

EFFICIENT ZERO-SHOT INPAINTING WITH DECOUPLED DIFFUSION GUIDANCE

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 Diffusion models have emerged as powerful priors for image editing tasks such
 012 as inpainting and local modification, where the objective is to generate realistic
 013 content that remains consistent with observed regions. In particular, zero-shot
 014 approaches that leverage a pretrained diffusion model, without any retraining, have
 015 been shown to achieve highly effective reconstructions. However, state-of-the-art
 016 zero-shot methods typically rely on a sequence of surrogate likelihood functions,
 017 whose scores are used as proxies for the ideal score. This procedure however
 018 requires vector-Jacobian products through the denoiser at every reverse step, in-
 019 troducing significant memory and runtime overhead. To address this issue, we
 020 propose a new likelihood surrogate that yields simple and efficient to sample Gaus-
 021 sian posterior transitions, sidestepping the backpropagation through the denoiser
 022 network. Our extensive experiments show that our method achieves strong ob-
 023 servation consistency compared with fine-tuned baselines and produces coherent,
 024 high-quality reconstructions, all while significantly reducing inference cost.

1 INTRODUCTION

029 We focus on *inpainting problems* in computer vision, which play a central role in applications ranging
 030 from photo restoration to content creation and interactive design. Given an image with prescribed
 031 missing pixels, the objective is to generate a semantically coherent completion while ensuring strict
 032 consistency with the observed region. The importance of this task has motivated extensive research,
 033 spanning both classical approaches and, more recently, generative modeling with diffusion models
 034 (Rombach et al., 2022; Esser et al., 2024; Batifol et al., 2025; Wu et al., 2025).

035 To address this problem, two main diffusion-based approaches have been popularized. The first relies
 036 on training *conditional diffusion models* tailored to a specific editing setup. These models directly
 037 approximate the conditional distribution of interest (Saharia et al., 2022; Wang et al., 2023a; Kawar
 038 et al., 2023; Huang et al., 2025) and take as side inputs additional information such as a mask, a text
 039 prompt, or reference pixels (Saharia et al., 2022; Wang et al., 2023a; Kawar et al., 2023; Huang et al.,
 040 2025). An alternative approach, which has recently attracted growing attention, is *zero-shot image*
 041 *editing*, requiring no extra training or fine-tuning. In this formulation, the task is cast as a Bayesian
 042 inverse problem: the pre-trained diffusion model serves as a prior, while a likelihood term enforces
 043 fidelity to the observations, and the resulting posterior distribution defines the reconstructions (Song &
 044 Ermon, 2019; Song et al., 2021b; Kadkhodaie & Simoncelli, 2020; Kawar et al., 2022; Lugmayr et al.,
 045 2022; Avrahami et al., 2022; Chung et al., 2023; Mardani et al., 2024; Rout et al., 2024a). Sampling
 046 from this posterior is achieved by approximating the score functions associated with the diffusion
 047 model adapted to this distribution. This *plug-and-play* paradigm has been investigated across a variety
 048 of inverse problems, from image restoration to scientific imaging, and has demonstrated strong editing
 049 performance without task-specific training.

050 While current zero-shot methods are appealing, they face a critical practical limitation. Implemen-
 051 tations of strong zero-shot posterior sampling with diffusion priors typically rely on the twisting
 052 function proposed by Ho et al. (2022); Chung et al. (2023); Song et al. (2023a), which corresponds to
 053 the likelihood evaluated at the denoiser’s output given the observation. Simulating the corresponding
 reverse diffusion process requires computing gradients of the denoiser with respect to its input. This
 in turn entails repeated backpropagation through the denoiser network and costly vector-Jacobian

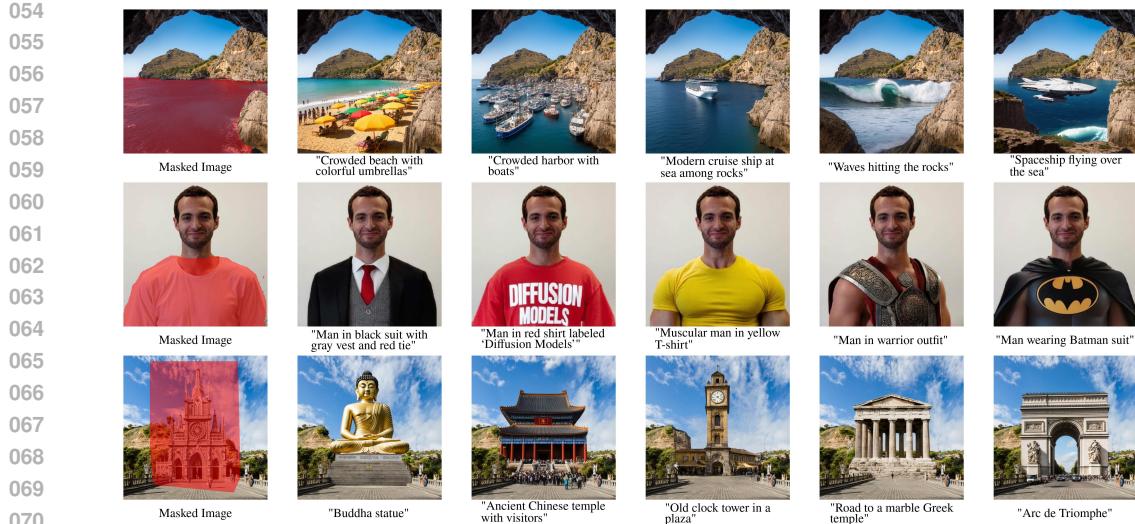


Figure 1: Zero-shot inpainting edits generated by DING (50 NFEs) for different masking patterns using Stable Diffusion 3.5 (medium). Given masked inputs (left column), the model fills the missing regions according to diverse textual prompts.

product (VJP) evaluations. This makes such methods computationally demanding, memory intensive, and often slower than training a dedicated conditional model.

Contributions. We propose a new *VJP-free* framework for zero-shot inpainting with a pre-trained diffusion prior. Our key idea is to approximate the intractable twisted posterior-sampling transitions by a closed-form mixture distribution that can be sampled exactly, thereby eliminating the need for VJP evaluations and backpropagation through the denoiser. Concretely, we modify the twisting function of Ho et al. (2022); Chung et al. (2023) so that it evaluates the denoiser at an independent draw from the pretrained transition. This decoupling breaks the dependency between the denoiser and the arguments of the transition density. As a result, our method provides posterior transitions that can be sampled efficiently for zero-shot inpainting with latent diffusion models. We demonstrate through extensive experiments on Stable Diffusion (SD) 3.5 that our method, coined DECOUPLED INPAINTING GUIDANCE (DING), consistently outperforms state-of-the-art guidance methods under low NFE budgets. It achieves, across three benchmarks, the best trade-off between fidelity to the visible content and realism of the reconstructions, while being both faster and more memory-efficient than competing approaches. Remarkably, even without any task-specific fine-tuning, it outperforms an SD 3 model that has been fine-tuned for image editing, confirming the effectiveness and practicality of our framework.

