



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/239024/>

Version: Accepted Version

Article:

Shamsolmoali, Pourya, Zareapoor, Masoumeh, Zhou, Huiyu et al. (2025) From Missing Pieces to Masterpieces: Image Completion with Context-Adaptive Diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence. ISSN: 0162-8828

<https://doi.org/10.1109/TPAMI.2025.3558092>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

From Missing Pieces to Masterpieces: Image Completion with Context-Adaptive Diffusion

Pourya Shamsolmoali, *Senior Member, IEEE*, Masoumeh Zareapoor, *Member, IEEE*, Huiyu Zhou, Michael Felsberg, *Senior Member, IEEE*, Dacheng Tao, *Fellow, IEEE*, and Xuelong Li, *Fellow, IEEE*

Abstract—Image completion is a challenging task, particularly when ensuring that generated content seamlessly integrates with existing parts of an image. While recent diffusion models have shown promise, they often struggle with maintaining coherence between known and unknown (missing) regions. This issue arises from the lack of explicit spatial and semantic alignment during the diffusion process, resulting in content that does not smoothly integrate with the original image. Additionally, diffusion models typically rely on global learned distributions rather than localized features, leading to inconsistencies between the generated and existing image parts. In this work, we propose ConFill, a novel framework that introduces a Context-Adaptive Discrepancy (CAD) model to ensure that intermediate distributions of known and unknown regions are closely aligned throughout the diffusion process. By incorporating CAD, our model progressively reduces discrepancies between generated and original images at each diffusion step, leading to contextually aligned completion. Moreover, ConFill uses a new Dynamic Sampling mechanism that adaptively increases the sampling rate in regions with high reconstruction complexity. This approach enables precise adjustments, enhancing detail and integration in restored areas. Extensive experiments demonstrate that ConFill outperforms current methods, setting a new benchmark in image completion.

Index Terms—Image completion, diffusion models, context-adaptive discrepancy

1 INTRODUCTION

IMAGE completion, also called image inpainting, is an important task in computer vision and image processing that aims to reconstruct missing or corrupted portions of an image such that the restored areas are both visually pleasing and coherent. This technique is widely used in various domains, including image editing [1], [2], image restoration [3], [4], and the removal of unwanted objects [5]. At present, several state-of-the-art methods use approaches such as GANs [6], [7], [8], transformers [9], [10], and autoregressive modeling [11], [12]. Despite their advancements, these techniques often experience unstable training [12] and struggle with filling large missing regions in images [2], [8]. Recently, there has been a growing interest among researchers in using diffusion models [13], [14] to address the above issue. These generative models transform noisy images into natural-looking images through sequential denoising steps. A common strategy in this domain is to use a pre-trained, static diffusion model, which eliminates the need for additional training and enhances the model’s flexibility and adaptability. These methods address image completion constraints by using various replacement strategies. For example, by replacing predicted image regions with corresponding parts from the reference image [15], or by incorporating corrupted regions from the reference image into the intermediate

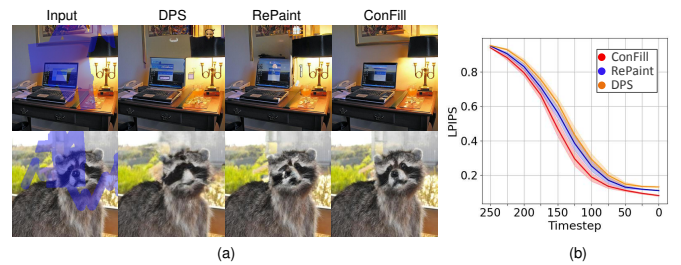


Fig. 1: (a) Generated images by DPS, RePaint, and our ConFill, using a fixed diffusion model on different masked inputs. (b) LPIPS score vs. denoising timesteps for a single input image.

denoising process [16], [17]. Similarly, [2] focuses on optimizing pixel distribution rather than denoising, aiming to reduce the number of iterations needed for training. Although these methods modify large portions of an image, they struggle to properly adjust the pixels in unknown areas to ensure consistency with the surrounding context, resulting in noticeable discontinuities, which significantly compromise the overall quality of the generated images [18]. Additionally, the sampling techniques used in [19], [20], [21] often compromise sample quality by introducing additional noise, further limiting their effectiveness. In image completion, the efficacy of diffusion models is significantly influenced by the discrepancy metrics they use [22], [23]. These metrics are important as they determine the fidelity of the transition from noisy image to an output that aligns with the original image regions. Traditional discrepancy metrics used in diffusion models, such as Euclidean distance, generally focus on minimizing statistical disparities without considering the complex spatial and structural details of the images. For example, two images might have similar pixel values but different textures or structures, leading to high Euclidean distance but visually incoherent outputs [24]. Moreover,

P. Shamsolmoali is with the Department of Computer Science, University of York, UK and SCEE, East China Normal University, China (pshams55@gmail.com)

M. Zareapoor (corresponding author) is with the Department of Electrical Engineering, Shanghai Jiao Tong University, China (mzarea222@gmail.com).

H. Zhou is with the School of Computing and Mathematical Sciences, University of Leicester, UK (hz143@leicester.ac.uk).

M. Felsberg is with Computer Vision Laboratory, Linköping University, Sweden (michael.felsberg@liu.se).

D. Tao is with the College of Computing and Data Science at Nanyang Technological University, Singapore (dacheng.tao@gmail.com).

X. Li is with the Institute of Artificial Intelligence (TeleAI) of China Telecom (xuelong_li@ieee.org).

P. Shamsolmoali and M. Zareapoor contributed equally.

the limitations of traditional discrepancy metrics become more visible when models are required to adapt to the specific needs of different regions of an image. Indeed, each part of an image may have unique characteristics such as texture, lighting, or context that require specific handling to achieve a realistic completion. DPS [25] and SMCDiff [18] use Bayesian optimization for conditioned sampling, however, their reliance on computationally intensive processes and approximations compromises both accuracy and visual quality. Fig. 1 illustrates examples of incoherent image completions for masked images. DPS shows a noticeable disparity between the left and right sides of the generated images. RePaint also has discontinuities in the reconstructions. MCG [26] uses manifold constraints, but this can restrict diffusion process, leading to underfitting and failing to capture complex data variability.

CoPaint [27] addresses this problem by modifying both the known and unknown regions to achieve more seamless integration. This approach alters key image features, such as edges that should remain intact, leading to errors that compromise the objectives of image completion. In these diffusion models, the focus is on learning the overall noise distribution, with limited attention to the alignment between consecutive denoising steps. While this assumption of treating each step as independent can be effective for completing small missing regions, it becomes problematic when dealing with large or complex areas. Without explicitly modeling the transitions between steps, inconsistencies can arise, leading to artifacts or unrealistic completions, particularly in regions requiring detailed reconstruction.

A potential solution is to optimize the mapping between inverse diffusion steps. Recently, [23] uses Brenier theorem [28], [29] to ensure structural fidelity during transformations between two probability distributions for generative task. Brenier theorem provides a globally optimal solution for mapping distributions, but it does not consider local coherence within patches of the image [30]. In diffusion-based image completion, maintaining smooth and coherent transitions across these local patches is essential. This gap highlights the need for a more flexible alignment process that can better handle the diverse patterns and textures within an image. In this paper, we introduce ConFill, a novel framework specially designed for image completion. ConFill overcomes the limitations of traditional discrepancy metrics in diffusion models and enhances both precision and contextual relevance in image reconstruction. At the core of ConFill is the Context-adaptive Discrepancy (CAD) method, which redefines the denoising process by minimizing inter-step discrepancies in the diffusion sequence, ensuring smoother transitions and improving alignment with the original image content. CAD incorporates a context-sensitive adaptation mechanism that dynamically adjusts to local image attributes throughout the diffusion process. This ensures that each image region is processed based on its unique characteristics, thereby enhancing the alignment of the reconstructed image with the original content. As a result, every step in the diffusion cycle is optimized for better contextual consistency. Moreover, ConFill enhances image reconstruction quality by selectively sampling from the posterior distribution of intermediate images, which helps reduce errors and improve the quality of the denoised output.

Indeed, ConFill through combination of CAD and dynamic sampling bridges the gap between local and global features, dynamically aligning the diffusion process with both the immediate context and the broader structure of the image. Additionally, the framework optimises the iterative process of final image generation, leading to more efficient computation and enhancing

the overall performance of the model.

In our experimental analysis, ConFill is evaluated on three challenging datasets using masks of varying shapes and sizes. The results demonstrate that ConFill outperforms existing diffusion-based methods in achieving superior fidelity and coherence in image completion tasks. These findings are supported by enhanced performance metrics in both qualitative and quantitative analyses, highlighting the model’s effectiveness in complex image completion. Specifically, ConFill achieves a relative improvement of 10.9% and 6.7% in LPIPS scores compared to the leading baselines, RePaint and CoPaint, respectively, on the ImageNet dataset. Our main contributions are as follows:

- We introduce ConFill, a unified learning framework that integrates our CAD model with an enhanced diffusion model, providing a robust solution for image completion. This integration enables the model to adaptively refine image completions by using contextual features, resulting in higher precision in denoising and improved coherence across the completed images.
- ConFill uses principles from the Brenier potential to determine the transformation that minimizes the discrepancy between latent distributions across different timesteps. This approach reduces mode mixing, resulting in high-fidelity samples that closely match the target data distribution and effectively address the inverse problem.
- We introduce a Dynamic Sampling mechanism that adjusts sampling density based on structural complexity. This strategy improves the focus and efficiency of the diffusion process, particularly in regions with higher complexity.

The rest of this manuscript is structured as follows: Section 2 discusses related studies. In Section 3, we introduce the details of ConFill. Section 4 discusses the experimental results and ablation study. Finally, Section 5 provides the concluding remarks.

2 RELATED WORK

2.1 Image Completion

Conventional image completion models are based on robust low-level assumptions, such as using local patches [31] to reconstruct missing regions in images. Recent advancements in CNN-based methods for image inpainting [6], [7], [12], [32], [33], [34], [35], [36], [37], [38] generally adopt encoder-decoder structures or variations of these stacked architectures. In particular, [32] introduces a framework based on GANs that uses both local and global discriminators. CoMod-GAN [39] is another remarkable method that enhances image generation quality by including a stochastic noise vector in the encoded representation. Building upon the GAN framework, innovative components have been integrated, including regional convolutions [40], regional normalization [41] and attentions [33], [36]. New advances such as the Vision Transformer [42] are also attracting significant interest in the image inpainting community [9], [10], [11], [43]. Meaningful priors, including semantic labels [44] and edge information [45], perform an important role in improving the process. To generate high-quality completions, multi-stage networks are a popular approach. For instance, [46] demonstrates refinement of low-resolution outputs through contextual aggregation. MAT [9] and ZITS [37] represent transformer-based image inpainting systems specifically designed for high-resolution images. However, their

multi-stage network structure tends to be slow for practical, real-world applications. In contrast, LaMa [6] introduces a one-stage approach by integrating multi-scale receptive fields to enhance efficiency. However, the speed of these systems is still limited by the conventional decoders they use.

