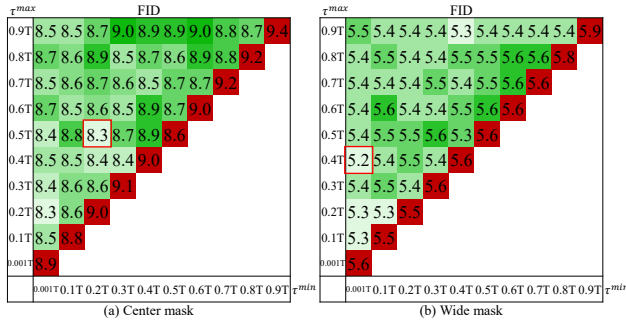


Figure 6: Analysis of τ^{min} and τ^{max} on Places2.




Figure 7: Comparison between theoretical schedule and practical process on different frequency.

Ablation Study

Hyper-parameters (τ^{min} and τ^{max}) Setting In Figure 6, we conduct a parameter analysis on Places 2 from $\tau^{min} = 0.001T$ to $\tau^{min} = 0.9T$ and $\tau^{max} = 0.001T$ to $\tau^{min} = 0.9T$. In particular, the blank value in Figure 6 is not valid since two parameters must satisfy $\tau^{min} \leq \tau^{max}$. From results, we can see that 1) compared with the setting of $\tau^{min} = \tau^{max}$ (*i.e.*, pixel-synchronous schedule), our AsyncDSB achieves more superior performance when τ^{min} is set to a smaller value, which further verifies the effectiveness of our pixel-asynchronous schedule; 2) our AsyncDSB achieves best performance when we set (τ^{min}, τ^{max}) to $(0.2T, 0.5T)$ for the regular masks like center and $(0.001T, 0.4T)$ for the irregular masks like wide. We guess that this phenomenon may be caused by the loss of different frequencies caused under different masks.

Hyper-parameter (τ^{min} and τ^{max}) and Mask Ratio Intuitively, a large mask ratio is easier to cause the loss of high-frequency information since most pixels are low-frequency. To this end, in Table 2, we report the optimal parameters for different mask ratio setting. We find that 1) large mask ratios tend to the large τ_{min} and τ_{max} , *i.e.*, tend to set earlier schedule for each pixel; 2) the performance achieve consistent improvement. This means exploring a mask-and-sample-specific pixel-asynchronous noise schedule is promising for enhancing inpainting performance. We will explore it in future work.

Pixel-asynchronous Schedule v.s. Synchronous Schedule

In Table 3, we add the image gradient into I^2SB (*i.e.*, regarding as UNet condition) as baseline, which can be regarded as a gradient-guided pixel-synchronous I^2SB and report its performance on CelebA-HQ. We find that our AsyncDSB exceeds the baseline by a large margin, which means that our pixel-asynchronous schedule is effective.

Can AsyncDSB solve the schedule-restoration mismatch

Mask Ratios	hyper-parameters		Improved FID
	τ_{min}	τ_{max}	
1% ~ 10%	0.001T	0.9T	7.6→7.5
11% ~ 20%	0.2T	0.9T	13.1→13.0
21% ~ 30%	0.4T	0.9T	17.6→17.5
31% ~ 40%	0.7T	0.9T	26.2→26.1

Table 2: Analysis of mask ratios on Places2.

Settings	FID↓			
	center	half	wide	narrow
I^2SB +gradient	2.0	5.8	3.0	1.7
Our AsyncDSB	1.9	5.2	2.8	1.6

Table 3: Analysis on asynchronous schedule.




Figure 8: Sampling process of our AsyncDSB.

issue? In Figure 7, we compare the restoration process (*i.e.*, the SSIM derivative) between I^2SB and our AsyncDSB on CelebA-HQ test set at different frequency pixels. From results, we can see that the schedule-restoration mismatch issue can be significantly alleviated after applying our AsyncDSB, which verifies our effectiveness.

How does our AsyncDSB work? In Figure 8, we random select two test image from CelebA-HQ and visualize its completed image gradient and restoration process. From Figure 8, we can see that 1) our AsyncDSB well competed the missing gradient map and 2) compared with pixel-synchronous schedule (*i.e.*, I^2SB), in our AsyncDSB, some high-frequency content (see the outline of hat at $t = 0.66T$) indeed do recover earlier than low-frequency details (see the face color at $t = 0.33T$).

Conclusion

In this paper, we find that existing diffusion Schrödinger bridge method suffers from a schedule-restoration mismatch issue, which is caused by setting all pixels to a synchronous noise schedule. To this end, we propose a schedule-asynchronous diffusion Schrödinger bridge for image inpainting. Its advantage is resorting to the asynchronous schedule at pixels, the temporal interdependence between pixels can be fully characterized for high-quality image inpainting. Experimental results show that our method achieves superior performance over previous state-of-the-art methods.

Acknowledgments

This work was supported by the Shenzhen Peacock Program under Grant No. ZX20230597, Stabilization Support Program (ZX20240460), NSFC under Grant No. 62272130 and Grant No. 62376072, and the Shenzhen Science and Technology Program under Grant No. KCXFZ20211020163403005.