2 BACKGROUND

Diffusion models Denoising diffusion models (DDMs) (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020) define a generative process for a data distribution p_0 by constructing a continuous path $(p_t)_{t \in [0,1]}$ of distributions between p_0 and $p_1 := \mathcal{N}(0, \mathbf{I}_d)$. More precisely, $p_t = \text{Law}(X_t)$, where

$$X_t = \alpha_t X_0 + \sigma_t X_1, \quad X_0 \sim p_0, \quad X_1 \sim p_1. \quad (2.1)$$

Here X_0 and X_1 are supposed to be independent and $(\alpha_t)_{t \in [0,1]}$ and $(\sigma_t)_{t \in [0,1]}$ are deterministic, non-increasing and non-decreasing, respectively, schedules with boundary conditions $(\alpha_0, \sigma_0) := (1, 0)$ and $(\alpha_1, \sigma_1) := (0, 1)$. Typical choices include the *variance-preserving schedule*, satisfying $\alpha_t^2 + \sigma_t^2 = 1$ (Ho et al., 2020; Dhariwal & Nichol, 2021), and the *linear schedule*, defined by $(\alpha_t, \sigma_t) = (1 - t, t)$ (Lipman et al., 2023; Esser et al., 2024; Gao et al., 2024). The path $(p_t)_{t \in [0,1]}$ defines an interpolation that gradually transforms the clean data distribution p_0 into the Gaussian reference distribution p_1 . To generate new samples, DDMs simulate a time-reversed Markov chain.

Given a decreasing sequence $(t_k)_{k=0}^K$ of time steps with $t_K = 1$ and $t_0 = 0$, reverse transitions are iteratively applied to map a sample from $p_{t_{k+1}}$ to one from p_{t_k} , thereby progressively denoising until convergence to the clean distribution p_0 .

The DDIM framework (Song et al., 2021a) introduces a general family of reverse transitions for denoising diffusion models. It defines a new schedule $(\eta_t)_{t \in [0,1]}$, satisfying $\eta_t \leq \sigma_t$ for all $t \in [0, 1]$, along with a family of transition densities given for $s < t$ by

$$p_{s|t}^\eta(\mathbf{x}_s | \mathbf{x}_t) = \mathbb{E} \left[q_{s|0,1}^\eta(\mathbf{x}_s | X_0, X_1) \mid X_t = \mathbf{x}_t \right], \quad (2.2)$$

where $q_{s|0,1}^\eta(\mathbf{x}_s | \mathbf{x}_0, \mathbf{x}_1) := N(\mathbf{x}_s; \alpha_s \mathbf{x}_0 + \sqrt{\sigma_s^2 - \eta_s^2} \mathbf{x}_1, \eta_s^2 \mathbf{I})$ and the random variables (X_0, X_t, X_1) are defined as in (2.1). By construction, this family satisfies the marginalization property $p_s(\mathbf{x}_s) = \int p_{s|t}^\eta(\mathbf{x}_s | \mathbf{x}_t) p_t(\mathbf{x}_t) d\mathbf{x}_t$ (Song et al., 2021a, Appendix B). Thus, $(p_{t_k|t_{k+1}}^\eta)_{k=0}^{K-1}$ defines a consistent set of reverse transitions, enabling stepwise sampling from the sequence $(p_{t_k})_{k=0}^K$. In practice, however, these transitions are intractable. A common approximation is to replace X_0 and X_1 in (2.2) by their conditional expectations (Ho et al., 2020; Song et al., 2021a). More precisely, let $\hat{\mathbf{x}}_0^\theta(\cdot, t)$ denote a parametric estimator of $\hat{\mathbf{x}}_0(\mathbf{x}_t, t) := \mathbb{E}[X_0 \mid X_t = \mathbf{x}_t]$. Since $\mathbb{E}[X_1 | X_t = \mathbf{x}_t] = (\mathbf{x}_t - \alpha_t \hat{\mathbf{x}}_0(\mathbf{x}_t, t)) / \sigma_t$, we set $\hat{\mathbf{x}}_1^\theta(\mathbf{x}_t, t) := (\mathbf{x}_t - \alpha_t \hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)) / \sigma_t$. Then the parametric model proposed by Ho et al. (2020); Song et al. (2021a) corresponds to approximating each $p_{t_k|t_{k+1}}^\eta$ by

$$p_{t_k|t_{k+1}}^{\eta, \theta}(\mathbf{x}_t | \mathbf{x}_{t_{k+1}}) := q_{t_k|0,1}^\eta(\mathbf{x}_t | \hat{\mathbf{x}}_0^\theta(\mathbf{x}_{t_{k+1}}, t_{k+1}), \hat{\mathbf{x}}_1^\theta(\mathbf{x}_{t_{k+1}}, t_{k+1})). \quad (2.3)$$

For $k = 0$, $p_{0|t_1}^{\eta, \theta}(\cdot | \mathbf{x}_{t_1})$ is simply defined as the Dirac mass at $\hat{\mathbf{x}}_0^\theta(\mathbf{x}_{t_1}, t_1)$. In the rest of the paper we omit the superscript η when there is no ambiguity.

Image editing In this work, we address the task of image editing via inpainting. We assume access to some reference image $\mathbf{x}_* \in \mathbb{R}^d$ that must be modified while remaining consistent with a prescribed set of observed pixels. Let $\mathbf{m} \subset \{1, \dots, d\}$ denote the index set of missing (masked) pixels, and let $\overline{\mathbf{m}} = \{1, \dots, d\} \setminus \mathbf{m}$ be the index set of observed (unmasked) pixels, with cardinality $|\overline{\mathbf{m}}| = d_y$. For any $\mathbf{x} \in \mathbb{R}^d$ and $\mathbf{i} \subset \{1, \dots, d\}$, we denote by $\mathbf{x}[\mathbf{i}] \in \mathbb{R}^{|\mathbf{i}|}$ the subvector formed by the components of \mathbf{x} with indices \mathbf{i} . The observation is thus given by $\mathbf{y} := \mathbf{x}_*[\overline{\mathbf{m}}]$, and the objective is to synthesize a reconstruction $\hat{\mathbf{x}}$ such that $\hat{\mathbf{x}}[\overline{\mathbf{m}}] \approx \mathbf{y}$ while generating the missing region $\hat{\mathbf{x}}[\mathbf{m}]$ in a realistic and semantically coherent manner with respect to the observed pixels. In the Bayesian formulation, the data distribution p_0 serves as a prior over natural images, while the observation model is encoded by a Gaussian likelihood on the observed coordinates:

$$\ell_0(\mathbf{y} | \mathbf{x}) = N(\mathbf{y}; \mathbf{x}[\overline{\mathbf{m}}], \sigma_y^2 \mathbf{I}_{d_y}). \quad (2.4)$$