While these approaches have demonstrated promising results in image completion tasks, they largely rely on supervised learning, where networks are trained on specific types of degradation. These approaches demand considerable computational resources and often struggle with masks not seen during training, resulting in limited generalization capabilities [2]. More recently, diffusion methods have been gaining popularity for their impressive performance in image generation tasks [47], [48], [49], [50], [51], [52], [53]. An additional advantage of these methods is their ability to perform image completion tasks effectively without the need for training on specific types of degradation [54]. In this section, we discuss existing diffusion-based image completion methods, classifying them into two primary categories: supervised and unsupervised approaches [55].

2.2 Diffusion Image Completion

Supervised diffusion-based image completion methods require training a diffusion model designed for completion tasks, with a focus on various forms of image degradation. In [56], the diffusion model receives the corrupted image at every stage of its training process. A similar approach is adopted by [57] and [58], in which a text-conditional diffusion model is specifically designed for completion tasks. In [59], an autoencoding model is integrated to compress the image space, followed by the concatenation of spatially aligned conditioning information with the model’s input. In contrast, [48] uses a non-expansive mapping approach to aggregate the degradation operation throughout the training process. In [15], a conditional diffusion model is proposed, designed to focus on refining and improving outputs from a deterministic predictor through specific training of the diffusion model. Additionally, [60] and [61] address a range of inverse problems by directly estimating conditional scores based on the specific measurement model without needing extra training. [26] addresses the inverse problems in diffusion models by integrating manifold constraints and [62] introduces an enhanced inference method by modeling variable noise levels at the pixel level. Recent studies have explored extending diffusion models with external guidance. ControlNet [63] integrates structured conditions like edge maps, enabling precise control over outputs. Similarly, PowerPaint [64] uses task-specific prompts to guide inpainting tasks. Both these approaches highlight the adaptability of diffusion models to specific tasks through conditioning. However, these approaches require training for specific types of degradation, leading to intensive computational demands and limited adaptability to degradation operators not used in the training dataset.

Unlike supervised approaches, unsupervised diffusion image completion approaches rely on pre-trained diffusion models, which eliminates the requirement for model structure adjustments. Our method is categorised as unsupervised diffusion. [54] proposes a modification to the DDPM (denoising diffusion probabilistic models) [13] sampling process, involving the integration of a noisy variant of the damaged image at every stage of the denoising sequence. This approach is similarly applied in [16] for text-driven image completion. In DDRM [55], a posterior diffusion process is formulated, demonstrating that its marginal

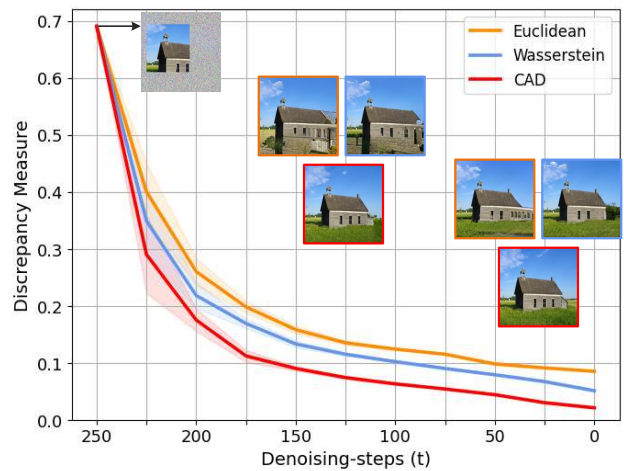


Fig. 2: The evolution of the discrepancy during the denoising process. This illustrates the mean pixel-wise discrepancy using Context-Adaptive Discrepancy (CAD) (red), Wasserstein Discrepancy (blue), and Euclidean Distance (orange).

probability distribution aligns with that of the DDPM framework. This denoising process can be approximated as a weighted-sum blending of the degraded image iteratively at each timestep. Even if these blending-based techniques are effective, the images they generate often lack harmony in the recovered parts, leading to inconsistencies in the generated images [17].

To address this challenge, RePaint [17] introduces a new resampling strategy that incorporates a bidirectional operation in diffusion time. This approach blends images from the current timestep t with a noisy version of the degraded image, enhancing the reconstruction. The resulting blended images are then used to generate outputs for the next timestep $t + 1$ through a one-step forward process. By incorporating this mechanism, RePaint effectively reduces the visual inconsistencies commonly associated with basic blending methods. In [18], it is demonstrated that methods based on simple blending can lead to irreducible approximation errors during the image generation process. To address this, a particle filtering-based approach (residual resampling) is proposed. In this method, for each timestep t , every generated image is resampled according to its likelihood of producing the known region of the degraded image at the previous timestep $t - 1$. This method aims to enhance the accuracy of the generation process by aligning it more closely with the known parts of the degraded image. [65] explores diffusion models through the lens of deep equilibrium analysis, leading to the development of an efficient method for inverting denoising diffusion implicit models [14]. This approach focuses on optimising memory usage to make diffusion models more effective. DDNM [15] employs a blending strategy that incorporates the damaged image directly into each timestep of the diffusion process without adding any noise. Meanwhile, DPS [25] addresses the task of completion by adopting a method similar to classifier-free guided diffusion [66]. This method is focused on approximating posterior sampling, introducing a refined strategy for image completion. In a similar direction, CoPaint [27] proposes an efficient inpainting framework based on [14]. By using Markov diffusion processes, this method enables a deterministic sampling strategy that accelerates inference while maintaining high-quality reconstructions. However, despite its efficiency, the reliance on deterministic sampling limits diversity

in plausible completions, especially in complex inpaintings. Different from these methods, our approach refines the completion process by progressively aligning the missing regions with the surrounding context through an adaptive refinement strategy. By incorporating structural and textural constraints into the denoising process, our method ensures consistency across timesteps while maintaining high fidelity in complex regions. This allows for more accurate reconstructions that better preserve global coherence and fine-grained details.

3 OUR METHOD

3.1 Problem Statement

Image completion is a process designed to fill in damaged or missing areas of an image x_0 with content that is contextually relevant and visually coherent. This objective can be written as

$$\mathcal{L}_{\text{completion}} = \|s(\hat{x}_0) - r_0\|_2^2 + F\|\hat{x}_0 - P_{\text{data}}\|_2^2, \quad (1)$$

where $s(\cdot)$ is an operator that extracts the known parts of the image r_0 , $\|\cdot\|_2$ denotes the L_2 norm and P_{data} represents the natural image prior, while F controls the trade-off between these two objectives. Minimizing this loss ensures that the reconstructed image \hat{x}_0 aligns with the known portions r_0 .

Diffusion models are widely used as effective methods for image generation, excelling at modeling complex data distributions. These models are particularly well-suited for image completion due to their iterative denoising process, which progressively refines noisy inputs into high-quality images. In image completion, diffusion models start with a noisy version of the image and iteratively refine it to produce a plausible output. In the forward diffusion process, noise is progressively added into the image,

$$q(x_{1:T}|x_0) = \prod_{t=1}^T p(x_t|x_{t-1}), \quad (2)$$

in which x_t denotes the noisy image at timestep t , and T represents the number of timesteps. Moreover, to control the noise level, we can define the forward diffusion as $q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I})$, in which \mathcal{N} denotes a Gaussian distribution, β_t denotes the variance of the noise added at timestep t , $\sqrt{1 - \beta_t}$ scales the previous image x_{t-1} , and \mathbf{I} indicates the identity matrix. In contrast, the goal of the reverse operation is to recover x_0 from x_T ,

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

in which μ_θ and σ_θ denote the learned parameters of the model. In pixel-wise comparison, current diffusion models typically use the Euclidean distance to measure the discrepancy between the generated \hat{x}_t and the reference x_t images as follows,

$$\mathcal{L}_{\text{diffusion}} = \mathbb{E}_{q(x_{1:T}|x_0)} \left[\sum_{t=1}^T \|x_t - \hat{x}_t\|_2^2 \right]. \quad (4)$$

While Euclidean distance is simple and computationally efficient, it has limitations for image completion: it treats all pixel differences equally, regardless of their spatial arrangement or context, leading to minimized local pixel values but compromised overall perceptual quality and coherence. Additionally, Euclidean distance fails to capture complex structural and semantic relationships within an image, overlooking higher-level features like edges, textures, and object boundaries (see Fig. 2). It also does not

Algorithm 1 CAD

Require: Reference image r_0 , pre-trained denoising functions g_θ , G , adaptability function parameters $V_{\text{ar}}(x)$ and ψ
Ensure: Transport map, minimizing CAD

- 1: Initialize transport map $\mathbf{f} : \mathcal{Y} \rightarrow Y$
- 2: Use adaptive function based on local image variance $\rho(x) = \exp(-\psi \cdot V_{\text{ar}}(x))$
- 3: **for** each cell C_i in \mathcal{Y} **do**
- 4: Compute the adaptive weighting from $w_i = \rho(x_i)$
- 5: Match w_i from region C_i to the corresponding weight at y_i
- 6: **end for**
- 7: Apply the cost function for contextual sensitivity: $c_C(x, y; C) = c(x, y) \times (1 + v \cdot e^{-\tau \cdot \|C(x) - C(r_0)\|})$
- 8: Calculate the total cost of $f(x): f(x) = \sum_{i \in \mathcal{I}} \int_{C_i} c_C(x, f(x)) d\varepsilon(x)$
- 9: Use $g_\theta^{(t)}(x)$ to refine intermediate reconstructions at each timestep t
- 10: **for** $t = T$ to 1 **do**
- 11: Compute the reconstructed image $\hat{x}_t = g_\theta^{(t)}(x_t)$
- 12: Update f by minimizing the total cost using w_i and \hat{x}_t
- 13: **end for**
- 14: **return** Transport map minimizing CAD (use as input for Algorithm 2)

consider the broader context, leading to inconsistencies where generated pixels appear out of place with their surroundings.

Another approach to addressing this limitation is the Wasserstein Discrepancy (WD) [67]. At the core of this method is the gradient of the Brenier potential, which defines the optimal transport map for transforming the degraded image distribution into its restored form with minimal cost. This alignment helps capture the underlying geometric structure, ensuring that the generated content is consistent with the reference image. WD is defined as:

$$W(\mathbb{P}, \mathbb{Q}) = \inf_{\gamma \in \Pi(\mathbb{P}, \mathbb{Q})} \mathbb{E}_{(x, y) \sim \gamma} [\|x - y\|], \quad (5)$$

in which \mathbb{P} and \mathbb{Q} denote probability distributions, and $\Pi(\mathbb{P}, \mathbb{Q})$ represents the set of all distributions that have \mathbb{P} and \mathbb{Q} as their marginal distributions. As shown in Fig. 2, while the WD outperforms the Euclidean distance by considering spatial relationships and the distribution of pixel values, it still does not minimize the discrepancy between the generated and reference images effectively. This indicates that although WD captures geometric structure, it does not fully address factors such as local coherence or semantic alignment.

