

SONIC: Spectral Optimization of Noise for Inpainting with Consistency

Seungyeon Baek¹, Erqun Dong¹, Shadan Namazifard¹, Mark J. Matthews^{2*},
and Kwang Moo Yi¹

¹ University of British Columbia, Canada
{syeonb, eqdong, kmyi}@cs.ubc.ca

shadan82@student.ubc.ca

² Google DeepMind, USA
mjmatthews@google.com

Abstract. We propose a novel training-free method for inpainting with off-the-shelf text-to-image models. While guidance-based methods in theory allow *generic* models to be used for inverse problems such as inpainting — in practice their effectiveness is limited, leading to the necessity of *specialized* inpainting-specific models. In this work, we argue the missing ingredient for training-free generic model usage is proper optimization of the *initial noise sample*. We optimize the initial noise to approximately reproduce the unmasked image, in as few as tens of optimization steps, then use it with a conventional training-free inpainting method. Critically, we propose two core ideas that make this possible: (i) we perform linear approximation that avoids the costly and often impractical unrolling required to relate the initial noise sample to model output—which potentially is why this relationship was previously overlooked; and (ii) perform spectral preconditioning by optimizing the initial noise sample in the spectral domain with Adam, which stabilizes the optimization. We demonstrate our method on various inpainting tasks, outperforming the state of the art. Project website: <https://ubc-vision.github.io/sonic/>

1 Introduction

Denoising diffusion models [20, 51] and more recently flow models [33, 35] have become the go-to solution for various inverse problems in Computer Vision [11]. For example, most modern image inpainters [4, 22] utilize them in one way or another, including those that use pre-trained models with guidance [4, 34, 42], and those that train a conditional inpainting specific model [24, 73]. Diffusion and flow methods have also been used for inverse problems in other domains, such as image super-resolution and deblurring [12, 28, 42].

Among the many, we are similarly drawn to off-the-shelf models [4, 12, 28, 42] for inpainting, due to the versatility afforded by bypassing expensive training [24, 70] and task-specific augmentations that may not always generalize.¹

* Participated in an advisory capacity only.

¹ We show later in experiments that, *e.g.*, BrushNet [24] produces inpainting mask specific results in some cases.




Fig. 1: Teaser – We propose a novel training-free method of inpainting that focuses exclusively on the initial noise sample. **(Top row)** We show the denoising result of an initial noise sample, as we optimize the seed noise using our method. We optimize the seed noise to faithfully regenerate the non-masked regions of the input image, so as to obtain more consistent inpainting results. **(Bottom row)** Inpainting results of competing methods, with our final result on the right.

Unfortunately, the performance of existing training-free methods [4, 12, 28, 42] is not as good as their more specialized counterparts [24]; see Fig. 1.

Here, we argue that the shortcomings of existing training-free methods, which rely on posterior sampling [12, 28] or conditioning [4, 42], are due to an oversight—optimizing the initial noise sample. The concept of optimizing the initial noise sample is not new in itself. Ever since deep generative models have been used to solve inverse problems [8], one of the very first approaches has been to treat generative models as transport functions and find the input condition (the initial noise sample for denoising models) that matches the observed data [8, 67].

Manipulating the initial noise sample for denoising models, however, requires back-propagating through the denoising chain [6]. This can be costly, both in terms of compute and memory requirements. Thus, with larger modern models, this back-propagation is deemed impractical, making way for guidance methods [5, 18, 37] that change the denoising trajectory during runtime. Still, recent works [32, 56] have demonstrated that, already at the initial denoising step, much of the structure of the final denoised outcome is decided. Indeed, in Fig. 2 we show the profound impact that the initial noise sample has on the final inpainted structure. Thus, the initial noise sample *must* be considered.

To consider the initial noise sample while circumventing the back-propagation problem, various methods have been proposed. Some choose the best initial noise sample among many [29, 32], train a network to generate an ideal noise sample [2, 72], or iteratively refine the noise to improve visual qualities of the output [60]. Others formulate the problem as posterior sampling [10, 12, 28, 41, 42, 47], including a *concurrent* work that drops the model Jacobian [46] and performs inversion, starting from a point generated by another inversion method [23].




Fig. 2: initial noise sample already determines structure – We show example outcomes of BLD [4] with StableDiffusion 3.5 [13] using different initial noise samples. Different initial noise denoise to different scene structure even with the *same* prompt, and *also* the final inpainting outcomes. Note how the content within the inpainted region is structurally similar to the denoised initial noise sample.

Fig. 3: Spectral preconditioning is important – We show examples of how the optimized initial noise sample denoises during the optimization process, when optimized to match the non-masked regions in Fig. 2, starting from **seed A**. Naive optimization also guides toward the desired scene composition, but spectral preconditioning via spectral optimization provides a more stable and robust outcome.

In this work, we propose an inpainting method that *optimizes* the initial noise sample with respect to the final generation outcome, based on two key ideas that allow *practical* and *stable* optimization.

Practical optimization through linearization. To avoid back-propagating through the denoising chain, we propose to *linearize* the entire denoising (flow) trajectory. A naive back-propagation through the denoising chain quickly becomes impractical due to memory and compute requirements. Instead, we approximate the entire denoising trajectory as a linear path, allowing us to optimize the initial noise sample *without back-propagating through the denoiser*.

Stable optimization through spectral preconditioning. Our second contribution is the spectral preconditioning of this optimization that allows for stable convergence. Various works on spectral and time-based guidance scaling have hinted at this [15, 48, 68]—that different frequencies in the latent space have different preferences for the rate of change. We thus precondition the optimization so that, effectively, each frequency can converge at its preferred pace. Implementing this can be effectively done by moving the optimization into the spectral domain, together with the Adam [31] optimizer that would divide each parameter (now each frequency) by their second moments. This produces significantly more stable optimization and convergence; see Fig. 3.

Implementation of this idea, however, requires care. We find that a naive implementation can lead to corruption of the initial noise sample, moving it away from the manifold of acceptable samples. To prevent this, we constrain gradient updates to observed data points only, *i.e.*, the unmasked pixels in inpainting. We further find that the robustness of our method allows it to operate purely

in latent space with simple nearest-neighbor fills for the masked regions when encoding masked images to the latent space.