The parameter $\sigma_y > 0$ serves as a relaxation factor: smaller values enforce strict adherence to the observation, while larger values permit controlled deviations from \mathbf{x}_* , thereby facilitating the reconstruction process. In this Bayesian framework, the target distribution from which we aim to sample is the *posterior distribution*

$$\pi_0(\mathbf{x}_0 | \mathbf{y}) \propto \ell_0(\mathbf{y} | \mathbf{x}_0) p_0(\mathbf{x}_0). \quad (2.5)$$

Inference-time guidance As observed in the seminal works of Song & Ermon (2019); Kadkhodaie & Simoncelli (2020); Song et al. (2021b); Kawar et al. (2021), approximate sampling from the posterior distribution can be performed by biasing the denoising process with guidance terms, without requiring any additional fine-tuning. The central idea is to modify the sampling dynamics of diffusion models on-the-fly so that the generated samples both satisfy the likelihood constraint $\ell_0(\mathbf{y} | \cdot)$ and remain plausible under the prior p_0 . More precisely, a standard approach is to approximate the iterative updates of a diffusion model defined to target the posterior $\ell_0(\mathbf{y} | \cdot)$. This in turn entails deriving an approximation of the posterior denoiser $\hat{\mathbf{x}}_0(\mathbf{x}_t, t | \mathbf{y}) := \int \mathbf{x}_0 \pi_{0|t}(\mathbf{x}_0 | \mathbf{x}_t, \mathbf{y}) d\mathbf{x}_0$, where $\pi_{0|t}(\mathbf{x}_0 | \mathbf{x}_t, \mathbf{y}) \propto \pi_0(\mathbf{x}_0 | \mathbf{y}) N(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})$. The denoiser $\hat{\mathbf{x}}_0(\cdot, t | \mathbf{y})$ is related to the prior denoiser via the identity

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, t | \mathbf{y}) = \hat{\mathbf{x}}_0(\mathbf{x}_t, t) + \alpha_t^{-1} \sigma_t^2 \nabla_{\mathbf{x}_t} \log \ell_t(\mathbf{y} | \mathbf{x}_t), \quad (2.6)$$

where the additional term is referred to as the *guidance term*; see (Daras et al., 2024, Eq. 2.15 and 2.17). It is defined as the logarithmic gradient of the *propagated likelihood*

$$\ell_t(\mathbf{y} | \mathbf{x}_t) := \int \ell_0(\mathbf{y} | \mathbf{x}_0) p_{0|t}(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0, \text{ with } p_{0|t}(\mathbf{x}_0 | \mathbf{x}_t) \propto p_0(\mathbf{x}_0) N(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I}); \quad (2.7)$$

162 **Algorithm 1** Posterior sampling with decoupled guidance

```

163
164 1: Input: decreasing timesteps  $(t_k)_{k=K}^0$  with  $t_K = 1$ ,  $t_0 = 0$ ; original image  $\mathbf{x}_*$ ; mask  $\mathbf{m}$ ;
165 2:  $\mathbf{y} \leftarrow \mathbf{x}_*[\mathbf{m}]$ ;  $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)$ 
166 3: for  $k = K - 1$  to 1 do
167 4:    $\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}_0^\theta(\mathbf{x}, t_{k+1})$ 
168 5:    $\hat{\mathbf{x}}_1 \leftarrow (\mathbf{x} - \alpha_{t_{k+1}} \hat{\mathbf{x}}_0) / \sigma_{t_{k+1}}$ 
169 6:    $\boldsymbol{\mu} \leftarrow \alpha_{t_k} \hat{\mathbf{x}}_0 + (\sigma_{t_k}^2 - \eta_k^2)^{1/2} \hat{\mathbf{x}}_1$ 
170 7:   /* Sampling (3.4) */
171 8:    $(\mathbf{w}, \mathbf{w}') \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0_d, \mathbf{I}_d)$ 
172 9:    $\mathbf{z} \leftarrow \boldsymbol{\mu} + \eta_k \mathbf{w}$ 
173 10:   $\hat{\mathbf{x}}_1^{\text{pxy}} \leftarrow (\mathbf{z} - \alpha_{t_k} \hat{\mathbf{x}}_0^\theta(\mathbf{z}, t_k)) / \sigma_{t_k}$ 
174 11:   $\gamma \leftarrow \eta_{t_k}^2 / (\eta_{t_k}^2 + \alpha_{t_k}^2 \sigma_{\mathbf{y}}^2)$ 
175 12:   $\mathbf{x}[\mathbf{m}] \leftarrow \boldsymbol{\mu}[\mathbf{m}] + \eta_k \mathbf{w}'[\mathbf{m}]$ 
176 13:   $\mathbf{x}[\overline{\mathbf{m}}] \leftarrow (1 - \gamma) \boldsymbol{\mu}[\overline{\mathbf{m}}] + \gamma (\alpha_{t_k} \mathbf{y} + \sigma_{t_k} \hat{\mathbf{x}}_1^{\text{pxy}}[\overline{\mathbf{m}}]) + \alpha_{t_k} \sigma_{\mathbf{y}} \sqrt{\gamma} \mathbf{w}'[\overline{\mathbf{m}}]$ 
177 14: end for
178 15: Return:  $\hat{\mathbf{x}}_0^\theta(\mathbf{x}, t_1)$ 

```

181 see (Daras et al., 2024, Equation 2.20). Since the pre-trained parametric approximation $\hat{\mathbf{x}}_0^\theta(\cdot, t)$ of the
182 prior denoiser $\hat{\mathbf{x}}_0(\cdot, t)$ is already available, estimating $\hat{\mathbf{x}}_0(\cdot, t|\mathbf{y})$ reduces to computing the intractable
183 score term $\nabla_{\mathbf{x}_t} \log \ell_t(\mathbf{y}|\mathbf{x}_t)$. A widely adopted approximation (Ho et al., 2022; Chung et al., 2023)
184 replaces $p_{0|t}(\cdot|\mathbf{x}_t)$ in (2.7) by a Dirac mass at the denoiser estimate $\hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)$, yielding

$$\hat{\ell}_t^\theta(\mathbf{y}|\mathbf{x}_t) := \ell_0(\mathbf{y}|\hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)). \quad (2.8)$$

185 This approximation is often combined with a suitable rescaling weight (possibly depending on \mathbf{x}_t);
186 see (Ho et al., 2022, Equation 8) and (Chung et al., 2023, Algorithm 1). Substituting this into the
187 identity (2.6) yields an approximation of the posterior denoiser, which in turn defines an approximate
188 diffusion model for $\pi_0(\cdot|\mathbf{y})$.
189

190 3 METHOD

191 The methods discussed in the previous section rely on the likelihood approximation (2.8), which is
192 then inserted into (2.6). However, computing this term requires differentiating through the denoisers
193 $\hat{\mathbf{x}}_0^\theta(\cdot, t_k)$ at each timestep t_k . This operation is computationally demanding: it increases memory
194 usage, slows down the sampling process, and reduces scalability. By contrast, fine-tuned conditional
195 diffusion models bypass these inference costs once training is complete, but at the expense of per-task
196 retraining. This highlights a fundamental trade-off: zero-shot posterior sampling eliminates the
197 need for retraining, but incurs substantial overhead during inference. Our goal is to bridge this gap
198 by designing a zero-shot posterior sampler that removes the need for backpropagation through the
199 denoiser while preserving the effectiveness of guidance.