3.2 Foundational Strategy

To address the above limitations, we introduce CAD, which integrates contextual information into the discrepancy measure as

$$W_C(x, y; C) = \|\nabla \phi(x) - y\|^2 \cdot f(C(x), C(y)), \quad (6)$$

in which $C(\cdot)$ denotes contextual features extracted from the image, such as texture or color gradients, $\phi(x)$ is the Brenier potential that defines an optimal transport map to ensure monotonic and cost-efficient displacement of probability mass between latent distributions. The discrepancy function adapts dynamically based on local and global features, aligning the generative process to a transport-based regularization, and $f(\cdot, \cdot)$ is a scaling function based on the differences between these features

$$f(C(x), C(y)) = 1 + v \cdot e^{-\tau \cdot \|C(x) - C(y)\|}, \quad (7)$$

where v and τ mainly control the sensitivity of the scaling function to contextual differences. This process ensures that the

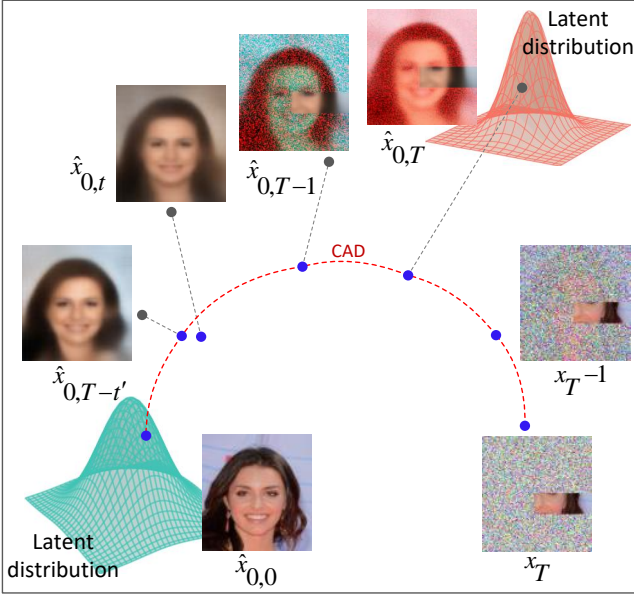


Fig. 3: The ConFill framework. The red curve represents the CAD inverse diffusion steps, which iteratively balance the distribution of image patches across the latent space. The green and red hills represent latent distributions closer to the final completed image and noisy initial steps, respectively.

discrepancy metric dynamically adjusts to both local and global image features, providing a more accurate measure of similarity. Therefore, the CAD loss can be written as

$$\mathcal{L}_{CAD} = \mathbb{E}_{q(x_{1:T}|x_0)} \left[\sum_{t=1}^T W_C(\hat{x}_t, x_0; C) \right]. \quad (8)$$

In our approach, we refine x_t by minimizing the CAD loss

$$\mathcal{L}_{CAD} = \|x_t - \mu_\theta(x_{t+1}, t+1)\|_2^2 + \gamma_t \cdot W_C(s(g_\theta(x_t)), r_0; C), \quad (9)$$

where γ_t is a weighting parameter, and g_θ is the denoising function. Then apply gradient descent to update x_t at each timestep ($x_{t-1} = x_t - \lambda \cdot \nabla_{x_t} \mathcal{L}_{CAD}$), where λ is the learning rate. Considering the diffusion process for all t , the denoising follows a deterministic process aimed at minimizing CAD between \hat{x}_T and the original data distribution, aligning the model with the target distribution. Therefore, CAD directly influences the optimization objective in Eq. (9) and estimation of \hat{x}_0 . Thus, the completion constraint $\mathcal{O} : s(\hat{x}_0)$, can be applied to \hat{x}_T . Therefore, addressing the image completion challenge depends on determining an optimal \hat{x}_T guided by the posterior distribution

$$p_\theta(\hat{x}_T | \mathcal{O}) = p_\theta(\hat{x}_T) \cdot \delta_{CAD}(s(\hat{x}_0) = r_0), \quad (10)$$

where δ_{CAD} is the delta function in the context of CAD, emphasizing the optimal alignment of \hat{x}_0 and r_0 . As established in Eq. (2) and Eq. (3), $p_\theta(\hat{x}_T)$ is a Gaussian distribution. Notably, $p_\theta(s(\hat{x}_0) = r_0 | \hat{x}_T)$ represents the conditional likelihood of $s(\hat{x}_0)$ at r_0 , given \hat{x}_T . With \hat{x}_T fixed and \hat{x}_0 as a deterministic function of \hat{x}_T , this measure is represented by a delta function $\delta_{CAD}(\cdot)$, which reflects the alignment between $s(\hat{x}_0)$ and r_0 . This function indicates perfect alignment, where the CAD is minimized, ensuring that the transformation between image distributions more effectively minimizes discrepancies. Consequently,

Algorithm 2 ConFill

Require: f (from Algorithm 1), r_0 , g_θ , \hat{t} , \mathcal{J} , G , and λ_t (at t)

- 1: **Initialize:** $\hat{x}_T \sim \mathcal{N}(0, \mathbf{I})$
- 2: Set $t = T$ and $j = \mathcal{J}$
- 3: **while** $t > 0$ **do**
- 4: **Step 1: CAD Refinement**
- 5: Compute transport potential: $\nabla \phi_t(\hat{x}_t)$ (Eq. (6))
- 6: Update \hat{x}_t by minimizing: $\text{CAD}(\hat{x}_t, r_0) + \lambda_t \|\nabla \phi_t(\hat{x}_t) - \hat{x}_t\|$ (Eq. (9))
- 7: Apply gradient penalty to enforce Lipschitz constraints.
- 8: **Step 2: Reverse Diffusion Update**
- 9: Compute \hat{x}_{t-1} via CAD-driven transition: $\hat{x}_{t-1} = g_\theta(t)(\hat{x}_t) - \lambda \nabla \phi_t(\hat{x}_t)$
- 10: Reduce timestep: $t \leftarrow t - 1$
- 11: **Step 3: Temporal Regression**
- 12: **if** $t \leq T - \hat{t}$ **and** $t \bmod \hat{t} = 0$ **then**
- 13: **if** $j > 0$ **then**
- 14: Predict $\hat{x}_{t+\hat{t}}$ via transition distribution $q(\hat{x}_{t+\hat{t}}|\hat{x}_t)$ (Eq. (19))
- 15: Adjust timestep: $t \leftarrow t + \hat{t} - 1$
- 16: Decrement jump count: $j \leftarrow j - 1$
- 17: **else**
- 18: Reset jump count: $j \leftarrow \mathcal{J}$
- 19: **end if**
- 20: **end if**
- 21: **end while**
- 22: **Output:** Reconstructed image \hat{x}_0

the logarithm of the posterior distribution in Eq. (10) can be directly estimated as follows

$$\log p_\theta(\hat{x}_T | \mathcal{O}) \approx -\frac{1}{2} \|\hat{x}_T\|_2^2 - \frac{1}{2\gamma_T^2} \text{CAD}(r_0, s(\mathcal{F}_\sigma(\hat{x}_T)))^2 + \Gamma, \quad (11)$$

where $\mathcal{F}_\sigma(\hat{x}_T) = \hat{x}_0$ represents that \hat{x}_0 is functionally dependent on \hat{x}_T . The Γ denotes the normalising constant, while γ_T represents the dispersion of the distribution, shaping the precision of the CAD-based approximation. As γ_T tends towards zero, the approximation in Eq. (11) becomes increasingly accurate. Indeed, optimally selecting γ_T can significantly enhance performance. Eq. (11) provides a basis for determining \hat{x}_T using the optimization method. The first part of this equation serves as a guiding framework, while the second part imposes a constraint to ensure accurate completion. To estimate \hat{x}_T , we use Wasserstein gradient flows [68] to generate random samples from complex distributions. This optimization effectively addresses the inconsistency issues observed in current methods [69].

Contextual Feature Extractor ($C(\cdot)$): To compute the context-adaptive discrepancy $c_C(x, y; C)$, the $C(\cdot)$ combines handcrafted and learned features. I) Handcrafted features are used to capture low-level structural and color information. These include edge-based features, texture-based features, and color distributions. Together, these features provide a robust representation of local structural patterns and variations in color across the image. II) Learned features are used to extract high-level semantic features, including object boundaries, relationships, and contextual patterns. Our workflow is illustrated in Fig. 3. In our model Algorithm 1 computes the transport map f by minimizing the CAD loss. This ensures alignment between the generated regions and the known parts of the image using adaptive weightings and cost functions. f serves as an intermediate representation and acts as input to Algorithm 2. This Algorithm takes f along with g_θ to iteratively refine the image reconstruction.

One of the challenges with our approach is its computational

demand. Specifically, generating the final output $\mathcal{F}_\sigma(\hat{x}_T)$ and calculating its gradient require executing both forward and backward passes through a denoising sequence. To address this, it is essential to improve the computational efficiency of the algorithm while preserving the efficacy of our strategy. Referring to the discussion in Sec. 3.2, a practical modification involves substituting the complex evaluation of $\mathcal{F}_\sigma(\hat{x}_T)$ in Eq. (11) with the gradual generation $g_\theta^{(T)}(\hat{x}_T)$. This adjustment expedites the approximation of the final output and streamlines the computation process. Therefore, we establish an estimated conditional distribution model for $s(\hat{x}_0)$ conditional on \hat{x}_T , represented as $p'_\theta(s(\hat{x}_0)|\hat{x}_T)$. This estimation is produced by a singular step generation, $s(g_\theta^{(T)}(\hat{x}_T))$, and incorporates a small Gaussian perturbation to model the variability.

$$p'_\theta(\hat{x}_T|s(\hat{x}_0)) \sim \mathcal{N}(s(\hat{x}_0); s(g_\theta^{(T)}(\hat{x}_T)), \gamma_T'^2 \mathbf{I}), \quad (12)$$

in which γ_T' denotes the parameter of standard deviation. Upon integrating this estimated distribution, we obtain an accurate approximation of the posterior distribution

$$\log p'_\theta(\hat{x}_T | \mathcal{O}) \approx -\frac{1}{2} \|\hat{x}_T\|_2^2 - \frac{1}{2\gamma_T'^2} \text{CAD}\left(r_0, s\left(g_\theta^{(T)}(\hat{x}_T)\right)\right)^2 + \Gamma', \quad (13)$$

in which Γ' acts as a normalisation factor. To effectively reduce the approximation gap, between $p_\theta(s(\hat{x}_0)|\hat{x}_T)$ and $p'_\theta(s(\hat{x}_0)|\hat{x}_T)$, we propose adjusting the variance parameter $\gamma_T'^2$ with respect to CAD. This ensures the estimated distribution more accurately reflects the true distribution, thereby enhancing the approximation.

$$\gamma_T'^2 = \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \text{CAD}\left(s\left(g_\theta^{(T)}(\hat{x}_T^{(i)})\right), s(\hat{x}_0^{(i)})\right), \quad (14)$$

we adjust $\gamma_T'^2$ based on the average CAD across samples, where \mathcal{N} represents the r_0 's dimension. Here, $\text{CAD}(\cdot, \cdot)$ measures the shortest distance between the transformed versions of \hat{x}_T and \hat{x}_0 . Maximizing Eq. (13) with respect to \hat{x}_T aims to balance the image completion constraint with prior regularization. Indeed, unlike Eq. (11), which minimizes γ_T , Eq. (14) adjusts γ_T' to be sufficiently large. This adjustment minimises the approximation error, thereby reducing the emphasis on the image completion constraint in Eq. (13). Consequently, this ensures a balanced approach to minimizing errors while maintaining image integrity.