To demonstrate the effectiveness of our method, we evaluate our method on three standard inpainting datasets: FFHQ [25], DIV2K [1], and BrushBench [24], each with different types of masks. On all three datasets, we outperform the state of the art in SSIM [59] and LPIPS [71] perceptual metrics, and Fréchet Inception Distance (FID) [19]. Our method is also often preferable in human preference scores such as HPS v2 [61] in terms of aesthetics.

2 Related Work

While various methods have been proposed to inpaint images, recent works have focused on utilizing the powerful generative prior of diffusion models [30, 39, 57, 62, 63, 65, 66, 69]. We provide a brief review of works that utilize diffusion models, both training-free (ones that utilize generic models) and those that are trained or fine-tuned to be specialized for inpainting; for a more comprehensive review, we refer the reader to [22]. We also review recent works that focus on the importance of the initial noise sample and its effect on the denoising process, and methods that aim to bypass back-propagation.

Training-free methods. One of the key benefits of denoising diffusion models [20, 51] is their ability to be ‘guided’ by additional conditioning information [5, 18, 21]. The robustness of diffusion models under different manipulations naturally enables their use for inpainting. Some methods [4, 36] *blend* the noise within the masked region with the unmasked ground-truth noised image to condition the generation process. Others extend the Markov chain settings to be conditional [27] or perform range-null space decomposition [58]. Some more recent methods aim to sample from the posterior [10, 12, 28, 41, 42, 47], such that the observed data points can be treated as the likelihood, while the generic diffusion model serves as the prior; other recent works focus on manipulating the geometry of the sampling trajectory [50].

A notable recent method along these lines is FLAIR [12], which formulates a variational method for posterior sampling. Interestingly, in their work, the final algorithm involves keeping track of *both* the final reconstruction and the *initial noise*. Another very recent work, DiNG [41], performs guidance with closed-form solution for the posterior, evaluated through auxiliary random samples.

While these models all have shown improved performance beyond the original latent blending work [36], they still suffer from the same problem—their inpainting results often have an image structure that is inconsistent with the target image; see Fig. 2. Methods that manipulate the initial noise sample [12, 41] suffer less from this, but as we show later in Sec. 4 with FLAIR [12], the effect is still present. Thus, a proper intervention on the initial noise sample is important.

Importance of initial noise sample. Recent works [2, 32, 38, 56, 72] have revealed how much influence the initial noise has on the final generation. Li *et al.* [32] demonstrate that the initial seed largely determines the general layout of the generated image, and mine seeds that are more likely to produce better

generations. Similarly, Ahn *et al.* [2] train a neural network to modify the initial noise sample to generate those that lead to better generations. Lyu *et al.* [38] generate a semantically meaningful initial seed with an automatic refinement pipeline. Very recently, an in-depth study [56] using Principal Component Analysis (PCA) revealed that the initial noise can reliably predict image layout. All of these works hint at the same thing—that the initial noise sample must be considered if we are to match the structure of the image during inpainting.

In theory, this initial noise sample can be directly optimized to achieve a certain goal, which conventional methods have attempted with Generative Adversarial Nets [8, 67], denoising diffusion models [17, 26, 53, 55], and flow models [6, 14, 16]. These methods, however, require back-propagating through the entire network, which is only suitable for smaller models or single-step generators.

Circumventing back-propagation. With the back-propagation being costly, common workarounds include using methods such as the Straight-Through Estimator [7], adopted by many zero-order optimization works [40, 54]. In the context of inpainting, recent works such as FlowChef [42] and DiNG [41] also suggest adopting the straight-through strategy, *i.e.*, dropping the vector-Jacobian product, although they aim towards optimizing the intermediate noise estimates, not the initial noise sample. A concurrent work [46] also drops the vector-Jacobian and optimizes for the initial noise sample, but as we will show later (Sec. 4.3), this alone does not provide stable optimization—our spectral preconditioning is required. This may be why their method starts from an inverted noise from a conventional method [23] before beginning the optimization, and requires a carefully tuned learning rate—our method requires none of these.

Trained methods. Various methods have been proposed that *train* or *fine-tune* existing models. Effectively, these models circumvent the initial noise problem by training inpainting models that ignore the structure imposed by the initial noise sample. Common strategies include providing an inpainting mask and the encoded masked images as additional input and finetuning a pretrained model [45] with randomly masked images [52]. These can include having inpainting as a part of a multi-task adaptation via fine-tuning and prompt learning [73]. Some separate masked-image features and noisy latents to make the task easier to learn [24], some utilize semantics [9], and others train inpainting adapters [10, 62]. These methods, however, require additional training which can be costly, and as we show empirically later in Sec. 4, can fail to generalize for mask types not seen during training.

3 Method

3.1 Overview

An overview of our method is provided in Fig. 4. Our method focuses solely on finding the initial noise sample that best fits a partial (masked) observation. Our key innovation is a linearization strategy that allows optimization of the




Fig. 4: Method overview – We optimize the initial noise sample in the spectral domain $\mathbf{X}_T$, starting from a random noise $\mathbf{x}_T$, such that our denoised latent matches the masked observation $\mathbf{y}$ in the latent space. To allow partial observations to be encoded, we use nearest-pixel filling before passing it into the encoder. We then compute the masked mean square error in the latent space, comparing it with a *fully denoised* latent and update $\mathbf{X}_T$ accordingly. Importantly, we linearize the entire T step denoising process, essentially disconnecting the gradient flow passing through it. This allows us to optimize the initial noise sample $\mathbf{X}_T$ *without back-propagating through the denoiser*.

initial noise *without back-propagating through the denoiser*. We first formalize the problem, then explain our linear approximation. We then introduce our spectral optimization, as well as subtle yet important implementation details.