200 **Reverse transitions for the posterior.** Our method builds upon the alternative sampling strategy
201 introduced in Wu et al. (2023); Zhang et al. (2023); Janati et al. (2024). Instead of initializing
202 the interpolation (2.1) with the prior $X_0 \sim p_0$, we consider the same process initialized from the
203 posterior distribution $X_0 \sim \pi_0(\cdot|\mathbf{y})$. This yields a new family of random variables whose marginals
204 are $\pi_t(\mathbf{x}_t|\mathbf{y}) := \int N(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I}_d) \pi_0(\mathbf{x}_0|\mathbf{y}) d\mathbf{x}_0$, in analogy with the prior family $(p_t)_{t \in [0,1]}$.
205 Moreover, the DDIM transitions associated with $(\pi_t(\cdot|\mathbf{y}))_{t \in [0,1]}$ are given by (Janati et al., 2025b,
206 Equation 1.17):

$$\pi_{s|t}^\eta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) \propto \ell_s(\mathbf{y}|\mathbf{x}_s) p_{s|t}^\eta(\mathbf{x}_s|\mathbf{x}_t), \quad (3.1)$$

207 which defines a valid Markov chain with marginals $(\pi_{t_k}(\cdot|\mathbf{y}))_{k=1}^K$. This chain defines a path between
208 the Gaussian reference $\mathcal{N}(0, \mathbf{I}_d)$ and the posterior distribution $\pi_0(\cdot|\mathbf{y})$. However, the presence of
209 the likelihood term $\ell_t(\mathbf{y}|\mathbf{x}_t)$ makes also these transitions intractable. To address this issue, prior
210 works (Zhang et al., 2023; Wu et al., 2023) introduced the surrogate transitions proportional to
211 $\mathbf{x}_s \mapsto \hat{\ell}_s^\theta(\mathbf{y}|\mathbf{x}_s) p_{s|t}^{\eta, \theta}(\mathbf{x}_s|\mathbf{x}_t)$, for fixed \mathbf{x}_t and \mathbf{y} , where $\hat{\ell}_t^\theta(\mathbf{y}|\cdot)$ are defined in (2.8). These transitions

216 are then approximated using either variational inference (Janati et al., 2024; Pandey et al., 2025) or
 217 sequential Monte Carlo methods (Wu et al., 2023). However, similar to the methods described in the
 218 previous section, these approximations rely on the approximate guidance term and thus suffer from
 219 inflated memory usage and higher runtime.
 220

221 **Our likelihood approximation.** To address this limitation, we draw inspiration from (2.8) to
 222 propose a lightweight approximation, designed to eliminate the need for VJP evaluations through
 223 the denoiser. Using the relation $\hat{\mathbf{x}}_0^\theta(\mathbf{x}_s, s) = (\mathbf{x}_s - \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{x}_s, s)) / \alpha_s$, we first rewrite the standard
 224 likelihood approximation (2.8) in terms of the noise prediction $\hat{\mathbf{x}}_1^\theta(\mathbf{x}_s, s)$ according to

$$\hat{\ell}_s^\theta(\mathbf{y}|\mathbf{x}_s) = \ell_0(\mathbf{y}|(\mathbf{x}_s - \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{x}_s, s)) / \alpha_s).$$

225 Based on this parametrization, we then introduce the following alternative approximation:
 226

$$\hat{\ell}_s^\theta(\mathbf{y}|\mathbf{x}_s, \mathbf{z}_s) := \ell_0(\mathbf{y}|(\mathbf{x}_s - \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{z}_s, s)) / \alpha_s), \quad (3.2)$$

227 where the noise predictor is evaluated at $\mathbf{z}_s \in \mathbb{R}^d$, which serves as a proxy for \mathbf{x}_s . A key feature of
 228 this decoupling is that it enables lightweight updates, avoids costly denoiser backpropagation, and
 229 still provides high-quality reconstructions. Then, similarly to (3.1), we define
 230

$$\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y}) \propto \hat{\ell}_s^\theta(\mathbf{y}|\mathbf{x}_s, \mathbf{z}_s) p_{s|t}^{\eta, \theta}(\mathbf{x}_s|\mathbf{x}_t).$$

231 This leads us to propose the surrogate
 232

$$\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) := \mathbb{E} [\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|Z_s, \mathbf{x}_t, \mathbf{y})], \quad (3.3)$$

233 where $Z_s \sim p_{s|t}^\theta(\cdot|\mathbf{x}_t)$, for (3.1). The transition $\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y})$ generally lacks a closed-form ex-
 234 pression; nevertheless, since it has a *mixture structure*, it allows for straightforward and efficient
 235 sampling. Sampling from $\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y})$ can be performed by first drawing Z_s from $p_{s|t}^\theta(\cdot|\mathbf{x}_t)$, and
 236 then sampling from $\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|Z_s, \mathbf{x}_t, \mathbf{y})$. Moreover, as we will now show, in the case of inpainting, the
 237 second step can be carried out exactly.
 238

239 Let $\mu_{s|t}^\theta(\mathbf{x}_t; \eta)$ denote the mean of the Gaussian reverse transition $p_{s|t}^{\eta, \theta}(\cdot|\mathbf{x}_t)$. In the case of inpainting
 240 (2.4), standard Gaussian conjugacy results (Bishop, 2006, Equation 2.116) show that $\hat{\pi}_{s|t}^\theta(\cdot|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y})$
 241 admits a closed-form Gaussian expression:
 242

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y}) &= \mathcal{N}(\mathbf{x}_s|\mathbf{m}; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 \mathbf{I}_{d-d_y}) \\ &\times \mathcal{N}(\mathbf{x}_s[\bar{\mathbf{m}}]; (1 - \gamma_{s|t}) \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}] + \gamma_{s|t} (\alpha_s \mathbf{y} + \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{z}_s, s)[\bar{\mathbf{m}}]), \alpha_s^2 \sigma_y^2 \gamma_{s|t} \mathbf{I}_{d_y}), \end{aligned} \quad (3.4)$$