3.3 Dynamic Sampling

To enhance the effect of CAD in regions with varying complexity, we introduce a novel dynamic sampling technique. This strategy dynamically modulates the sampling density by using an adaptable function, which evaluates the textural and structural complexity inherent in each image region. Textural complexity is quantified as the local variance of pixel intensities within a region as $\text{Var}(x) = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2$, where $\mu = \frac{1}{N} \sum_{i=1}^N x_i$, and x_i represents the pixel intensities. Moreover, the structural complexity is measured using edge density (ED), defined as the proportion of edge pixels in the region. The two measures are combined into the adaptability function $\varrho(x)$, which assigns weights based on the relative contributions of textural and structural complexity:

$$\varrho(x) = \hat{\alpha} \cdot \exp\left(-\psi \cdot \frac{\text{Var}(x) \cdot N}{N-1}\right) + \hat{\beta} \cdot \text{ED}(x), \quad (15)$$

where $\hat{\alpha}$ and $\hat{\beta}$ control the relative importance of textural and structural complexity, $\text{Var}(x)$ represents the local variance within a specific region of the image, and ψ is a scaling parameter that adjusts sensitivity to textural variations. The dynamic sampling process enhances CAD computation by integrating the adaptability

function as a weighting factor within the discrepancy measure. Consequently, CAD can be computed as

$$\text{CAD}(x, y) = \sum_{i=1}^N \varrho(x_i) \cdot W(p_{x_i}, p_{y_i}), \quad (16)$$

where $W(p_{x_i}, p_{y_i})$ is the discrepancy between the distributions p_{x_i} and p_{y_i} at points x_i and y_i . During the image completion process, the sampling rate for each pixel or region is adjusted based on $\varrho(x)$, effectively increasing the density of sample points in regions with higher complexity $\mathcal{S}(x) = \lceil \mathfrak{m} \cdot \varrho(x) \rceil$, where $\mathcal{S}(x)$ represents the number of samples, and \mathfrak{m} is the maximum samples per region. By dynamically adjusting the sampling based on local image characteristics, CAD enables more precise control over the restoration process. The CAD framework balances local and global features by capturing edges and textures through the handcrafted component of $C(\cdot)$ and ensuring global semantic consistency using high-level features from a pre-trained network. This dual strategy dynamically adapts the discrepancy metric, aligning fine-grained details with broader structural coherence for superior completion quality.

3.4 Gradual Approximation

To further refine our stochastic diffusion process, we gradually refine the approximation error at each stage of the process. This ensures that the approximated generation closely aligns with the reference image r_0 . Our diffusion model draws samples $\hat{x}_{0:T}$ from an estimated posterior distribution $p'_\theta(\hat{x}_{0:T}|\mathcal{O})$, in order to reduce the gap between the estimated distribution and the true distribution of $\hat{x}_{0:T}$ given \mathcal{O} , which can be expressed as

$$p'_\theta(\hat{x}_{0:T} | \mathcal{O}) = p'_\theta(\hat{x}_T | \mathcal{O}) \cdot \prod_{t=1}^T p'_\theta(\hat{x}_{t-1} | \hat{x}_t, \mathcal{O}), \quad (17)$$

in which $p'_\theta(\hat{x}_T | \mathcal{O}_T)$ is defined based on minimising the CAD, as indicated in Eq. (13). To compute $p'_\theta(\hat{x}_{t-1} | \hat{x}_t, \mathcal{O})$, we apply a series of Gaussian approximation distributions to minimize the distance to the true conditional distributions at each step t , similar to the approach described in Eq. (12).

$$p'_\theta(\hat{x}_t \rightarrow s(\hat{x}_0)) \sim \mathcal{N}\left(\text{CAD}(g_\theta^{(t)}(\hat{x}_t) \rightarrow s(\hat{x}_0)), \mathbf{I} \cdot \gamma_t'^2\right). \quad (18)$$

For minimising the gradual approximation error, we define $\gamma_t'^2$ as per Eq. (14), with t replacing T . This adjustment allows us to more accurately compute $p'_\theta(\hat{x}_{t-1}|\hat{x}_t, \mathcal{O})$ as

$$\log p'_\theta(\hat{x}_{t-1} | \hat{x}_t, \mathcal{O}) \approx -\frac{1}{2\sigma_t^2} \|\hat{x}_{t-1} - \hat{\kappa}_t\|_2^2 - \frac{1}{2\gamma_{t-1}'^2} \text{CAD}\left(r_0 - s\left(g_\theta^{(t-1)}(\hat{x}_{t-1})\right)\right)^2 + \Gamma', \quad (19)$$

where the derivation of the second term can be written as

$$\hat{\kappa}_t = g_\theta^{(t)}(\hat{x}_t) \cdot \sqrt{\alpha_{t-1}} + \left(\frac{\hat{x}_t - g_\theta^{(t)}(\hat{x}_t) \cdot \sqrt{\alpha_t}}{\sqrt{1 - \alpha_t}}\right) \cdot \sqrt{1 - \alpha_{t-1} - \sigma_t^2}, \quad (20)$$

where $\alpha_t = 1 - \beta_t$ denotes the noise level at each timestep and the standard deviation of the noise distribution is σ . For generating the final completion result, we use a step-by-step optimization method aimed at selecting sequences of $\hat{x}_{0:T}$ that increase the probability $p'_\theta(\hat{x}_{0:T}|\mathcal{O})$ mentioned in Eq. (17). This process starts by choosing an initial \hat{x}_T based on Eq. (13). Using \hat{x}_T , the subsequent sample \hat{x}_{t-1} is obtained by optimizing the objective function outlined in Eq. (19). These iterative steps are designed



Fig. 4: Qualitative comparison on the CelebA-HQ. The facial images generated by ConFill not only present more distinctive facial features but also show a higher degree of similarity and coherence with the original images compared to those produced by other baseline models.

to satisfy the completion criteria within the denoising diffusion approach. Fig. 3 shows that the CAD-based approximation error decreases as t reduces, indicating that the algorithm progressively improves fidelity to the completion constraint with each step. Specifically, when t reaches 1, by assigning $\sigma_1 = 0$ and reducing γ_1 towards zero, we can minimise the approximation error to an extent where $g_{\theta}^{(1)}(\hat{x}_1)$ approximates \hat{x}_0 with high accuracy. This method ensures high quality generation while efficiently reducing transport costs with each iteration.

While our algorithm reduces approximation error during the final denoising phase, early inaccuracies still affect image quality by influencing the prior distribution for later steps. To address this, during the initial denoising phase, in which the significance of approximation errors is higher, we adopt a multi-step approximation approach. This involves approximating \hat{x}_0 by traversing through multiple deterministic denoising steps at selected time intervals. As outlined in [15], [17], [27], we can integrate the time travel method to bolster the self-consistency of intermediate samples. This method involves returning previous denoising steps by adding noise to the intermediate images. At the denoising time step $T - \hat{t}$, we consider a predefined range of time steps denoted as ρ . Indeed, rather than proceeding to $T - \hat{t} - 1$, we backtrack to time $T - 1$ and sample a new \hat{x}_{T-1} according to $q(\hat{x}_{T-1}|\hat{x}_{T-\hat{t}})$, and repeat the denoising steps from that point. This iterative process repeats for \mathcal{J} rounds, alternating between rewinding and denoising. The sequence progresses through steps $T - 1$ to $T - \hat{t}$ and then $T - \hat{t} - 1$ to $T - 2\hat{t}$ repeating iteratively until reaching time zero. The details of our model’s procedure are illustrated in Algorithm 2.

4 EXPERIMENTS

4.1 Setup Details

Datasets and models: Our methodology is evaluated on three widely recognised image datasets: CelebA-HQ [70], Places2 [71], and ImageNet-1K [72]. Given that diffusion models typically require square-shaped images, we pre-process the datasets by




Fig. 5: Qualitative comparison on the Places2. ConFill outperforms the other models, providing accurate and seamless restoration. The inpainted regions blend naturally with the surrounding landscape, maintaining consistent texture and structure across the images.

cropping all images to a uniform size of 256×256 pixels to ensure compatibility with pre-trained diffusion models. We use the pre-trained diffusion model from [17] and [66] for the CelebA, Places2 and ImageNet. Hyperparameter selection is conducted using the first five images from the validation set. Evaluation is performed on the images from the test sets, similar to the methodology outlined by [17]. For the contextual feature extractor $C(\cdot)$ we use Sobel and Gabor filters for edge and texture extraction. Additionally, a ResNet18 is used to extract high-level semantic features. The network is pre-trained on ImageNet and fine-tuned on the target dataset.

Masks: We use masks from [6] and [17]. These masks are created by uniformly sampling from polygonal chains that have been expanded by a randomly determined width, combined with rectangles of diverse aspect ratios and orientations.

Metrics: The quality of the completion results is evaluated through a combination of objective and subjective metrics. We use the Fréchet Inception Distance (FID), Learned Perceptual Image Patch Similarity (LPIPS), and Structural Similarity Index Measure (SSIM). For each image, two completed images are produced, and we report the mean scores. Our evaluation focuses on the degree to which the completed images achieve a natural appearance free of artefacts, closely mirror the original reference image in the areas that were not masked, and maintain a cohesive visual flow throughout the entire image. This set of criteria is designed to assess the quality of the completion results, focusing on realism, fidelity to the original scenes, and internal consistency.

Baselines comparisons and implementation details: Our analysis focuses on diffusion approaches, which have demonstrated superior performance in comparison to non-diffusion model approaches [17]. To this end, we evaluate against established baselines, including DDRM [55], RePaint [17], CoPaint [27], DPS [25], I²SB [61], ControlNe [63], and PowerPaint [64] to highlight our comparative study. We follow the original implementations for selecting the NFE and hyper-parameters for each method.

TABLE 1: Quantitative comparisons on three (CelebA, Places2, and ImageNet-1k) datasets with various mask types. Results are based on FID (\downarrow), LPIPS (\downarrow) and SSIM (\uparrow). For each mask type and metric, **bold** indicates the best result and under the second best.