Formalization. Let $\mathbf{y}$ be a corrupted (masked) image and $\mathbf{A}$ be the corruption operator (mask), often denoted as a matrix in a slight abuse of notation. Let $\mathbf{x}_T$ be the initial noise of a T -step denoising process, such that we can write the denoising process as $\mathbf{x}_0 = \mathcal{D}_T(\mathbf{x}_T)$, where $\mathcal{D}_T$ is the T -step denoiser (*e.g.*, a flow model [13]) and $\mathbf{x}_0$ is the fully denoised image. We aim to find the initial noise $\mathbf{x}_T^*$ such that, when denoised $\mathcal{D}_T(\mathbf{x}_T^*)$, then corrupted, it matches the observed image $\mathbf{y}$. We formulate this as an optimization problem in a least-squares sense:

$$\mathbf{x}_T^* = \arg \min_{\mathbf{x}_T} \|\mathbf{y} - \mathbf{A}\mathcal{D}_T(\mathbf{x}_T)\|_2^2 \quad (1)$$

In this work, we focus on inpainting, thus $\mathbf{A}$ is simply a binary mask. Note that, contrary to other methods, we do not alter the denoising process itself, but *only the initial noise* $\mathbf{x}_T$. While a noise that exactly satisfies $\mathbf{y} = \mathbf{A}\mathcal{D}_T(\mathbf{x}_T^*)$ may not exist, our method still provides an *enhanced* starting point for *any* denoising process.

Naively optimizing $\mathbf{x}_T$ would mean back-propagating through $\mathcal{D}_T(\cdot)$, an expensive iterative process requiring unrolling the entire sequence, costly both in memory and compute—in fact, with an NVidia RTX 5090, back-propagating through more than one denoising step for StableDiffusion 3.5 [13] is not feasible. We thus propose a linear approximation to circumvent this issue.

3.2 Optimizing without heavy back-propagation

Denoising trajectory linearization. We approximate the denoising trajectory $\mathbf{x}(t)$, with a linear equation of the form:

$$\mathbf{x}(t) \approx \hat{\mathbf{x}}(t) = [\mathcal{D}_T(\mathbf{x}_T) - \mathbf{x}_T]_{\text{sg}} \left(1 - \frac{t}{T}\right) + \mathbf{x}_T, \quad (2)$$

where $t \in [0, T]$ and $[\cdot]_{\text{sg}}$ is the stop-gradient operator, implying that we consider the term $\mathcal{D}_T(\mathbf{x}_T) - \mathbf{x}_T$ to be constant throughout the denoising process. This is a reasonable assumption, given that modern flow models [33, 35] are generally trained to produce linear trajectories.²

Optimization objective in the spatial domain. Given this linearization, we can now replace $\mathcal{D}_T(\mathbf{x}_T)$ in Eq. (1) with $\hat{\mathbf{x}}(0)$ from Eq. (2), and formulate a loss that optimizes for $\mathbf{x}_T^*$:

$$\mathcal{L}_{\text{linear}} = \left\| \mathbf{y} - \mathbf{A} \left([\mathcal{D}_T(\mathbf{x}_T) - \mathbf{x}_T]_{\text{sg}} + \mathbf{x}_T \right) \right\|_2^2. \quad (3)$$

Note that Eq. (3) is differentiable with respect to $\mathbf{x}_T$, *without* any need for back-propagation through $\mathcal{D}_T(\cdot)$. Eq. (3) is, in fact, the straight-through estimator [7] that drops Jacobians, and a concurrent work [46] that directly drops Jacobians also arrives at Eq. (3). Direct application, however, as we show in Fig. 3 and in Sec. 4.3, does not produce stable optimization without the spectral optimization introduced next in Sec. 3.3. The concurrent work [46] opts for a carefully tuned learning rate, which is not necessary for our method.

3.3 Optimizing in the spectral domain

When directly optimizing for $\mathbf{x}_T^*$ via Eq. (3) we observe what appear to be regional instabilities in the *spectra* of $\mathbf{x}_T$ —denoised image flickering between optimization steps, showing different levels of detail, *i.e.*, *spatial frequencies* converging at different paces. A recent work [16] that optimizes initial noise without linearization, back-propagating through individual steps, also notes instability and opts for careful control of optimization steps. A lower learning rate can avoid this instability, but drastically increases compute expense and leads to under-convergence, producing blurry results. These observations, combined with recent works that treat different frequencies of the latent space differently during guidance [15, 48, 68], motivate us to precondition the optimization based on frequency.

Spectral preconditioning with Adam [31]. Spectral preconditioning can be achieved by moving the optimization into the spectral domain and using the Adam optimizer [31]. Instead of treating $\mathbf{x}_T$ as a trainable parameter, we optimize its spectral representation $\mathbf{X}_T$, defined as the Fourier transform of $\mathbf{x}_T$, *i.e.*, $\mathbf{x}_T = \mathcal{F}^{-1}(\mathbf{X}_T)$. Then we use the Adam optimizer, which divides each

² We provide an illustrative example in **Supplementary Material**.

dimension—now frequencies in the spectral domain—by their second moment, achieving *spectrally preconditioned* optimization. This spectral preconditioning allows more stable convergence behavior when optimizing, as shown in Fig. 5. While this may seem like a small change, it has a significant impact on the quality of inpainting outcomes—to a degree where spatial optimization often results in inpainting failures. We empirically ablate this choice in Sec. 4.3 and provide further theoretical justification in the **Supplementary Material**.

Thus, our final optimization objective is:

$$\mathcal{L}_{\text{spectral}} = \left\| \mathbf{y} - \mathbf{A} \left([\mathcal{D}_T(\mathbf{x}_T) - \mathcal{F}^{-1}(\mathbf{X}_T)]_{\text{sg}} + \mathcal{F}^{-1}(\mathbf{X}_T) \right) \right\|_2^2, \quad (4)$$

which we optimize for $\mathbf{X}_T$ with Adam [31].

3.4 Implementation details

Spatially constraining the optimization. A critical component to consider is that our optimized initial noise sample must remain within the manifold of acceptable initial noise samples. A naive implementation of our method, even with spectral optimization, can still lead to corruption because of the continuous frequency formulation of the optimization target—where a single frequency component affects all pixels. Thus, the latents corresponding to *unobserved* regions, without observations to guide them, can easily diverge away from the Gaussian distribution that the denoising model expects the initial noise to follow.