243 with $\gamma_{s|t} := \eta_s^2 / (\eta_s^2 + \alpha_s^2 \sigma_y^2)$. A derivation is provided in Appendix A.1. Thus, a sample X_s from
 244 (3.3) can be drawn exactly by, first, generating a realization \mathbf{z}_s of $Z_s \sim p_{s|t}^{\eta, \theta}(\cdot|\mathbf{x}_t)$ and, second,
 245 sampling $X_s[\mathbf{m}]$ and $X_s[\bar{\mathbf{m}}]$ conditionally independently from the two Gaussian distributions in
 246 (3.4); see Algorithm 1 for a pseudocode of this approach, which we refer to as DING (see Section 1).
 247

248 **Practical implementation.** A key practical feature of our method is that it depends on a single
 249 hyperparameter: the sequence $(\eta_t)_{t \in [0, 1]}$ of standard deviations, which controls the level of stochas-
 250 ticity in the DDIM reverse process. This choice is particularly critical in the low-NFE regime, where
 251 only a few function evaluations are available and the variance schedule strongly influences both
 252 observation fidelity and perceptual quality. In all experiments, we adopt the schedule $\eta_t = \sigma_t(1 - \alpha_t)$.
 253 An ablation study of this parameter is reported in Section 4.
 254

255 Beyond this hyperparameter, an important practical consideration is that most large-scale diffusion
 256 models for high-resolution image generation operate in a compressed latent space rather than in pixel
 257 space (Rombach et al., 2022; Esser et al., 2024). To apply our algorithm in this setting, we must
 258 therefore formulate the inpainting task in the latent domain. Denote by Enc the encoder, \mathbf{X}_* the
 259 pixel-space ground-truth image and \mathbf{M} the corresponding pixel-space mask. Following Avrahami et al.
 260 (2022), we set $\mathbf{x}_* := \text{Enc}(\mathbf{X}_*)$, the observation to $\mathbf{y} := \mathbf{x}_*[\mathbf{m}]$ where \mathbf{m} is a downsampled version of
 261 the pixel-space mask \mathbf{M} . Since the encoder reduces spatial resolution by a fixed factor (e.g., $s = 8$ in
 262 Esser et al. (2024)), we construct the latent mask \mathbf{m} by average pooling the binary pixel-space mask
 263 \mathbf{M} with kernel and stride s . Each latent site is assigned the fraction of unmasked pixels within its
 264 receptive field. These fractional values are then thresholded (typically at 0.5) to produce a binary
 265 mask; in other words, a latent site is marked as observed if the majority of its underlying pixels are
 266 unmasked. In practice, the mask \mathbf{m} is provided as a single-channel image and broadcast across all
 267 latent channels when applied to \mathbf{x}_* . Finally, we apply Algorithm 1 with $(\mathbf{x}_*, \mathbf{y}, \mathbf{m})$ thus defined in
 268 the latent space.
 269

270 **Related methods.** Our work shares similarities with various recent approaches to zero-shot diffusion
 271 guidance, which we now briefly review. The closest line of work comprises variants of the *replacement*
 272 *method* (Song & Ermon, 2019; Song et al., 2021b), which follows the same structure as Algorithm 1.
 273 In these schemes, the masked coordinates of the state are updated according to the standard DDIM
 274 transition (Line 11), while the unmasked coordinates are replaced by a direct update that enforces
 275 consistency with the observation \mathbf{y} (Line 12). In its simplest form, the method performs ancestral
 276 sampling with the transition

$$277 \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) = N(\mathbf{x}_s[\mathbf{m}]; \boldsymbol{\mu}_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 \mathbf{I}_{d_{\mathbf{y}}}) N(\mathbf{x}_s[\bar{\mathbf{m}}]; \alpha_s \mathbf{y}, \sigma_s^2 \mathbf{I}_{d_{\mathbf{y}}}), \quad (3.5)$$

278 *i.e.*, the unmasked state is set to a noisy version of the observation $\alpha_s \mathbf{y} + \sigma_s W_s$, where $W_s \sim$
 279 $\mathcal{N}(0, \mathbf{I}_{d_{\mathbf{y}}})$; see (Song & Ermon, 2019, Algorithm 2) and (Song et al., 2021b, Appendix I.2). Avrahami
 280 et al. (2022) extended this approach to the latent domain using a downsampled mask. The method
 281 was later refined in RePaint (Lugmayr et al., 2022), which improves sample quality by performing
 282 multiple back-and-forth updates: after applying the replacement step from t_{k+1} to t_k , a forward
 283 noising step is applied from t_k back to t_{k+1} , and this cycle is repeated several times. Several works
 284 have combined the replacement method with sequential Monte Carlo (SMC) sampling (Trippe et al.,
 285 2023; Cardoso et al., 2023; Dou & Song, 2024; Corenflos et al., 2025; Zhao, 2025). In particular,
 286 Cardoso et al. (2023) update the unmasked coordinates of each particle using a Gaussian transition
 287 whose mean is a convex combination of the DDIM mean and the rescaled observation $\alpha_{t_k} \mathbf{y}$. In the
 288 inpainting framework, the recently proposed PnP-Flow (Martin et al., 2025) reduces to using similar
 289 transitions without relying on SMC, *i.e.*, by using a single particle. We explicitly compare the update
 290 rules in (Cardoso et al., 2023; Martin et al., 2025) with ours in Appendix A.2, where we also discuss
 291 additional related work.

292 4 EXPERIMENTS

294 In this section, we extensively evaluate the inpainting performance of DING when used with different
 295 large-scale models. We benchmark its performance on multiple datasets against several state-of-the-
 296 art baselines. We further analyze the relevance of our modeling choices, specifically the formulation
 297 of the approximation in (3.2) and the schedule of DDIM standard deviations $(\eta_t)_{t \in [0,1]}$, through a
 298 series of targeted ablations.

300 **Models and datasets.** We evaluate our method on Stable Diffusion 3.5 (medium) (Esser et al.,
 301 2024). We set the CFG scale to 2. Our experiments cover three datasets: FFHQ (Karras et al., 2019),
 302 DIV2K (Agustsson & Timofte, 2017), and PIE-Bench (Ju et al., 2024). For FFHQ, we use the first
 303 5k images and condition generation on the prompt “*a high-quality photo of a face*”. For DIV2K,
 304 we include both training and validation splits (900 images in total), and generate captions for each
 305 image using BLIP-2 (Li et al., 2023); see Appendix B for details. All FFHQ and DIV2K images are
 306 resized to a resolution of 768×768 . The PIE-Bench dataset contains 700 images of resolution
 307 512×512 , each paired with an inpainting mask and an edit caption. After removing cases where the
 308 mask completely covers the image, the resulting evaluation set contains 556 images.