Mask	ConFill (ours)			PowerPaint			I ² SB			DPS			RePaint			CoPaint			ControlNet		
	FID	LPIPS	SSIM	FID	LPIPS	SSIM	FID	LPIPS	SSIM	FID	LPIPS	SSIM	FID	LPIPS	SSIM	FID	LPIPS	SSIM	FID	LPIPS	SSIM
<i>CelebA-HQ</i>																					
Narrow	7.08	0.034	0.896	8.65	0.058	0.859	8.03	0.051	0.864	8.72	0.063	0.854	<u>7.32</u>	0.045	0.881	<u>7.39</u>	<u>0.041</u>	<u>0.887</u>	8.77	0.049	0.851
Wide1	8.05	0.064	0.862	9.18	0.097	0.810	8.79	0.087	0.835	9.26	0.085	0.821	8.54	0.077	0.857	<u>8.48</u>	<u>0.070</u>	0.851	8.93	0.091	0.818
Wide2	8.13	<u>0.076</u>	0.854	9.36	0.102	0.804	8.85	0.095	0.830	9.31	0.093	0.824	8.61	0.073	0.854	8.70	0.082	0.842	9.28	0.097	0.811
Half ver.	12.94	0.187	0.735	16.97	0.236	0.708	16.56	0.214	0.719	16.46	0.198	0.721	<u>16.38</u>	<u>0.193</u>	<u>0.725</u>	16.49	0.203	0.720	16.54	0.205	0.720
Half hor.	12.67	<u>0.184</u>	0.742	16.74	0.229	0.716	16.48	0.207	0.725	16.37	0.192	0.729	<u>16.34</u>	0.187	0.729	<u>16.34</u>	0.180	0.735	16.43	0.210	0.727
Expand	36.79	<u>0.440</u>	0.483	54.22	0.507	0.447	53.51	0.498	0.443	49.31	0.435	0.476	<u>50.83</u>	0.476	0.463	<u>48.75</u>	0.472	0.470	52.61	0.480	0.454
Avg	14.27	<u>0.164</u>	0.760	19.18	0.204	0.724	18.70	0.192	0.736	18.24	0.178	0.738	18.02	0.175	0.750	<u>17.69</u>	<u>0.173</u>	<u>0.751</u>	18.76	0.189	0.730
<i>Places2</i>																					
Narrow	8.91	0.063	0.847	10.54	0.106	0.805	10.56	0.094	0.817	10.45	0.108	0.803	9.35	0.072	0.834	<u>9.26</u>	<u>0.067</u>	<u>0.841</u>	10.57	0.108	0.801
Wide1	10.27	<u>0.110</u>	0.829	12.20	0.143	0.773	12.03	0.123	0.804	12.02	0.121	0.796	11.73	0.116	0.820	<u>11.67</u>	0.106	0.824	12.13	0.139	0.780
Wide2	10.39	0.120	0.812	12.49	0.161	0.770	12.13	0.138	0.797	12.23	0.142	0.789	12.05	0.128	0.806	<u>11.86</u>	<u>0.124</u>	0.802	12.42	0.154	0.773
Half ver.	14.02	0.270	0.663	19.22	0.290	0.647	19.67	0.291	0.644	20.18	0.306	0.635	15.49	0.286	0.651	<u>15.41</u>	<u>0.283</u>	0.658	19.04	0.289	0.651
Half hor.	13.76	0.258	0.676	19.34	0.282	0.641	19.86	0.284	0.651	20.09	0.294	0.643	<u>15.30</u>	0.269	0.662	15.39	<u>0.264</u>	0.660	19.15	0.278	0.647
Expand	48.65	0.543	0.417	57.74	0.602	0.368	58.59	0.609	0.374	57.83	0.581	0.388	<u>60.47</u>	0.618	0.354	<u>57.47</u>	<u>0.591</u>	0.378	58.04	0.611	0.359
Avg	17.65	0.227	0.707	21.93	0.263	0.667	22.14	0.257	0.681	21.13	0.259	0.675	20.73	0.248	0.688	<u>20.17</u>	<u>0.239</u>	<u>0.694</u>	21.89	0.262	0.669
<i>ImageNet</i>																					
Narrow	11.15	<u>0.071</u>	0.812	12.81	0.124	0.762	12.74	0.116	0.769	12.69	0.125	0.768	<u>11.75</u>	0.083	0.793	11.97	0.068	0.804	12.85	0.126	0.760
Wide1	13.24	0.121	0.785	14.28	0.141	0.756	14.16	0.137	0.760	14.21	0.135	0.754	13.98	0.125	0.781	<u>13.93</u>	0.132	0.772	14.25	0.138	0.758
Wide2	13.37	<u>0.134</u>	0.776	14.42	0.157	0.739	14.35	0.151	0.749	14.46	0.158	0.741	14.27	0.146	0.764	<u>14.11</u>	0.129	0.759	14.40	0.154	0.741
Half ver.	15.59	0.283	0.614	22.67	0.338	0.562	21.63	0.308	0.594	22.47	0.322	0.578	17.43	0.292	0.604	<u>17.38</u>	0.298	0.598	22.80	0.342	0.553
Half hor.	15.24	0.271	0.622	22.53	0.324	0.582	21.84	0.302	0.603	22.41	0.317	0.594	17.36	0.289	0.607	<u>17.29</u>	0.284	0.607	22.61	0.330	0.570
Expand	50.35	0.574	0.349	63.21	0.674	0.271	60.86	0.627	0.307	59.97	0.605	0.328	<u>62.85</u>	0.659	0.298	<u>59.43</u>	0.618	0.314	63.15	0.670	0.277
Avg	19.83	0.237	0.660	24.99	0.293	0.612	24.26	0.274	0.630	24.36	0.277	0.625	22.94	0.266	0.640	<u>22.35</u>	<u>0.254</u>	<u>0.643</u>	25.02	0.295	0.610



Fig. 6: Qualitative comparison on the ImageNet-1k. ConFill outperforms other methods in terms of preserving structural and details.

Our model has been developed using PyTorch, and we use four RTX 3080-TI GPUs for training and conducting all experiments. In the case of RePaint, we applied the publicly available code, following the configurations outlined in the original paper. We adapted the implementation of all other methods from RePaint, ensuring that the hyperparameters remained consistent with those detailed in their papers. For ConFill, specifically, we set the number of gradient descent steps, $G = 2$. We apply CAD to balance the initial token distribution before generating the final completed image. Additionally, we alternatively applied it to the intermediate steps to refine token assignments. Here, \hat{x}_t is iteratively optimised to maximise the posterior probability at timesteps. This process is equivalent to minimising the loss for \hat{x}_t .

We have observed that directly optimising \hat{x}_t using the gradient often leads to NaN values. To mitigate this issue, following [27] we fine-tune the learning rate by introducing an additional term to optimise the training process more effectively. Therefore, we start with an initial learning rate of 0.01 and then adjust it using

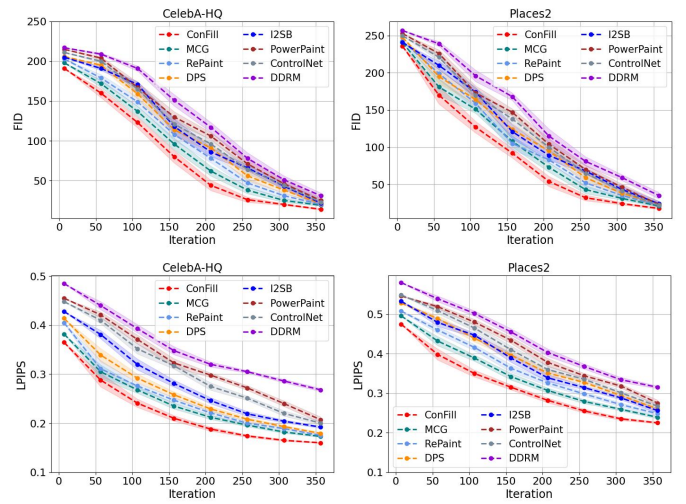


Fig. 7: Evaluation of FID (top row) and LPIPS (bottom row) scores for ConFill and baseline models on the CelebA-HQ and Places2 datasets. ConFill consistently outperforms the other models in both fidelity and perceptual similarity at all iteration levels.

an adaptive learning rate to fine-tune our learning process. We initially set $\nu = 0.1$ and $\tau = 0.02$. Moreover, ConFill’s uses a travel frequency of $\mathcal{J} = 1$ and a time travel interval of $\hat{t} = 10$.

4.2 Comparison to other Established Methods

This section presents a detailed comparison between ConFill and recent diffusion-based image completion methods.

Qualitative Comparisons: In Figs. 3, 4, and 5, we show sample images generated using datasets from CelebA-HQ, Places2, and ImageNet-1k, respectively. Two key observations can be drawn from these results. Firstly, ConFill generates more realistic content with better coherence and semantic consistency compared to baseline methods, particularly for larger masks like Expand and Half. For example, in the fifth row of Fig. 5, where only half of a desk is visible in the masked region, most baseline methods fail to produce satisfactory completions. In contrast, ConFill successfully reconstructs a desk that aligns with the input image in both size and style. Additionally, while some approaches, like DPS and CoPaint, do create relatively realistic images, our method




Fig. 8: The results of the user study comparing our model, ConFill, with baseline models CoPaint, RePaint, and DPS across various mask ratios on three different datasets: CelebA, Places2, and ImageNet. ConFill consistently outperforms the other models across all datasets, as demonstrated by the higher human preference rates, especially at lower mask ratios. Notably, ConFill achieves a +21% increase on the Places2 dataset and a +29% increase on the ImageNet dataset compared to the best-performing baseline models.




Fig. 9: A comparison of image inpainting results on the facial images with large unseen holes (left) and a corresponding evaluation of FID scores with varying mask ratios (right). ConFill demonstrates superior performance in generating realistic and coherent facial features, with smoother integration and fewer artifacts than DPS, RePaint and CoPaint. As the mask ratio increases, ConFill consistently achieves lower FID scores, indicating its effectiveness in preserving image quality under increasingly challenging conditions.

consistently generates images with greater consistency. As an example, in the last row of Fig. 5, the sunset view produced by our method shows greater consistency with the input image, whereas the versions produced by DPS and CoPaint are less detailed. Our model demonstrates robust performance, enabling the generation of high-quality outputs that remain coherent with the existing image content. This efficacy is attributed to our model’s innovative refining posterior distribution estimation mechanism during denoising. This mechanism starts with the visible pixels and methodically diffuses essential information throughout the image, assuring flawless integration of the generated content with the known regions of the image. This approach is particularly effective for facial images that show strong inherent characteristics like symmetry and structural features. For example, in the fourth line of Fig. 4, where the bottom half of the face is masked, our ConFill successfully uses visible features to generate a realistic and consistent result. This capacity is also demonstrated in natural scenes, for instance, in the fourth example in Fig. 6, our method accurately restores the structure of a flower, outperforming other methods. Overall, these results indicate the effectiveness of our approach across a variety of contexts. Our framework reconstructs fine-grained textures and edges (local information) while maintaining overall semantic and structural coherence (global information), even in regions with large occlusions.

Quantitative Comparisons: Table 1 reports the quantitative results of ConFill, as well as various baseline methods, across CelebA-HQ, Places2, and ImageNet-1k datasets using different mask types. Here are the notable observations derived from our analysis. Firstly, concerning the evaluation metrics, ConFill consistently surpasses other baseline methods, demonstrating an aver-

age decrease in FID score by 2.1%, 2.62%, and 2.53% compared to the top-performing baseline, CoPaint, across the CelebA-HQ, Places2, and ImageNet-1k datasets, respectively. Furthermore, the images generated by ConFill show improved coherence and outperform in several other respects, including naturalness and complying with the image completion constraint. Moreover, it is important to observe that ConFill performance advantage is particularly notable on the Places2 and ImageNet-1k datasets. This is likely because these datasets contain more complex images, making any imperfections, including incoherence, more obvious. For example, on Places2 with an expand mask, ConFill has 3.4% lower FID score compared to CoPaint. On the other hand, on ImageNet-1k our model has an average LPIPS score of 0.242, which is 4.8% and 8.6% higher than CoPaint and RePaint, respectively. As demonstrated in Figs. 5 and 6, ControlNet and PowerPaint struggle when synthesising images with large holes, whereas ConFill produces realistic images with fewer artefacts than other methods.

Fig. 7 highlights the superior performance of ConFill compared to the state-of-the-art models. ConFill converges faster than baseline models due to CAD, which efficiently guides the diffusion process to produce high-quality samples that align closely with the target distribution in fewer iterations. Notably, after just 50k iterations, ConFill achieves a 62-point improvement in FID on Places2, demonstrating its effectiveness in enhancing image quality and perceptual similarity with greater efficiency.