We thus opt to *freeze* latents in these regions, as they *already* follow a Gaussian distribution, by masking their gradient updates. As shown in Fig. 6, without such treatment, optimized noise may introduce visible artifacts, especially near the border.

Masking in the latent space. Another computational bottleneck, aside from back-propagation through the denoiser or the flow estimation, is the Variational Auto-Encoder (VAE) that maps images to the latent space. Thanks to the robustness of our method, we find that we can safely perform all operations in the latent space, by simply infilling masked region pixels with their nearest-neighbor color; see Fig. 4. While this does not look like a proper inpainted image, we empirically find this method sufficient. We compare this simple infill strategy against the theoretical *upper-bound* of ground-truth pixels in Sec. 4.3.

Other details. We implement our method with StableDiffusion 3.5 [13], and with BLD-SD3.5 [4, 13] as the inpainter starting from our optimized initial noise sample. To optimize the initial noise sample, we use the Adam [31] optimizer with a learning rate of 3.0, and with default parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We use a classifier-free guidance [21] scale of 2.0, following FLAIR [12]. We use an image resolution of 1024×1024 for all our experiments to match the expected input size of SD 3.5. We also use $T = 20$ for our denoiser.




Fig. 5: Spectral preconditioning – We plot the average convergence trajectory in terms of root mean-squared error (RMSE) when inverting 100 randomly generated images starting from a random noise sample (see Sec. 4.3 for details). Using the Adam optimizer in the spectral domain, *i.e.*, with our spectral conditioning, converges well regardless of the learning rate, while Adam in the spatial domain performs significantly worse. SGD performs the same regardless of the domain, as simply moving to the spectral domain, for SGD, is mathematically identical to being in the spatial domain.




Fig. 6: Spatial constraining – We show how the initial noise sample denoises, similarly to Fig. 3, but with and without the spatial constraint. As shown, without spatial constraining, the optimization diverges from the initial noise sample manifold, causing the visible colour shifts, artifacts, and mask borders as optimization progresses.

4 Results

We first discuss our experimental setup, then present our results, and finally ablate our design choices.

4.1 Experimental setup

Datasets and masks. We evaluate our method on three standard datasets, using different types of masks on each.

- **FFHQ [25]**: A standard dataset often used for evaluating image inpainting methods, containing 70,000 high-quality face photographs with diverse age, pose, and accessories. We follow the protocol in FLAIR [12], using the first 1000 images and applying a large, rectangular mask with a fixed size and position that covers approximately half of the subject’s face.
- **DIV2K [1]**: A set of 1,000 2K high-resolution natural images. Again, we follow the protocol of FLAIR [12] by taking the 800 images from the training set, and center-cropping a square region from the image. Then, we apply the same masking pattern to all samples, consisting of six randomly positioned rectangular regions.³
- **BrushBench [24]**: 600 natural and artificial images with human-annotated masks based on segmentation and captions. We use the provided inpaint segmentation masks and prompts.

³ We use the exact same locations as FLAIR [12].

Prompts. As diffusion/flow-based inpainting methods require prompts, we systematically generate prompts through a Vision Language Model (VLM) [3] for FFHQ [25] and DIV2K [1].

We provide the masked images and prompt the VLM to infer the contents of the image and describe them in two to three concise sentences. The exact prompt used for the VLM can be found in **Supplementary Material**. We use the generated prompts for *all* methods, including ours. For BrushBench [24] we simply use the standard dataset-provided prompts.

Baselines. We compare our method against the following baselines:

- **BLD-SD3.5** [4, 13]: Adopted by diffusers [43] as the default inpainting pipeline, it is a training-free method that ‘blends’ the denoising process in the masked region with the ground-truth denoising trajectory of the unmasked region.
- **FlowChef** [42]: Improves trajectories and stability for flow/diffusion models when solving inverse problems.
- **FlowDPS** [28]: A flow-based posterior sampling method for inverse problems with explicit guidance. Following FLAIR [12], we implement an inpainting solver for this method.
- **FLAIR** [12]: A recent training-free inverse problem solver that uses a variational formula involving the update of both initial noise and final estimates.
- **BrushNet** [24]: A trained method that specializes in inpainting. We use the Stable Diffusion 1.5 [45] checkpoint from the official implementation, and their standard classifier-free guidance scale of 7.5. As BrushNet offers two model subtypes: one trained with random masks, and another with segmentation masks. We use the former for FFHQ [25] and DIV2K [1] while using the latter for BrushBench [24].

For fairness, as in FLAIR [12], we use the *same* StableDiffusion 3.5 [13] (‘medium’ variant) base model for training-free methods (BLD-SD3.5 [4, 13], FlowChef [42], FlowDPS [28], FLAIR [12], and ours). We set the classifier-free guidance scale to 2.0 for these methods, as in FLAIR [12]. Additionally, as the baselines generate varying resolution outputs, we process images at their native resolution (varying from 512×512 to 768×768) and resize to 1024×1024 for evaluation.

For each posterior baseline, we report performance using their suggested number of function evaluations (NFE), and performance at 400 NFE, which is what we use for our method. Note, however, that as our method does not back-propagate through the denoiser, our runtime is comparable to methods with a lower NFE—taking approximately a minute to inpaint one image on a GeForce RTX 5090 GPU, roughly the same time as FLAIR [12] using 50 NFE. We include the wall clock time table in the **Supplementary Material**.

Evaluation metrics. To quantify the quality of the inpainting results, we use the standard metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [59], Learned Perceptual Image Patch Similarity (LPIPS) [71], Fréchet Inception Distance (FID) [19], as well as CLIP score (Contrastive Language-Image Pretraining) [44]. We also use metrics that are designed to




Fig. 7: Qualitative results – We show results across different methods and datasets. We denote the masked areas with black rectangles on the ground truth images. Our method provides the best inpainting results, especially aligning well with the layout and also the lighting conditions of the input image. Zoom in for better detail.

mimic human preference: Image Reward (IR) [64], Human Preference Score v2 (HPS v2) [61], and Aesthetic Score (AS) [49].