309 **Evaluation and masks.** For FFHQ and DIV2K, we evaluate inpainting performance under four
 310 rectangular masking configurations: (i) right half of the image (*Half*), (ii) upper half (*Top*), (iii)
 311 lower half (*Bottom*), and (iv) a central 512×512 square (*Center*). In contrast, PIE-Bench provides
 312 irregular masks with diverse spatial patterns; see Appendix C for examples. Unless otherwise
 313 stated, we set $\sigma_y = 0.01$ across all tasks. Since exact posterior sampling is infeasible, we assess
 314 inpainting quality using proxy metrics. To measure distributional alignment with the dataset, we
 315 report both FID and patch FID (pFID) (Chai et al., 2022), the latter offering finer granularity for
 316 high-resolution evaluation. Following the standard FID protocol, we extract 10 random 256×256
 317 patches per image, yielding a total of 50k patches. To quantify consistency with the observed
 318 content, we compute context PSNR (cPSNR), defined as the PSNR over the unmasked region only.
 319 We further report LPIPS (Zhang et al., 2018) relative to the ground truth to evaluate perceptual
 320 similarity, which is especially relevant for FFHQ where facial symmetries make reconstructions
 321 visually close to the reference. For PIE-Bench, which includes edit captions, we additionally report
 322 CLIP-Score (Radford et al., 2021) on both the full image (CLIP) and the edited region (CLIP-ED),
 323 following Ju et al. (2024). Together, these metrics provide a comprehensive evaluation of inpainting
 324 quality. While each captures a different aspect of performance, none should be interpreted in isolation.

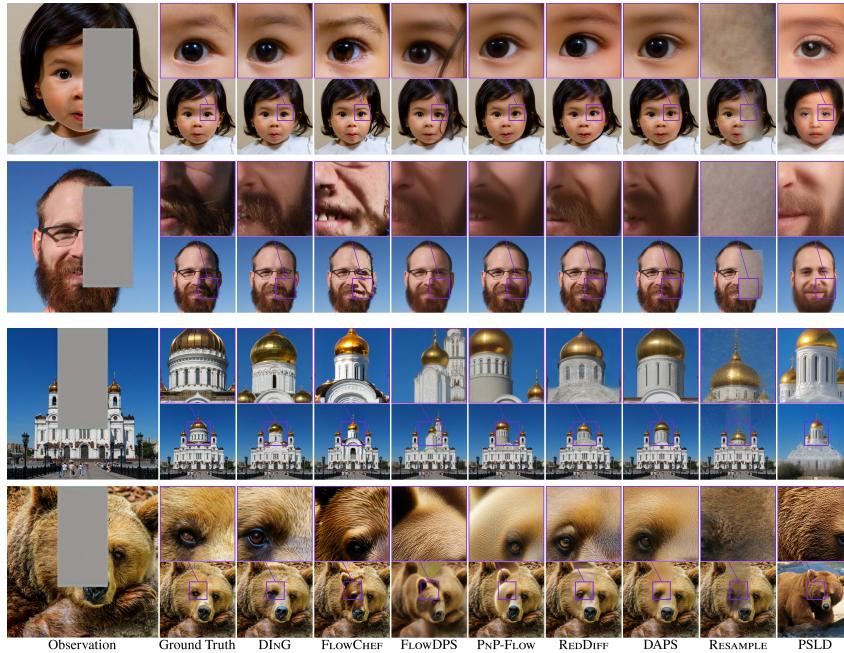


Figure 2: Examples of reconstructions on FFHQ and DIV2K with 50 NFEs.

Baselines. We compare against seven state-of-the-art baselines: FLOWCHEF (Patel et al., 2024), FLOWDPS (Kim et al., 2025), DAPS (Zhang et al., 2025), REDDIFF (Mardani et al., 2024), RESAMPLE (Song et al., 2024), PSLD (Rout et al., 2024b), PNP-FLOW (Martin et al., 2025), DIFFPIR (Zhu et al., 2023), DDNM (Wang et al., 2023b) and BLENDED-DIFF (Avrahami et al., 2023). For the main comparison, all methods are evaluated under a fixed budget of 50 NFEs. Since our method requires two denoiser evaluations per diffusion step, we use 25 steps to match this budget. We focus on this low-NFE regime as it reflects realistic settings, where inference is constrained by latency and compute. To ensure fairness, all methods are run in the latent space with downsampled masks, and extensive hyperparameter tuning is performed for each baseline on each dataset. For baselines that require VJP or backpropagation through the denoiser, we report their actual runtime and memory costs, ensuring that comparisons reflect effective inference cost rather than nominal NFE counts. Average runtime and memory usage across all the experiments, measured on H100 GPUs, are provided in Table 1.

4.1 MAIN RESULTS

Tables 2 and 3 summarize the results on FFHQ, DIV2K and PIE-Bench, respectively. On FFHQ (Table 3), DING achieves the best performance on all masks and almost all the metrics. In particular, it improves both pFID and FID by significant margins over the strongest competing method FLOWCHEF. It also obtains the highest cPSNR scores, indicating a faithful reconstruction of the visible content, while simultaneously producing visually coherent completions with the lowest LPIPS. On DIV2K, the comparison is more nuanced. DING consistently attains the best FID and LPIPS across all four masks, while remaining competitive on pFID and comparable to strong baselines on cPSNR. On PIE-Bench, DING achieves the best results on all metrics except cPSNR and CLIP. We note, however, that although FLOWDPS and PSLD obtain slightly higher CLIP

Table 1: Memory and runtime.

Method	Time (s)	Mem. (GB)
BLENDED-DIFF	3.0	22.09
DAPS	9.1	22.09
DDNM	3.1	22.09
DIFFPIR	3.1	22.09
FLOWCHEF	3.0	22.09
FLOWDPS	3.0	22.10
PnP-FLOW	3.1	22.09
PSLD	7.4	24.49
REDDIFF	3.1	22.09
RESAMPLE	8.1	24.50
DING (ours)	2.9	22.09

Table 3: Results on PIE-Bench with 556 samples and 50 NFEs.

Method	FID	pFID	cPSNR	LPIPS	CLIP	CLIP-ED
BLENDED-DIFF	65.5	27.0	26.60	0.31	26.32	23.15
DAPS	65.9	30.2	27.08	0.34	25.57	21.75
DDNM	61.4	26.9	27.29	0.31	26.27	22.96
DIFFPIR	63.5	25.4	26.98	0.30	26.21	23.04
FLOWCHEF	68.3	27.4	26.84	0.30	26.02	22.47
FLOWDPS	74.6	42.7	22.05	0.45	26.35	22.79
PnP-FLOW	66.8	32.1	26.90	0.34	25.62	21.02
PSLD	94.1	56.2	14.25	0.65	26.32	21.81
REDDIFF	69.5	35.2	24.34	0.37	25.27	21.18
RESAMPLE	71.0	33.9	24.45	0.35	25.71	22.03
DING (ours)	61.4	24.7	27.03	0.30	26.30	23.36