4.3 Human Perceptual

We conducted a subjective analysis through a human evaluation with different mask ratios. Participants are shown a reference

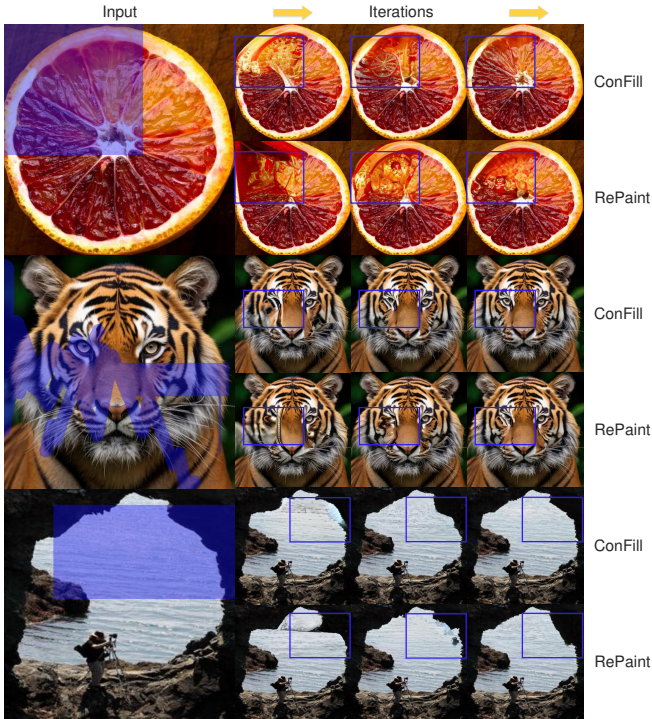


Fig. 10: The image completion results of our ConFill and RePaint across various iterations. Initially, both models generate artifact-laden outputs. However, as iterations progress, ConFill consistently produces sharper and more accurate reconstructions compared to RePaint. For example, our model produces a more realistic reconstruction of the missing fruit segment, with improved color and texture, while RePaint retains artifacts and shows less consistent color matching.

image with masked parts, alongside two completed versions: one by ConFill and the other by a baseline method. They chose the higher-quality image based on predefined criteria, focusing on the coherence of the results. Feedback from 46 participants was collected, evaluating 2,240 image pairs from CelebA-HQ, Places2, and ImageNet datasets. The vote difference, calculated as a percentage, indicates the preference for ConFill over baseline models. In Fig. 8, ConFill was 21% more likely to be selected over CoPaint on Places2 with a mask ratio of 0.8-1.0 and 29% more likely over RePaint on ImageNet with a mask ratio of 0.4-0.6.

4.4 Adaptation to Unseen Mask Types

Based on the guidelines of the NTIRE 2022 Image Inpainting Challenge [73], we developed a test set comprising 5,000 samples, each with an average missing ratio of approximately 60%. This set includes mask types such as Image Expansion and Nearest Neighbour, which were not included during the training phase. As illustrated in Table 2, other state-of-the-art models, particularly DPS, RePaint and CoPaint show performance declines. In contrast, our method shows excellent generalisation ability for these previously unseen mask types. Our method achieved a superior FID score of 19.63, representing a 17.8% improvement over the second-best approach, CoPaint, which scored 23.87. This success is due to our CAD-based denoising mechanism, which iteratively refines the model’s predictions, leading to more accurate image completion. The percentage improvement is calculated by $\left(\frac{\text{Baseline performance} - \text{new performance}}{\text{Baseline performance}}\right) \times 100$. Additionally, we evaluate time efficiency by calculating the average runtime required to sample an image. The results indicate that ConFill achieves faster

TABLE 2: Quantitative analysis and inpainting speed on previously unseen mask types. ConFill achieves a 17.8% improvement in FID, 12.2% in LPIPS, and 4.0% in SSIM compared to the closest competitor, CoPaint.

Method	FID↓	LPIPS↓	SSIM↑	Runtime↓
ConFill (ours)	19.63	0.173	0.752	109.24
CoPaint [27]	23.87	0.197	0.723	117.96
PowerPaint [64]	27.46	0.221	0.704	274.83
RePaint [17]	25.32	0.213	0.708	210.58
ControlNet [63]	27.83	0.235	0.697	223.26
DPS [25]	28.76	0.238	0.685	453.05
I ² SB [61]	29.25	0.241	0.680	178.67
DDRM [55]	43.51	0.306	0.629	162.41




Fig. 11: Ablation analysis using a half-mask on the CelebA and Places2 datasets. Shows the performance of ConFill with and without dynamic sampling (DS) and compares it with the other models. Results show ConFill with DS performs strongly even at lower NFEs.

sampling compared to other methods. The longer sampling times in RePaint is due to its resampling mechanism that needed to synchronize condition and generation, while CoPaint requires an additional computation due to its extra backward pass and the need for more optimization steps. ControlNet introduces conditional control to diffusion, adding moderate computational overhead, while PowerPaint, relies on task prompts, making it more time-intensive. In contrast, ConFill reduces inference time by using an optimized reverse diffusion schedule and a lightweight denoising architecture, which minimizes computational overhead without compromising image quality.

Fig. 9 (left) demonstrates that the baseline models struggle to accurately reconstruct images with extremely large unseen masks, leading to outputs with notable distortions. In contrast, ConFill consistently produces reconstructions of superior quality. Additionally, Fig. 9 (right) illustrates the relationship between FID scores and the mask ratio. While RePaint and CoPaint show FID scores comparable to our model at the minimal mask ratio, ConFill surpasses them by a greater margin as the mask size increases, indicating its effectiveness for unseen, larger missing regions.

4.5 Iterative Performance Analysis

Fig. 10 shows the incremental results across different iterations, indicating that our method reaches promising performance with fewer iterations, demonstrating a faster performance than RePaint. This suggests that our model is not only faster but also more efficient in terms of iteration count needed to reach a refined result, demonstrating its potential for applications that require high-quality image restoration with fewer computational resources.

4.6 Ablation Analysis

We investigate the contribution of CAD to our model by comparing the performance of our ConFill (which uses CAD), with a variant that uses standard WD and L_2 . We also examine the

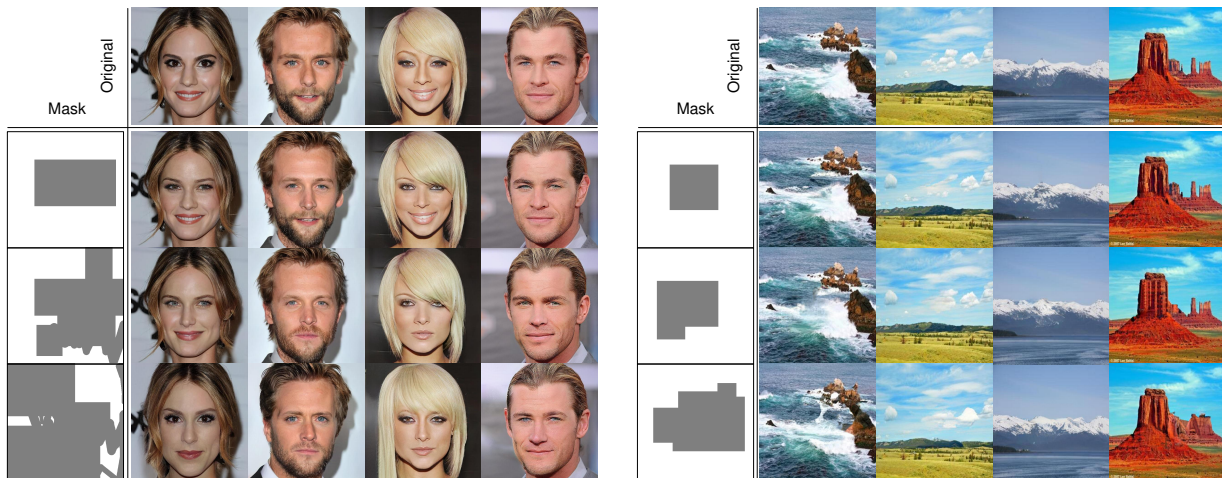


Fig. 12: Image completion results of ConFill for the CelebA and Places2 datasets using various mask types (small to large size).

TABLE 3: Ablation Analysis. Performance evaluation of our model using CAD, standard WD and Euclidean distance with varying gradient descent steps (G). In this evaluation, we use the test set of Places2 with wide mask. The best settings are highlighted.

Method	G	FID↓	LPIPS↓	SSIM↑
ConFill w CAD (ours)	1	10.48	0.129	0.808
	2	10.39	0.120	0.812
	5	10.56	0.133	0.801
ConFill w WD	1	12.35	0.152	0.779
	2	11.97	0.141	0.791
	5	12.43	0.154	0.769
ConFill w $L2$	1	12.28	0.154	0.784
	2	12.32	0.150	0.784
	5	12.61	0.166	0.759

TABLE 4: Impact of time travel interval (\hat{t}) and frequency (\mathcal{J}). The time travel interval of 10 and a frequency of 1 provide the most effective balance between image quality and structural similarity.

Method	\hat{t}	FID↓	LPIPS↓	SSIM↑
ConFill	2	10.48	0.126	0.807
	5	10.44	0.123	0.809
	10	10.39	0.120	0.812
	15	10.41	0.119	0.810
\mathcal{J}		FID↓	LPIPS↓	SSIM↑
ConFill	1	10.39	0.120	0.812
	2	10.40	0.121	0.810
	5	10.42	0.122	0.809

effects of three hyperparameters: gradient descent steps (G), time travel interval (\hat{t}) and its frequency (\mathcal{J}). Our experiments are conducted on Places2 with wide mask and the results are detailed in Table 3 and Table 4. Our model uses G -step gradient descent approach to optimize \hat{x}_t at every timestep. However, as shown in Table 3, increasing G does not improve performance or make the generated images better match the real ones. Because the optimization of \hat{x}_t focuses on minimising the mean square error in the known regions (not masked), as detailed in the second term of Eq. (19). CAD inherently provides smoother and more meaningful gradients during optimization. This makes each gradient descent step more effective at reducing the gap between generated images and real ones while also adapting to local contextual differences. However, a higher number of gradient steps can lead to overfitting. Moreover, the analysis shows that ConFill with WD provides more stable and consistent performance across all gradient steps compared to ConFill with $L2$. While $L2$ performs slightly better at

TABLE 5: Performance comparison on different dynamic sampling configurations. Evaluation on CelebA with an expanded mask.

Implementation	FID↓	LPIPS↓	SSIM↑
Combined sampling	14.27 (+5.2%)	0.164 (+6.9%)	0.760 (+6.3%)
Texture-driven sampling	15.86	0.178	0.753
Edge-driven sampling	16.21	0.182	0.748

TABLE 6: The impact of v , which controls the scaling of the distance function based on the contextual information and τ affects sensitivity to contextual differences, is evaluated using the ImageNet test set with a wide mask. The best settings are highlighted, with improvements (%) shown between the most and least effective settings.

Method	v	FID↓	LPIPS↓	SSIM↑
ConFill	0.05	13.43	0.130	0.764
	0.10	13.24 (+1.4%)	0.121 (+6.9%)	0.785 (+2.7%)
	0.15	13.26	0.122	0.783
τ		FID↓	LPIPS↓	SSIM↑
ConFill	0.01	13.26	0.121	0.786
	0.02	13.24	0.121	0.785
	0.03	13.27	0.120	0.784

the initial step ($G = 1$), its performance degrades as G increases, whereas WD maintains higher stability.