References

- Ballester, C.; Bertalmio, M.; Caselles, V.; Sapiro, G.; and Verdera, J. 2001. Filling-in by joint interpolation of vector fields and gray levels. *IEEE transactions on image processing*, 10(8): 1200–1211.
- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3): 24.
- Bendel, M.; Ahmad, R.; and Schniter, P. 2023. A Regularized Conditional GAN for Posterior Sampling in Image Recovery Problems. *Advances in Neural Information Processing Systems*, 36.
- Bertalmio, M.; Sapiro, G.; Caselles, V.; and Ballester, C. 2000. Image inpainting. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 417–424.
- Bertalmio, M.; Vese, L.; Sapiro, G.; and Osher, S. 2003. Simultaneous structure and texture image inpainting. *IEEE transactions on image processing*, 12(8): 882–889.
- Chen, T. 2023. On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972*.
- Chen, Y.; Georgiou, T. T.; and Pavon, M. 2016. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169: 671–691.
- Chu, T.; Chen, J.; Sun, J.; Lian, S.; Wang, Z.; Zuo, Z.; Zhao, L.; Xing, W.; and Lu, D. 2023. Rethinking fast fourier convolution in image inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23195–23205.
- Criminisi, A.; Pérez, P.; and Toyama, K. 2003. Object removal by exemplar-based inpainting. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, II–II. IEEE.
- Criminisi, A.; Pérez, P.; and Toyama, K. 2004. Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on image processing*, 13(9): 1200–1212.
- De Bortoli, V.; Thornton, J.; Heng, J.; and Doucet, A. 2021. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34: 17695–17709.
- Dong, Q.; Cao, C.; and Fu, Y. 2022. Incremental transformer structure enhanced image inpainting with masking positional encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11358–11368.
- Feng, T.; Feng, W.; Li, W.; and Lin, D. 2022. Cross-image context for single image inpainting. *Advances in Neural Information Processing Systems*, 35: 1474–1487.
- Guo, Q.; Liu, S.; Yu, Y.; and Luo, P. 2023. Rethinking the Noise Schedule of Diffusion-Based Generative Models.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Li, B.; Xue, K.; Liu, B.; and Lai, Y.-K. 2023. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, 1952–1961.
- Li, W.; Lin, Z.; Zhou, K.; Qi, L.; Wang, Y.; and Jia, J. 2022a. Mat: Mask-aware transformer for large hole image inpainting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10758–10768.
- Li, X.; Guo, Q.; Lin, D.; Li, P.; Feng, W.; and Wang, S. 2022b. Misf: Multi-level interactive siamese filtering for high-fidelity image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1869–1878.
- Lin, S.; Liu, B.; Li, J.; and Yang, X. 2024. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5404–5411.
- Liu, G.-H.; Vahdat, A.; Huang, D.-A.; Theodorou, E. A.; Nie, W.; and Anandkumar, A. 2023. I2SB: image-to-image Schrödinger bridge. In *Proceedings of the 40th International Conference on Machine Learning*, 22042–22062.
- Liu, H.; Wang, Y.; Qian, B.; Wang, M.; and Rui, Y. 2024. Structure Matters: Tackling the Semantic Discrepancy in Diffusion Models for Image Inpainting. *arXiv preprint arXiv:2403.19898*.
- Liu, X.; and Wu, L. 2023. Learning diffusion bridges on constrained domains. In *international conference on learning representations (ICLR)*.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11461–11471.
- Nazeri, K.; Ng, E.; Joseph, T.; Qureshi, F. Z.; and Ebrahimi, M. 2019. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Pavon, M. 2002. Quantum Schrödinger bridges. In *Directions in mathematical systems theory and optimization*, 227–238. Springer.
- Peluchetti, S. 2023. Diffusion bridge mixture transports, Schrödinger bridge problems and generative modeling. *Journal of Machine Learning Research*, 24(374): 1–51.
- Sargsyan, A.; Navasardyan, S.; Xu, X.; and Shi, H. 2023. Mi-gan: A simple baseline for image inpainting on mobile devices. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7335–7345.
- Shi, Y.; De Bortoli, V.; Campbell, A.; and Doucet, A. 2023. Diffusion Schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36.
- Strasman, S.; Ocello, A.; Boyer, C.; Corff, S. L.; and Lemaire, V. 2024. An analysis of the noise schedule for score-based generative models. *arXiv preprint arXiv:2402.04650*.

Su, X.; Song, J.; Meng, C.; and Ermon, S. 2022. Dual Diffusion Implicit Bridges for Image-to-Image Translation. In *The Eleventh International Conference on Learning Representations*.

Wang, W.; Niu, L.; Zhang, J.; Yang, X.; and Zhang, L. 2022. Dual-path image inpainting with auxiliary gan inversion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11421–11430.

Xia, B.; Zhang, Y.; Wang, S.; Wang, Y.; Wu, X.; Tian, Y.; Yang, W.; and Van Gool, L. 2023. Diffir: Efficient diffusion model for image restoration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13095–13105.

Zeng, Y.; Lin, Z.; Lu, H.; and Patel, V. M. 2021. Cr-fill: Generative image inpainting with auxiliary contextual reconstruction. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14164–14173.

Zhang, G.; Ji, J.; Zhang, Y.; Yu, M.; Jaakkola, T. S.; and Chang, S. 2023. Towards coherent image inpainting using denoising diffusion implicit models.

Zomet. 2003. Learning how to inpaint from global image statistics. In *Proceedings Ninth IEEE international conference on computer vision*, 305–312. IEEE.