For the task of inpainting, where many different answers can exist, we warn that the PSNR metric should be considered with care—*lower PSNR does not necessarily mean worse performance*, as also discussed in [12]. Instead, all metrics should be considered altogether. Still, if a few metrics were to be weighed strongly, we argue that LPIPS and FID are the *most important* for evaluating inpainting quality as they evaluate how similar the inpainted content is to the ground truth via ‘perceived’ similarity. This tolerates minor differences in exact structure and color, which are acceptable in the context of inpainting. CLIP score gauges how well the inpainting results match the original content *semantically*, but should also be read carefully, as the inpainted content may be *geometrically* misaligned with the image. IR, HPS v2, and AS are also imperfect metrics, as they are designed to measure image quality, not inpainting quality. We thus further conduct a single-blind two-alternative choice study with $N=37$ raters, where our method is preferred in 90% of comparisons. We include detailed results in the **Supplementary Material**.

4.2 Results

We present a qualitative comparison in Fig. 7. We compare our method against FLAIR [12], BrushNet [24], and BLD-SD3.5 [4, 13]. We provide additional qual-

Table 1: Quantitative results on FFHQ [25] – Our method outperforms all other methods in terms of SSIM, LPIPS, and FID by a large margin. Our method also performs best in terms of simulated human preference metrics, IR and HPS v2, and second best in AS. Note that while PSNR is lower, this does not mean our method provides worse results, as for inpainting, an image different from ground-truth may still be correct. FlowDPS [28] performs best in terms of CLIP scores, but notice the significantly higher LPIPS and FID scores, indicating poor inpainting quality.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FID $\downarrow$	IR $\uparrow$	HPS v2 $\uparrow$	CLIP $\uparrow$	AS $\uparrow$
BrushNet [24]	21.929	0.759	0.237	17.950	0.275	0.244	24.804	5.713
BLD-SD3.5 [4, 13]	17.592	0.824	0.180	25.842	0.038	0.213	22.456	5.136
FLAIR [12]	23.100	0.823	0.266	18.982	0.220	0.240	24.122	5.735
FLAIR [12] (400 NFEs)	24.099	0.829	0.292	19.873	0.154	0.234	23.765	5.537
FlowChef [42]	19.624	0.828	0.193	20.706	-0.021	0.214	22.178	5.074
FlowChef [42] (400 NFEs)	19.336	0.824	0.197	21.191	-0.015	0.214	22.066	5.029
FlowDPS [28]	20.215	0.760	0.374	48.333	0.025	0.234	25.076	5.792
FlowDPS [28] (400 NFEs)	23.243	0.786	0.389	68.644	0.008	0.232	25.014	5.924
Our method	22.562	0.857	0.121	12.883	0.280	0.245	24.042	5.848

itative results in the **Supplementary Material**. We also provide per-dataset quantitative summaries in Tabs. 1 to 3.

FFHQ [25] – Fig. 7 and Tab. 1. As seen in Fig. 7, our method significantly outperforms other baselines, followed by FLAIR [12]. BrushNet [24] is great at hiding seams, as it is a trained method, but the content it inpaints can misalign with the full image. In the case of BLD-SD3.5 [4, 13], it completely fails to align the global structure, resulting in an inpainting failure. These findings are also reflected in the quantitative metrics presented in Tab. 1, with our method outperforming all other methods in terms of SSIM, LPIPS, FID, IR, and HPS v2. In terms of FID, our method outperforms all other methods by a large margin, with BrushNet [24] being the second best and FLAIR [12] being the third best. As BrushNet [24] does not have seams, it shows up favorably in FID, and other aesthetic metrics. But in terms of SSIM and LPIPS, where structure is considered, the method performs significantly worse. While we perform worse in terms of PSNR and CLIP score, we emphasize again that these in isolation *do not* directly correlate with inpainting performance.

DIV2K [1] – Fig. 7 and Tab. 2. Our method performs best for this dataset as well. FLAIR [12] struggles and produces blurry patches, which may be why it shows high PSNR in Tab. 2, but lower SSIM and LPIPS scores. BrushNet [24], in this example, produces caption-like inpainting results in horizontal stripe masks, hinting that its inpainting outcomes are strongly correlated with the shape of the mask. The quantitative results in Tab. 2 further show how this method performs significantly worse than other methods for this dataset. BLD-SD3.5 [4, 13] does not show apparent artifacts, but as can be seen in the square mask in the bottom right, or the vertical mask on the left, its inpainting results do not align well with the full image. Again, for this dataset, our method outperforms all other methods in terms of SSIM, LPIPS, FID, and HPS v2.

BrushBench [24] – Fig. 7 and Tab. 3. Again, our method performs best. FLAIR [12], for this dataset, is unable to match the global structure of the image

Table 2: Quantitative results on DIV2K [1] – Our method performs best in terms of SSIM, LPIPS, and FID by a large margin. While FLAIR [12] shows high PSNR, as shown in Fig. 7, this could be attributed to the fact that it produces blurry patches. FLAIR [12] also has high IR and AS, and BLD-SD3.5 [4,13] and BrushNet [24] show higher CLIP scores, but they all perform significantly worse in terms of LPIPS and FID, indicating poor inpainting quality.

Method	PSNR↑	SSIM↑	LPIPS↓	FID ↓	IR ↑	HPS v2↑	CLIP↑	AS↑
BrushNet [24]	19.414	0.575	0.308	29.004	0.255	0.249	27.666	5.406
BLD-SD3.5 [4, 13]	19.768	0.789	0.144	25.328	0.389	0.255	27.752	5.883
FLAIR [12]	23.128	0.760	0.298	22.421	0.290	0.248	27.586	5.909
FLAIR [12] (400 NFEs)	23.250	0.768	0.294	22.147	0.339	0.251	27.555	5.929
FlowChef [42]	20.067	0.748	0.237	30.481	0.169	0.240	27.277	5.062
FlowChef [42] (400 NFEs)	19.938	0.747	0.238	30.769	0.177	0.241	27.322	5.084
FlowDPS [28]	20.596	0.628	0.424	46.611	0.179	0.237	27.583	5.802
FlowDPS [28] (400 NFEs)	21.728	0.650	0.420	67.893	0.041	0.231	27.522	5.514
Our method	21.284	0.803	0.115	17.699	0.309	0.253	27.613	5.927

Table 3: Quantitative results on BrushBench [24] – Our method outperforms all baselines in terms of SSIM, LPIPS, and FID. For other metrics, ours performs third-best. Unsurprisingly, BrushNet [24] performs second best in terms of FID, and best in terms of IR, HPS v2, and AS, as this dataset was developed by BrushNet [24]. But in terms of LPIPS and FID, they are worse than ours, in line with Fig. 7. FLAIR [12] again shows high PSNR and SSIM, but performs worse in all other metrics.