In Table 4, we evaluate the impact of the time travel interval \hat{t} and frequency \mathcal{J} on model performance. The results show that \hat{t} has minimal impact when $\hat{t} \geq 5$. Additionally, unlike RePaint, which uses a time travel frequency of $\mathcal{J} = 9$, our approach performs well with $\mathcal{J} = 1$. This suggests that our model more effectively enforces image completion constraints compared to the simple replacement techniques used in RePaint. To evaluate the role of different criteria for dynamic sampling, we analyze different sampling configurations: I) Texture-driven sampling ($\hat{\alpha} = 1, \hat{\beta} = 0$) where sampling weights are only based on local intensity variance. II) Edge-driven sampling ($\hat{\alpha} = 0, \hat{\beta} = 1$) where sampling weights are only based on edge density to focus on regions with prominent spatial structures. III) Combined sampling ($\hat{\alpha} = 0.5, \hat{\beta} = 0.5$) where sampling weights are computed as a balanced combination of both measures. Table 5 summarizes the results for each configuration. The combined approach improves reconstruction quality, particularly in regions with complex textures.

In Table 6 we evaluate the impact of the parameters v and τ . The results indicate that v has a significant effect on FID

TABLE 7: Performance comparison of contextual feature extractor implementations. Evaluation on ImageNet with an expanded mask.

Implementation	FID↓	LPIPS↓	SSIM↑
Combined features	50.35 (+5.2%)	0.574 (+6.9%)	0.349 (+6.3%)
Handcrafted features	55.34	0.657	0.308
Learned features	52.95	0.614	0.327

TABLE 8: Impact of different sample size on FID, LPIPS, and SSIM metrics for different methods. Evaluation on Places2 with an expanded mask.

Sample size	Method	FID↓	LPIPS↓	SSIM↑
500 images	RePaint	58.72	0.614	0.357
	PowerPaint	56.90	0.598	0.371
	CoPaint	56.37	0.588	0.381
	ConFill	47.12	0.540	0.419
1000 images	RePaint	58.21	0.612	0.359
	PowerPaint	56.34	0.599	0.373
	CoPaint	56.05	0.585	0.383
	ConFill	46.58	0.540	0.420

and LPIPS as it directly scales the distance metric. For instance, adjusting v from 0.05 to 0.10 results in a 1.4% improvement in FID (from 13.43 to 13.24) and a 6.9% improvement in LPIPS (from 0.130 to 0.121). However, values beyond 0.1 do not result in further notable changes. Meanwhile, τ controls the sensitivity of the scaling function to contextual differences. Fig. 11 compares the performance of different methods in terms of LPIPS across varying numbers of function evaluations (NFE) on CelebA-HQ and Places2. ConFill consistently achieves the lowest LPIPS values across all NFEs, demonstrating its effectiveness in minimizing perceived differences regardless of the NFE. ConFill without dynamic sampling (DS) maintains low LPIPS values but is slightly less effective compared to ConFill with DS. RePaint starts with a relatively high LPIPS value but shows significant improvement as NFEs increase. Despite this improvement, RePaint still falls below ConFill. To analyze the impact of $C(\cdot)$ in the image completion process, we evaluate its performance under different configurations. Handcrafted features only, learned features only, and a combination of handcrafted and learned features. Table 7 reports the results, demonstrating that combining handcrafted and learned features consistently outperforms the individual ones in terms of perceptual quality and contextual coherence. The results of our model using various mask types are shown in Fig. 12. ConFill connects different conditioned instances, ranging from small to large completions and generates plausible and coherent content for the missing areas.

In Table 8, we analyse the performance of different models as the evaluation set size increases. Across all metrics, we observe consistent trends. FID decreases with larger evaluation sets, reflecting improved alignment with the data distribution. LPIPS and SSIM remain relatively stable, with minor improvements, indicating the reliability and robustness of the metrics across varying sample sizes. These trends demonstrate that the evaluation protocol effectively captures the performance of the models, providing a comprehensive understanding of their behavior as sample sizes increase.

4.7 Comparison to Non-diffusion Models

To evaluate the performance of ConFill against non-diffusion models, we selected two state-of-the-art methods: LaMa [6] and MAT [9], which are benchmarks in GAN and Transformer-based

TABLE 9: Quantitative analysis with non-diffusion models on Places2 with an expanded mask. Compared to MAT, ConFill shows a 3.4% improvement in FID, and 1.8% in SSIM.

Method	FID↓	LPIPS↓	SSIM↑
ConFill (ours)	48.65 (+3.4%)	0.543	0.417 (+1.8%)
LaMa [6]	55.46	0.558	0.398
MAT [9]	50.31	0.540	0.410

architectures, respectively. By comparing ConFill with these models, we aim to show its strengths, particularly in visual quality and coherence, as shown in Table 9. Specifically, ConFill achieves an FID score of 50.35, outperforming LaMa (56.46) and MAT (52.31), demonstrating a stronger alignment with the original image distribution. Furthermore, in terms of LPIPS, MAT with a small margin performs better than ConFill which is due to its specific training with large masks. However, ConFill achieves a higher SSIM score of 0.417, outperforming the other models. These results highlight ConFill’s robustness in generating visually coherent and high-quality images, especially for large masked regions where maintaining visual coherence is particularly challenging. This is achieved by integrating context-adaptive discrepancy, which effectively balances the trade-offs between reconstruction fidelity and perceptual quality.

5 CONCLUSION

In this paper, we proposed ConFill, a novel approach that adopts the principled application of CAD within a diffusion model framework to address the challenges of image completion tasks. ConFill introduces a novel strategy for generative inverse problems by effectively minimizing the CAD between the latent representations of noisy and original data distributions. This refined approach ensures smooth integration of reconstructed regions with their surrounding context and dynamically adapts to local image features. By iteratively refining the posterior distribution estimation during the denoising process, ConFill substantially reduces errors and imposes strict penalties on discrepancies with the reference image. The integration of CAD enhances this process, allowing the model to adjust its computations based on the contextual differences between image segments, significantly improving the fidelity and coherence of the completed images. Additionally, ConFill incorporates a dynamic sampling strategy that adjusts the sampling process based on contextual information, further improving the efficiency and quality of image completion. Our extensive evaluations on three datasets have demonstrated ConFill’s superiority over current state-of-the-art methods.

ACKNOWLEDGMENT

This project is partially supported by the University of York Research Grant No. M0259102, and the National Research Foundation, Singapore, under its NRF Professorship Award No. NRF-P2024-001. Also supported by the Swedish Research Council (2022-04266) and KAW (DarkTree project; 2024.0076).

REFERENCES


- [1] H. Ling, K. Kreis, D. Li, S. W. Kim, A. Torralba, and S. Fidler, “Editgan: High-precision semantic image editing,” *Proc. Advances Neural Inf. Process. Syst.*, vol. 34, pp. 16331–16345, 2021.
- [2] W. Li, X. Yu, K. Zhou, Y. Song, Z. Lin, and J. Jia, “Image inpainting via iteratively decoupled probabilistic modeling,” *arXiv preprint arXiv:2212.02963*, 2022.