378 Table 2: **Top**: Quantitative results on FFHQ 768×768 with 5k samples. **Bottom**: DIV2K 768×768 with 900
379 samples. For FID, pFID, and LPIPS, the lower the better. For cPSNR, the higher the better. **50 NFEs were used.**

Method	Half				Center				Top				Bottom			
	FID	pFID	cPSNR	LPIPS	FID	pFID	cPSNR	LPIPS	FID	pFID	cPSNR	LPIPS	FID	pFID	cPSNR	LPIPS
FFHQ 768×768																
BLENDDED-DIFF	23.5	16.3	31.32	0.38	35.3	36.7	31.54	0.33	32.8	15.8	32.05	0.38	43.7	19.8	30.85	0.37
DAPS	17.9	25.1	30.50	0.36	35.1	54.5	31.15	0.32	30.1	30.5	31.54	0.39	52.8	27.6	30.30	0.34
DDNM	12.3	13.68	31.27	0.33	24.4	34.8	31.61	0.27	22.3	23.2	31.82	0.36	38.3	19.6	30.51	0.32
DIFFPIR	12.1	11.23	30.91	0.36	19.4	19.6	31.67	0.30	19.7	14.1	32.07	0.36	30.7	11.4	30.74	0.35
FLOWCHEF	20.2	16.5	30.41	0.36	29.3	35.0	31.00	0.31	27.8	21.1	31.05	0.36	35.9	22.7	29.94	0.35
FLOWDPS	36.2	49.0	26.72	0.46	49.5	79.9	23.36	0.53	52.1	74.2	24.15	0.56	72.3	71.5	23.06	0.55
PNP-FLOW	20.5	33.4	30.62	0.37	36.6	65.1	31.67	0.32	33.6	42.7	31.54	0.38	56.8	33.4	29.95	0.33
PSLD	116.3	73.8	6.89	0.81	98.0	69.1	10.09	0.73	120.6	75.4	7.06	0.81	107.0	70.4	6.46	0.81
REDDIFF	28.5	37.9	27.39	0.39	30.7	41.8	27.85	0.32	33.0	41.1	27.92	0.41	76.4	41.3	26.96	0.39
RESAMPLE	32.4	48.8	28.53	0.44	53.8	103.4	28.46	0.40	63.2	56.2	29.02	0.44	97.8	57.0	28.06	0.44
DING (ours)	9.6	6.6	31.03	0.33	15.5	14.0	31.38	0.27	19.7	12.5	31.64	0.34	29.6	8.6	30.50	0.32
DIV2K 768×768																
BLENDDED-DIFF	43.6	12.9	26.03	0.37	54.8	20.2	26.43	0.35	44.8	13.2	25.28	0.39	48.1	13.1	26.85	0.38
DAPS	51.0	38.4	25.92	0.46	74.8	67.6	26.14	0.44	54.8	41.0	25.22	0.44	61.2	39.7	26.71	0.50
DDNM	42.5	21.2	26.03	0.41	57.7	38.5	26.61	0.37	45.7	21.3	25.36	0.42	49.6	23.2	26.81	0.45
DIFFPIR	41.1	12.9	26.09	0.37	52.8	21.4	26.58	0.34	43.5	13.1	25.36	0.39	44.9	14.9	26.91	0.39
FLOWCHEF	43.3	12.2	25.78	0.36	53.6	22.3	26.27	0.32	45.0	13.8	25.09	0.37	46.9	13.2	26.57	0.37
FLOWDPS	50.8	33.2	21.30	0.49	70.3	62.8	18.38	0.63	64.1	57.9	17.43	0.65	64.2	57.3	19.06	0.63
PNP-FLOW	54.2	42.7	26.00	0.46	79.7	71.4	26.63	0.44	57.1	33.5	25.19	0.44	64.7	50.8	26.61	0.50
PSLD	66.4	32.3	6.15	0.79	66.9	35.7	9.89	0.72	66.5	31.9	6.35	0.79	66.3	32.6	6.41	0.78
REDDIFF	54.2	45.7	22.64	0.49	77.4	69.8	23.25	0.46	57.7	40.9	22.17	0.48	60.6	46.3	23.41	0.52
RESAMPLE	52.7	34.1	23.33	0.47	80.8	63.8	23.76	0.43	56.1	33.1	22.84	0.47	60.6	41.0	24.06	0.48
DING (ours)	39.2	13.0	25.90	0.35	50.7	19.5	26.41	0.31	41.4	13.7	25.19	0.37	43.4	13.4	26.72	0.37

396 scores, they perform markedly worse on fidelity and perceptual quality metrics, suggesting that their
397 improvements in CLIP may reflect metric hacking rather than genuine reconstruction quality. We
398 provide qualitative comparisons of the reconstruction in Figure 7 and Appendix C.

400 We now compare DING with a Stable Diffusion 3 model fine-tuned for inpainting¹, trained on
401 12M images at 1024×1024 resolution. To ensure fairness, both models are evaluated under
402 the same runtime budget (2.2s), which corresponds to 56 NFEs for DING and 28 NFEs for
403 the fine-tuned baseline. We also provide the results for the finetuned model using 56 NFEs.
404 The results are given in Tables 5 and 4. Across
405 FFHQ, DIV2K, and PIE-Bench, DING
406 consistently outperforms the fine-tuned SD3 model
407 on all reported metrics. The gains are especially
408 pronounced in cPSNR, where DING achieves
409 8–10 dB higher fidelity to the observed pixels.
410 This indicates that our method preserves the
411 known content far more accurately while still producing realistic completions, as confirmed by
412 lower FID and LPIPS. On PIE-Bench, DING further improves over the fine-tuned baseline on
413 every metric, including perceptual ones (pFID, LPIPS), while also yielding stronger text–image
414 alignment (CLIP, CLIP-ED). See Figure 7 for a qualitative comparison of the reconstructions. These
415 results demonstrate that, even without task-specific fine-tuning on a large amount of images, DING
416 not only matches but surpasses a specialized SD3 inpaint model. Overall, these results show that our
417 method provides the strongest overall trade-off between realism and fidelity under low NFE budgets.

Table 4: Results on the PIE-Bench with 556 samples.

Method	FID	pFID	cPSNR	LPIPS	CLIP	CLIP-ED
SD3 Inpaint (28)	68.7	30.5	18.85	0.34	26.37	23.10
SD3 Inpaint (56)	68.2	30.3	19.03	0.34	26.38	23.03
DING (ours)	63.6	24.6	26.98	0.30	26.63	23.70

Table 5: DING compared to SD3 fine-tuned (28 and 56 NFEs) for inpainting tasks.