Method	PSNR↑	SSIM↑	LPIPS↓	FID ↓	IR ↑	HPS v2↑	CLIP↑	AS↑
BrushNet [24]	18.668	0.740	0.199	50.082	1.246	0.271	26.551	6.366
BLD-SD3.5 [4, 13]	18.234	0.854	0.169	51.661	1.211	0.267	27.262	6.072
FLAIR [12]	19.986	0.855	0.201	51.644	1.007	0.255	26.655	5.978
FLAIR [12] (400 NFEs)	20.154	0.855	0.190	51.535	1.025	0.255	26.455	5.878
FlowChef [42]	18.711	0.856	0.161	53.568	0.975	0.253	26.680	5.826
FlowChef [42] (400 NFEs)	18.688	0.853	0.161	52.281	1.025	0.256	26.767	5.877
FlowDPS [28]	18.269	0.775	0.360	67.731	0.895	0.248	27.045	5.909
FlowDPS [28] (400 NFEs)	18.558	0.795	0.338	80.544	0.832	0.241	26.742	5.813
Our Method	19.014	0.861	0.153	48.072	1.188	0.263	26.890	6.061

well and often produces a disjoint inpainting outcome. BrushNet [24], which is the method developed for this dataset, provides results that seamlessly align with the local surroundings, but often misaligns the rough global structure. BLD-SD3.5 [4,13], for this example, produces arguably the second best results, which aligns with the FID-based evaluation in Tab. 3. Our method, quantitatively, outperforms all methods in terms of SSIM, LPIPS, FID, and performs third-best in terms of other metrics. We emphasize again, that each metric alone does not paint a complete picture of the inpainting quality, and all metrics should be considered together. Among them, we believe LPIPS and FID *best* represent the performance of each method, as they compare the perceived similarity between the inpainting results and the ground truth, tolerating minor differences in exact structure and color.

4.3 Ablation study

We perform ablation studies to motivate our design choices. We summarize the quantitative results in Tab. 4, where we compare the inpainting performance

Table 4: Ablation study – We report the evaluation metrics for FFHQ [25] dataset with variants of our method: using spatial optimization, *i.e.*, without spectral preconditioning; removing spatial constraining of gradient updates; and using ground-truth image as input to the encoder, which shows the upper-bound performance of operating in the latent space. Both spectral preconditioning and spatial constraining are critical. Using ground-truth image for the encoder shows similar performance as our simple nearest-neighbor fill-ins, showing that this simple strategy is sufficient.

Variant	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FID $\downarrow$	IR $\uparrow$	HPS v2 $\uparrow$	CLIP $\uparrow$	AS $\uparrow$
w/o spectral preconditioning	22.120	0.853	0.127	14.579	0.216	0.239	23.654	5.760
w/o spatial constraining	18.863	0.833	0.165	33.269	-0.259	0.223	22.676	5.365
Optimizing with the GT	22.569	0.856	0.121	12.773	0.280	0.245	24.074	5.850
Our method	22.562	0.857	0.121	12.883	0.280	0.245	24.042	5.848

without spectral conditioning and spatial constraining. When removing spectral conditioning, *i.e.*, optimizing in the spatial domain, we use a learning rate of 0.05 for Adam [31]. Observe that both spectral preconditioning and spatial constraining are critical to the performance of the method. Also, optimizing with the nearest-neighbor fill-in produces similar metrics as optimizing with the ground truth image, which would be the upper-bound performance. This confirms the robustness of our method and that the simple fill-in strategy is sufficient.

For further analysis on the convergence enhancement of spectral preconditioning, we synthesize images whom we *know* their initial noise sample of, and plot their convergence in Fig. 5. Specifically, we use the first 100 prompts from DIV2K [1], with various learning rates, using StableDiffusion 3.5 [13] with CFG=2.0, T=20. For all methods, we use the same random starting point, drawn from $\mathcal{N}(0, 1)$. As shown, Adam [31] in spectral domain converges significantly faster, allows larger learning rates thanks to the preconditioning. Adam in spatial domain converges much slower, and becomes unstable if too large learning rates are used. As Fourier transform is linear, SGD in spectral domain has no impact on optimization—they both perform worse.

More ablations. In the **Supplementary Material**, we provide qualitative comparisons with randomly selected samples for objective comparison; and ablation studies on the number of optimization steps used, which show diminishing gains as we increase the number of optimization iterations.

5 Conclusion

We have presented a novel training-free method for inpainting that optimizes the initial noise sample. We have shown that by finding the right initial noise with our method, we can achieve inpainting results that outperform the state of the art. The key ideas that enable this are the linear approximation strategy that allows optimization of the initial noise without back-propagation through the denoiser, and spectral preconditioning, realized by optimizing in the spectral domain, ensuring fast and stable convergence. Spatially constraining the updates further ensures that the optimization does not stray from the initial noise sample manifold. We conduct experiments on three standard inpainting datasets

FFHQ [25], DIV2K [1], and BrushBench [24], demonstrating the effectiveness of our method.

Limitations and future work. While we have demonstrated our method purely with an inpainting task, we believe our method of linearization and spectral preconditioning could, in theory, be applied to other inverse problems such as super-resolution, deblurring, and video inpainting.⁴ We hope our work can serve as a starting point for future research that directly utilizes generic flow models for other inverse problems.

Acknowledgements

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, NSERC Alliance Advantage Grant, Google, the Digital Research Alliance of Canada, and the Advanced Research Computing and the Computer Science department at the University of British Columbia.