- [3] R. He, J. Cao, L. Song, Z. Sun, and T. Tan, "Adversarial cross-spectral face completion for nir-vis face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 5, pp. 1025–1037, 2020.
- [4] B. Li, X. Liu, P. Hu, Z. Wu, J. Lv, and X. Peng, "All-in-one image restoration for unknown corruption," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 17 452–17 462.
- [5] A. B. Yildirim, V. Baday, E. Erdem, A. Erdem, and A. Dundar, "Inst-inpaint: Instructing to remove objects with diffusion models," *arXiv preprint arXiv:2304.03246*, 2023.
- [6] R. Suvorov, E. Logacheva, A. Mashikhin, A. Remizova, A. Ashukha, A. Silvestrov, N. Kong, H. Goka, K. Park, and V. Lempitsky, "Resolution-robust large mask inpainting with fourier convolutions," in *Proc. Winter Conf. Appl. Comput. Vis.*, 2022, pp. 2149–2159.
- [7] Z. Lu, J. Jiang, J. Huang, G. Wu, and X. Liu, "Glama: Joint spatial and frequency loss for general image inpainting," in *Proc. Conf. Comput. Vis. Pattern Recognit. Workshop*, 2022, pp. 1301–1310.
- [8] Y. Zeng, J. Fu, H. Chao, and B. Guo, "Aggregated contextual transformations for high-resolution image inpainting," *IEEE Trans. Vis. Comput. Graph.*, 2022.
- [9] W. Li, Z. Lin, K. Zhou, L. Qi, Y. Wang, and J. Jia, "Mat: Mask-aware transformer for large hole image inpainting," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 10 758–10 768.
- [10] P. Shamsolmoali, M. Zareapoor, and E. Granger, "Transinpaint: Transformer-based image inpainting with context adaptation," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 849–858.
- [11] Z. Wan, J. Zhang, D. Chen, and J. Liao, "High-fidelity pluralistic image completion with transformers," in *Proc. Int. Conf. Comput. Vis.*, 2021.
- [12] P. Shamsolmoali, M. Zareapoor, H. Zhou, D. Tao, and X. Li, "Vtae: Variational transformer autoencoder with manifolds learning," *IEEE Trans. Image Process.*, 2023.
- [13] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Proc. Advances Neural Inf. Process. Syst.*, vol. 33, pp. 6840–6851, 2020.
- [14] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *Proc. Int. Conf. Learn. Represent.*, 2021.
- [15] Y. Wang, J. Yu, and J. Zhang, "Zero-shot image restoration using denoising diffusion null-space," *Proc. Int. Conf. Learn. Represent.*, 2023.
- [16] O. Avrahami, D. Lischinski, and O. Fried, "Blended diffusion for text-driven editing of natural images," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 18 208–18 218.
- [17] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "Repaint: Inpainting using denoising diffusion probabilistic models," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022.
- [18] B. L. Trippe, J. Yim, D. Tischer, D. Baker, T. Broderick, R. Barzilay, and T. Jaakkola, "Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem," *Proc. Advances Neural Inf. Process. Syst. workshop*, 2022.
- [19] Z. Kong and W. Ping, "On fast sampling of diffusion probabilistic models," *Proc. Int. Conf. Mach. Learn. workshop*, 2021.
- [20] A. Jonnarth and M. Felsberg, "Importance sampling cams for weakly-supervised segmentation," in *Proc. Int. Conf. Acoust. Speech Signal Process.*, 2022, pp. 2639–2643.
- [21] L. Liu, Y. Ren, Z. Lin, and Z. Zhao, "Pseudo numerical methods for diffusion models on manifolds," *Proc. Int. Conf. Learn. Represent.*, 2022.
- [22] A. Y. Tong, G. Hugué, A. Natik, K. MacDonald, M. Kuchroo, R. Coifman, G. Wolf, and S. Krishnaswamy, "Diffusion earth mover's distance and distribution embeddings," in *Proc. Int. Conf. Mach. Learn.*, 2021.
- [23] Z. Li, S. Li, Z. Wang, N. Lei, Z. Luo, and D. X. Gu, "Dpm-ot: a new diffusion probabilistic model based on optimal transport," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 22 624–22 633.
- [24] S. Ozair, C. Lynch, Y. Bengio, A. Van den Oord, S. Levine, and P. Sermanet, "Wasserstein dependency measure for representation learning," *Proc. Advances Neural Inf. Process. Syst.*, vol. 32, 2019.
- [25] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," *Proc. Int. Conf. Learn. Represent.*, 2023.
- [26] H. Chung, B. Sim, D. Ryu, and J. C. Ye, "Improving diffusion models for inverse problems using manifold constraints," *Proc. Advances Neural Inf. Process. Syst.*, vol. 35, pp. 25 683–25 696, 2022.
- [27] G. Zhang, J. Ji, Y. Zhang, M. Yu, T. S. Jaakkola, and S. Chang, "Towards coherent image inpainting using denoising diffusion implicit models," *Proc. Int. Conf. Mach. Learn.*, 2023.
- [28] Y. Lu, L. Chen, A. Saidi, and X. Gu, "Brenier approach for optimal transportation between a quasi-discrete measure and a discrete measure," 2019, pp. 204–212.
- [29] F.-P. Paty, A. d'Aspremont, and M. Cuturi, "Regularity as regularization: Smooth and strongly convex brenier potentials in optimal transport," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 1222–1232.
- [30] N. Papadakis, "Optimal transport for image processing," Ph.D. dissertation, Université de Bordeaux; Habilitation thesis, 2015.
- [31] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, "Patch-match: A randomized correspondence algorithm for structural image editing," *ACM Trans. Graph.*, vol. 28, no. 3, p. 24, 2009.
- [32] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and locally consistent image completion," *ACM Trans. Graph.*, vol. 36, no. 4, 2017.
- [33] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5505–5514.
- [34] H. Zhang, Z. Hu, C. Luo, W. Zuo, and M. Wang, "Semantic image inpainting with progressive generative networks," in *Proc. ACM Int. Conf. Multimedia*, 2018, pp. 1939–1947.
- [35] H. Liu, B. Jiang, Y. Song, W. Huang, and C. Yang, "Rethinking image inpainting via a mutual encoder-decoder with feature equalizations," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 725–741.
- [36] Y. Zeng, Z. Lin, H. Lu, and V. M. Patel, "Cr-fill: Generative image inpainting with auxiliary contextual reconstruction," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14 164–14 173.
- [37] Q. Dong, C. Cao, and Y. Fu, "Incremental transformer structure enhanced image inpainting with masking positional encoding," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 11 358–11 368.
- [38] J. Jain, Y. Zhou, and N. Yu, "Keys to better image inpainting: Structure and texture go hand in hand," in *Winter Conf. Appl. Comput. Vis.*, 2023.
- [39] S. Zhao, J. Cui, Y. Sheng, Y. Dong, X. Liang, E. I. Chang, and Y. Xu, "Large scale image completion via co-modulated generative adversarial networks," *Proc. Int. Conf. Learn. Representations*, 2021.
- [40] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, "Image inpainting for irregular holes using partial convolutions," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 85–100.
- [41] T. Yu, Z. Guo, X. Jin, S. Wu, Z. Chen, W. Li, Z. Zhang, and S. Liu, "Region normalization for image inpainting," in *Proc. AAAI*, 2020.
- [42] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *Proc. Int. Conf. Learn. Representations*, 2020.
- [43] Y. Yu, F. Zhan, R. Wu, J. Pan, K. Cui, S. Lu, F. Ma, X. Xie, and C. Miao, "Diverse image inpainting with bidirectional and autoregressive transformers," in *Proc. Int. Conf. Multimedia*, 2021, pp. 69–78.
- [44] L. Liao, J. Xiao, Z. Wang, C.-W. Lin, and S. Satoh, "Image inpainting guided by coherence priors of semantics and textures," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 6539–6548.
- [45] K. Nazeri, E. Ng, T. Joseph, F. Z. Qureshi, and M. Ebrahimi, "Edgeconnect: Generative image inpainting with adversarial edge learning," *Proc. Int. Conf. Comput. Vis. Workshop*, 2019.
- [46] Z. Yi, Q. Tang, S. Azizi, D. Jang, and Z. Xu, "Contextual residual aggregation for ultra high-resolution image inpainting," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7508–7517.
- [47] S. Bond-Taylor, P. Hessey, H. Sasaki, T. P. Breckon, and C. G. Willcocks, "Unleashing transformers: Parallel token prediction with discrete absorbing diffusion for fast high-resolution image generation from vector-quantized codes," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 170–188.
- [48] H. Chung, B. Sim, and J. C. Ye, "Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022.
- [49] A. Bansal, E. Borgnia, H.-M. Chu, J. S. Li, H. Kazemi, F. Huang, M. Goldblum, J. Geiping, and T. Goldstein, "Cold diffusion: Inverting arbitrary image transforms without noise," *arXiv preprint arXiv:2208.09392*, 2022.
- [50] H. Liu, Y. Wang, M. Wang, and Y. Rui, "Delving globally into texture and structure for image inpainting," in *Proc. ACM Int. Conf. Multimedia*, 2022, pp. 1270–1278.
- [51] X. Su, J. Song, C. Meng, and S. Ermon, "Dual diffusion implicit bridges for image-to-image translation," *Proc. Int. Conf. Learn. Represent.*, 2023.
- [52] Z. Wu, J. Hu, W. Lu, I. Gilitshenskiy, and A. Garg, "Slotdiffusion: Object-centric generative modeling with diffusion models," *Proc. Advances Neural Inf. Process. Syst.*, 2023.
- [53] Q. Zhang, M. Tao, and Y. Chen, "gddim: Generalized denoising diffusion implicit models," *Proc. Int. Conf. Learn. Represent.*, 2023.
- [54] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," *Proc. Advances Neural Inf. Process. Syst.*, 2019.
- [55] B. Kawar, M. Elad, S. Ermon, and J. Song, "Denoising diffusion restoration models," *Proc. Advances Neural Inf. Process. Syst.*, 2022.
- [56] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, "Palette: Image-to-image diffusion models," in *Proc. ACM SIGGRAPH Conf.*, 2022, pp. 1–10.

- [57] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, "Glide: Towards photorealistic image generation and editing with text-guided diffusion models," *Proc. Int. Conf. Learn. Represent.*, 2021.
- [58] S. Xie, Z. Zhang, Z. Lin, T. Hinz, and K. Zhang, "Smartbrush: Text and shape guided object inpainting with diffusion model," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 22 428–22 437.
- [59] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022.
- [60] J. Song, A. Vahdat, M. Mardani, and J. Kautz, "Pseudoinverse-guided diffusion models for inverse problems," in *Proc. Int. Conf. Learn. Represent.*, 2023.
- [61] G.-H. Liu, A. Vahdat, D.-A. Huang, E. A. Theodorou, W. Nie, and A. Anandkumar, "I²sb: Image-to-image schrodinger bridge," *Proc. Int. Conf. Mach. Learn.*, 2023.
- [62] H. Chung, J. Kim, and J. C. Ye, "Direct diffusion bridge using data consistency for inverse problems," *Proc. Advances Neural Inf. Process. Syst.*, vol. 36, 2024.
- [63] L. Zhang, A. Rao, and M. Agrawala, "Adding conditional control to text-to-image diffusion models," in *Proc. Int. Conf. Comput. Vis.*, 2023, pp. 3836–3847.
- [64] J. Zhuang, Y. Zeng, W. Liu, C. Yuan, and K. Chen, "A task is worth one word: Learning with task prompts for high-quality versatile image inpainting," in *Proc. Eur. Conf. Comput. Vis.*, 2025, pp. 195–211.
- [65] A. Pokle, Z. Geng, and J. Z. Kolter, "Deep equilibrium approaches to diffusion models," *Proc. Advances Neural Inf. Process. Syst.*, vol. 35, pp. 37 975–37 990, 2022.
- [66] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Proc. Advances Neural Inf. Process. Syst.*, 2021.
- [67] V. M. Panaretos and Y. Zemel, "Statistical aspects of wasserstein distances," *Annu. Rev. Stat. Appl.*, vol. 6, no. 1, pp. 405–431, 2019.
- [68] P. Mokrov, A. Korotin, L. Li, A. Genevay, J. M. Solomon, and E. Burnaev, "Large-scale wasserstein gradient flows," *Proc. Advances Neural Inf. Process. Syst.*, vol. 34, pp. 15 243–15 256, 2021.
- [69] J. Fan, Q. Zhang, A. Taghvaei, and Y. Chen, "Variational wasserstein gradient flow," *Proc. Int. Conf. Mach. Learn.*, 2022.
- [70] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. Int. Conf. Comput. Vis.*, 2015, pp. 3730–3738.
- [71] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, "Places: A 10 million image database for scene recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 6, pp. 1452–1464, 2017.
- [72] 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," *Int. J. Comput. Vis.*, vol. 115, no. 3, 2015.
- [73] A. Romero, A. Castillo, J. Abril-Nova, R. Timofte, R. Das, S. Hira, Z. Pan, M. Zhang, B. Li, D. He *et al.*, "Ntire 2022 image inpainting challenge: Report," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, 2022.



Pourya Shamsolmoali (Senior Member, IEEE) received the PhD degree in computer science from Shanghai Jiao Tong University. He has been a visiting researcher with Linköping University, INRIA-France, and ETS-Montreal. He is currently a lecturer with the University of York. His main research focuses on machine learning, computer vision, and image processing.



Masoumeh Zareapoor (Member, IEEE) received the PhD degree in computer science from Jamia University, India. She is an Associate Researcher at Shanghai Jiao Tong University and was a visiting researcher at Northwestern Polytechnical University, Xi'an. Additionally, she was a research engineer at Huawei, Shanghai. Her main research focuses on machine learning and computer vision.



Huiyu Zhou received the PhD degree in computer vision from Heriot-Watt University, Edinburgh, U.K. He currently is a full professor with the School of Computing and Mathematical Sciences, University of Leicester, U.K. He has published widely in the field. His research work has been or is being supported by U.K. EPSRC, ESRC, AHRC, MRC, EU, Royal Society, Leverhulme Trust, Puffin Trust, Invest NI and industry.



Michael Felsberg (Senior Member, IEEE) received the PhD degree from Kiel University, Germany, in 2002, and the docent degree from Linköping University, in 2005. He is a full professor with Linköping University, Sweden, since 2008. He received the DAGM Olympus award in 2005 and is fellow of the IAPR and ELLIS. His research interests include video object and instance segmentation, classification, registration of point clouds, as well as efficient machine learning techniques for incremental, few-shot, and long-tailed settings.



Dacheng Tao (IEEE Fellow) is currently a Distinguished University Professor in the College of Computing & Data Science at Nanyang Technological University. He mainly applies statistics and mathematics to artificial intelligence, data science, and his research is detailed in one monograph and over 200 publications in prestigious journals and proceedings at leading conferences, with best paper awards, best student paper awards, and test-of-time awards. His publications have been cited over 112K times and he has an h-index of 160+ in Google Scholar. He received the 2015 and 2020 Australian Eureka Prize, the 2018 IEEE ICDM Research Contributions Award, the 2019 Diploma of The Polish Neural Network Society, and the 2021 IEEE Computer Society McCluskey Technical Achievement Award. He is a Fellow of the Australian Academy of Science, AAAS, ACM and IEEE.



Xuelong Li (Fellow, IEEE) is the CTO and Chief Scientist of China Telecom, where he founded the Institute of Artificial Intelligence (TeleAI) of China Telecom.