References

1. Agustsson, E., Timofte, R.: NTIRE 2017 challenge on single image super-resolution: Dataset and study. In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh. (2017)
2. Ahn, D., Kang, J., Lee, S., Min, J., Kim, M., Jang, W., Cho, H., Paul, S., Kim, S., Cha, E., Jin, K.H., Kim, S.: A noise is worth diffusion guidance. In: The Fourteenth International Conference on Learning Representations (2026)
3. Anthropic: Claude sonnet 4.5. <https://www.anthropic.com/news/claude-sonnet-4-5> (2025), accessed: 2026-06-30
4. Avrahami, O., Fried, O., Lischinski, D.: Blended latent diffusion. ACM Trans. Graph. **42**(4) (Jul 2023)
5. Bansal, A., Chu, H.M., Schwarzschild, A., Sengupta, S., Goldblum, M., Geiping, J., Goldstein, T.: Universal guidance for diffusion models. In: Int. Conf. Learn. Represent. (2024)
6. Ben-Hamu, H., Puny, O., Gat, I., Karrer, B., Singer, U., Lipman, Y.: D-flow: Differentiating through flows for controlled generation. In: Proceedings of the 41st International Conference on Machine Learning (2024)
7. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432 (2013)
8. Bora, A., Jalal, A., Price, E., Dimakis, A.G.: Compressed sensing using generative models. In: International conference on machine learning. pp. 537–546. PMLR (2017)
9. Chen, Y., Chen, J., Pan, Y., Li, Y., Yao, T., Chen, Z., Mei, T.: Improving text-guided object inpainting with semantic pre-inpainting. In: Eur. Conf. Comput. Vis. (2024)

⁴ We show preliminary video inpainting results in the **Supplementary Material**.

10. Corneanu, C., Gadde, R., Martinez, A.M.: Latentpaint: Image inpainting in latent space with diffusion models. In: IEEE Winter Conf. Appl. Comput. Vis. (2024)
11. Daras, G., Chung, H., Lai, C.H., Mitsufuji, Y., Ye, J.C., Milanfar, P., Dimakis, A.G., Delbracio, M.: A survey on diffusion models for inverse problems. arXiv (2024)
12. Erbach, J., Narnhofer, D., Dombos, A., Schiele, B., Lenssen, J.E., Schindler, K.: Solving inverse problems with flair. In: Adv. Neural Inform. Process. Syst. (2025)
13. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., Podell, D., Dockhorn, T., English, Z., Rombach, R.: Scaling rectified flow transformers for high-resolution image synthesis. In: International Conference on Machine Learning (2024)
14. Eyring, L., Karthik, S., Roth, K., Dosovitskiy, A., Akata, Z.: Reno: Enhancing one-step text-to-image models through reward-based noise optimization. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) *Advances in Neural Information Processing Systems*. vol. 37, pp. 125487–125519. Curran Associates, Inc. (2024)
15. Gao, Z., Song, J., Zhang, Z., Deng, J., Patras, I.: Frequency-guided diffusion for training-free text-driven image translation. In: Int. Conf. Comput. Vis. (2025)
16. Guo, X., Liu, J., Cui, M., Li, J., Yang, H., Huang, D.: Initno: Boosting text-to-image diffusion models via initial noise optimization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9380–9389 (2024)
17. Guo, Y., Yuan, H., Yang, Y., Chen, M., Wang, M.: Gradient guidance for diffusion models: An optimization perspective. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024)
18. He, Y., Murata, N., Lai, C.H., Takida, Y., Uesaka, T., Kim, D., Liao, W.H., Mitsufuji, Y., Kolter, J.Z., Salakhutdinov, R., Ermon, S.: Manifold preserving guided diffusion. In: Int. Conf. Learn. Represent. (2024)
19. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Adv. Neural Inform. Process. Syst. (2017)
20. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
21. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications (2021)
22. Huang, Y., Huang, J., Liu, Y., Yan, M., Lv, J., Liu, J., Xiong, W., Zhang, H., Cao, L., Chen, S.: Diffusion model-based image editing: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **47**(6) (Jun 2025)
23. Jiao, G., Huang, B., Wang, K.C.J., Liao, R.: Uniedit-flow: Unleashing inversion and editing in the era of flow models. In: The Fourteenth International Conference on Learning Representations (2026)
24. Ju, X., Liu, X., Wang, X., Bian, Y., Shan, Y., Xu, Q.: Brushnet: A plug-and-play image inpainting model with decomposed dual-branch diffusion. In: Eur. Conf. Comput. Vis. (2024)
25. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: IEEE Conf. Comput. Vis. Pattern Recog. (2019)
26. Karunratanakul, K., Preechakul, K., Aksan, E., Beeler, T., Suwajanakorn, S., Tang, S.: Optimizing diffusion noise can serve as universal motion priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1334–1345 (2024)

27. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. In: *Adv. Neural Inform. Process. Syst.* (2022)
28. Kim, J., Kim, B.S., Ye, J.C.: Flowdps: Flow-driven posterior sampling for inverse problems. In: *Int. Conf. Comput. Vis.* (2025)
29. Kim, K., Kim, S.: Model already knows the best noise: Bayesian active noise selection via attention in video diffusion model. In: *The Fourteenth International Conference on Learning Representations* (2026)
30. Kim, S., Suh, S., Lee, M.: Rad: Region-aware diffusion models for image inpainting. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2025)
31. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *Int. Conf. Learn. Represent.* (2015)
32. Li, S., Le, H., Xu, J., Salzmann, M.: All seeds are not equal: Enhancing compositional text-to-image generation with reliable random seeds. In: *Int. Conf. Learn. Represent.* (2025)
33. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *The Eleventh International Conference on Learning Representations* (2023)
34. Liu, K.H., Yang, C.K., Chen, M.H., Liu, Y.L., Lin, Y.Y.: Corrfill: Enhancing faithfulness in reference-based inpainting with correspondence guidance in diffusion models. In: *IEEE Winter Conf. Appl. Comput. Vis.* (2025)
35. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: *Int. Conf. Learn. Represent.* (2023)
36. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2022)
37. Luo, G., Darrell, T., Wang, O., Goldman, D.B., Holynski, A.: Readout guidance: Learning control from diffusion features. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2024)
38. Lyu, Y., Wu, Y., Lin, Y., Du, B.: Is-diff: Improving diffusion-based inpainting with better initial seed. *arXiv* (2025)
39. Manukyan, H., Sargsyan, A., Atanyan, B., Wang, Z., Navasardyan, S., Shi, H.: Hd-painter: High-resolution and prompt-faithful text-guided image inpainting with diffusion models. In: *Int. Conf. Learn. Represent.* (2025)
40. Milanfar, P.: Rendition: Reclaiming what a black box takes away. *SIAM Journal on Imaging Sciences* **11**(4), 2722–2756 (2018)
41. MOUFAD, B., Janati, Y., Shouraki, N.B., Durmus, A.O., Hirtz, T., Moulines, E., Olsson, J.: Efficient zero-shot inpainting with decoupled diffusion guidance. In: *The Fourteenth International Conference on Learning Representations* (2026)
42. Patel, M., Wen, S., Metaxas, D.N., Yang, Y.: Flowchef: Steering of rectified flow models for controlled generations. In: *Int. Conf. Comput. Vis.* (2025)
43. von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Nair, D., Paul, S., Berman, W., Xu, Y., Liu, S., Wolf, T.: Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers> (2022), accessed: 2026-06-30
44. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning* (2021)
45. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)

46. Ronai, O., Kulikov, V., Michaeli, T.: Flowopt: Fast optimization through whole flow processes for training-free editing. arXiv (2025)
47. Rout, L., Raoof, N., Daras, G., Caramanis, C., Dimakis, A.G., Shakkottai, S.: Solving linear inverse problems provably via posterior sampling with latent diffusion models. In: Adv. Neural Inform. Process. Syst. (2023)
48. Sadat, S., Vontobel, T., Salehi, F., Weber, R.M.: Guidance in the frequency domain enables high-fidelity sampling at low cfg scales. arXiv (2025)
49. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: Laion-5b: an open large-scale dataset for training next generation image-text models. In: Adv. Neural Inform. Process. Syst. (2022)
50. Shamsolmoali, P., Zareapoor, M., Zhou, H., Felsberg, M., Tao, D., Li, X.: From missing pieces to masterpieces: Image completion with context-adaptive diffusion. IEEE Trans. Pattern Anal. Mach. Intell. (2025)
51. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=St1giarCHLP>
52. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: IEEE Winter Conf. Appl. Comput. Vis. (2022)
53. Tang, Z., Peng, J., Tang, J., Hong, M., Wang, F., Chang, T.H.: Inference-time alignment of diffusion models with direct noise optimization. In: Forty-second International Conference on Machine Learning (2025)
54. Tao, X., Zhou, C., Shen, X., Wang, J., Jia, J.: Zero-order reverse filtering. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 222–230 (2017)
55. Wallace, B., Gokul, A., Ermon, S., Naik, N.: End-to-end diffusion latent optimization improves classifier guidance. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7280–7290 (October 2023)
56. Wang, Q., Wan, Z., Belkin, M., Wang, Y.: Seeds of structure: Patch pca reveals universal compositional cues in diffusion models. In: Adv. Neural Inform. Process. Syst. (2025)
57. Wang, Y., Cao, C., Yu, J., Fan, K., Xue, X., Fu, Y.: Towards enhanced image inpainting: Mitigating unwanted object insertion and preserving color consistency. In: IEEE Conf. Comput. Vis. Pattern Recog. (2025)
58. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. In: Int. Conf. Learn. Represent. (2023)
59. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: From error visibility to structural similarity. IEEE Trans. Image Process. **13** (05 2004)
60. Wu, T., Si, C., Jiang, Y., Huang, Z., Liu, Z.: Freeinit: Bridging initialization gap in video diffusion models. In: European conference on computer vision. pp. 378–394. Springer (2024)
61. Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., Li, H.: Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. arXiv (2023)
62. Xie, L., Pakhomov, D., Wang, Z., Wu, Z., Chen, Z., Zhou, Y., Zheng, H., Zhang, Z., Lin, Z., Zhou, J., Dong, C.: Turbofill: Adapting few-step text-to-image model for fast image inpainting. In: IEEE Conf. Comput. Vis. Pattern Recog. (2025)

63. Xie, S., Zhang, Z., Lin, Z., Hinz, T., Zhang, K.: Smartbrush: Text and shape guided object inpainting with diffusion model. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
64. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: Learning and evaluating human preferences for text-to-image generation. In: Adv. Neural Inform. Process. Syst. (2023)
65. Yang, B., Gu, S., Zhang, B., Zhang, T., Chen, X., Sun, X., Chen, D., Wen, F.: Paint by example: Exemplar-based image editing with diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
66. Yang, S., Chen, X., Liao, J.: Uni-paint: A unified framework for multimodal image inpainting with pretrained diffusion model. In: ACM Int. Conf. Multimedia (2023)
67. Yeh, R.A., Chen, C., Yian Lim, T., Schwing, A.G., Hasegawa-Johnson, M., Do, M.N.: Semantic image inpainting with deep generative models. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5485–5493 (2017)
68. Yu, J., Wang, Y., Zhao, C., Ghanem, B., Zhang, J.: Freedom: Training-free energy-guided conditional diffusion model. In: Int. Conf. Comput. Vis. (2023)
69. Zhang, G., Ji, J., Zhang, Y., Yu, M., Jaakkola, T., Chang, S.: Towards coherent image inpainting using denoising diffusion implicit models. In: International Conference on Machine Learning (2023)
70. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Int. Conf. Comput. Vis. (2023)
71. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: IEEE Conf. Comput. Vis. Pattern Recog. (2018)
72. Zhou, Z., Shao, S., Bai, L., Zhang, S., Xu, Z., Han, B., Xie, Z.: Golden noise for diffusion models: A learning framework. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17688–17697 (2025)
73. Zhuang, J., Zeng, Y., Liu, W., Yuan, C., Chen, K.: A task is worth one word: Learning with task prompts for high-quality versatile image inpainting. In: Eur. Conf. Comput. Vis. (2024)