Method	Half				Center				Top				Bottom			
	FID	pFID	cPSNR	LPIPS												
FFHQ 512×512																
SD3 Inpaint (28)	23.5	10.7	21.69	0.37	62.1	33.9	22.18	0.31	34.7	17.8	21.64	0.36	42.4	16.5	21.78	0.37
SD3 Inpaint (56)	23.7	10.3	21.53	0.37	63.7	34.4	21.94	0.31	35.4	16.5	21.41	0.36	43.8	16.8	21.53	0.36
DING (ours)	9.3	5.8	31.40	0.32	20.2	15.5	31.39	0.28	17.3	8.4	31.96	0.33	33.8	12.2	31.27	0.34
DIV2K 512×512																
SD3 Inpaint (28)	45.9	15.0	17.95	0.40	54.2	22.1	18.57	0.36	48.8	16.6	18.12	0.42	51.0	17.5	18.95	0.41
SD3 Inpaint (56)	45.1	14.0	17.91	0.40	54.2	20.5	18.63	0.36	48.6	16.1	18.16	0.41	50.3	17.2	19.02	0.41
DING (ours)	41.5	14.2	26.09	0.37	52.4	21.5	26.53	0.33	43.7	13.8	25.47	0.38	45.4	15.4	26.94	0.38

4.2 ABLATIONS

429 **Doubled NFE per diffusion step.** Because the x_1 -predictor must be evaluated at the proxy variable
430 (Line 9 in Algorithm 1), our algorithm requires two NFEs per diffusion step. An immediate question
431

¹<https://huggingface.co/alimama-creative/SD3-Controlnet-Inpainting>

is whether this overhead is needed. To explore this, we introduce a variant in which the noise prediction from the previous step is reused instead of being recomputed at the proxy. We coin this variant as Delayed DING, where Line 9 is replaced by $\hat{\mathbf{x}}_1^{\text{PXY}} \leftarrow (\mathbf{x} - \sigma_{t_k} \hat{\mathbf{x}}_1(\mathbf{x}, t_{k+1}))/\alpha_{t_k}$, and we further set $\eta_t = \sigma_t(1 - \alpha_t)$, which we found to yield the best performance in this setting. A quantitative comparison with the original DING is reported in Table 6, showing that while Delayed DING reduces the NFE cost per step, it consistently underperforms across metrics and masking patterns, indicating that the doubled NFE is necessary to retain the full effectiveness of our approach.

Table 6: Delayed DING compared to DING on FFHQ (5k samples) and DIV2K (900 samples) with 50 NFEs.

Method	Half				Center				Top				Bottom			
	pFID	FID	cPSNR	LPIPS	pFID	FID	cPSNR	LPIPS	pFID	FID	cPSNR	LPIPS	pFID	FID	cPSNR	LPIPS
FFHQ 768 × 768																
Delayed DING	9.1	7.4	29.21	0.33	21.3	20.7	29.90	0.26	15.7	9.9	29.84	0.33	31.0	12.3	28.88	0.33
DING	9.6	6.6	31.03	0.33	15.5	14.0	31.38	0.27	19.7	12.5	31.64	0.34	29.6	8.6	30.50	0.32
DIV2K 768 × 768																
Delayed DING	43.9	15.9	24.88	0.36	55.6	24.7	25.52	0.32	45.3	16.8	24.36	0.38	47.8	14.6	25.64	0.38
DING	39.2	13.0	25.90	0.35	50.7	19.5	26.41	0.31	41.4	13.7	25.19	0.37	43.4	13.4	26.72	0.37

DDIM schedule. Here we proceed to compare the behavior of our algorithm under different schedules $(\eta_t)_{t \in [0,1]}$. For this purpose we compare against some natural candidates. (A): we consider the DDPM schedule used in (Ho et al., 2020) and which corresponds to using in (2.2) a standard deviation that depends on both s and t , i.e., $\eta_s(t) = \sigma_s(\sigma_t^2 - (\alpha_t/\alpha_s)^2\sigma_s^2)^{1/2}/\sigma_t$. (B): as we cannot use deterministic sampling in our approach, we rescale the DDPM schedule (A) with 0.01 to approach deterministic sampling. (C): $\eta_s = \sigma_s$, which is the maximum allowed standard deviation in (2.2). In this scenario, the transition is $p_{s|t}^{\theta}(\mathbf{x}_s | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_s; \hat{\mathbf{x}}_0^{\theta}(\mathbf{x}_t, t), \sigma_s^2 \mathbf{I}_d)$ and resembles the prior transition used in Martin et al. (2025). (D): $\eta_s = \sigma_s \sqrt{1 - \alpha_s}$, which corresponds to a slower decay of the standard deviation compared to our default choice $\eta_s = \sigma_s(1 - \alpha_s)$. The results on FFHQ with 5k samples and 50 NFEs are reported in Table 7. We observe that the rescaled DDPM schedule (B) degrades significantly across all metrics, while (A) and (C) yield nearly identical performance. This suggests that maintaining sufficient stochasticity at the beginning of the diffusion process is crucial for strong performance. Among the alternatives, (D) performs best; still, it is outperformed by our default schedule, confirming the benefit of a faster decay of (η_t) .

Table 7: Ablation results for the DDIM schedule (η_t) on FFHQ 768 × 768 with 5k samples and 50 NFEs.

Method	Half				Center				Top				Bottom			
	FID	pFID	cPSNR	LPIPS	FID	pFID	cPSNR	LPIPS	FID	pFID	cPSNR	LPIPS	FID	pFID	cPSNR	LPIPS
(A)	13.9	14.0	31.19	0.36	19.1	25.5	31.50	0.30	21.8	18.9	31.80	0.38	35.1	15.1	30.70	0.35
(B)	21.5	18.7	26.06	0.41	29.0	31.9	26.23	0.35	31.7	21.1	26.64	0.41	48.4	28.6	25.56	0.40
(C)	13.9	14.2	31.19	0.36	19.1	25.5	31.50	0.30	21.8	19.0	31.80	0.38	35.1	15.6	30.70	0.35
(D)	10.2	10.7	31.33	0.33	16.7	19.0	31.70	0.27	19.6	15.7	31.95	0.35	31.6	12.0	30.81	0.32
Default	9.6	6.6	31.03	0.33	15.5	14.0	31.38	0.27	19.7	12.5	31.64	0.34	29.6	8.6	30.50	0.32

5 CONCLUSION

We have introduced DING, a novel diffusion-based method for zero-shot inpainting that operates fully in the latent space and enables fast, memory-efficient inference under low-NFE budgets. Through extensive experiments across multiple benchmarks, we have shown that DING consistently outperforms existing zero-shot approaches and even surpasses a fine-tuned Stable Diffusion 3 model for image editing, despite requiring no expensive training. Notably, our method produces globally coherent reconstructions while preserving the visible content with high fidelity.

Limitations and future directions. While these results highlight the effectiveness and practicality of DING, several avenues remain open. An important limitation of our current approach is that performance does not monotonically improve as the compute budget increases. Ideally, one would like reconstruction accuracy to keep improving with additional sampling steps, potentially beyond the standard diffusion horizon, but we observe diminishing returns due to the limitations of our current DDIM schedule. Addressing this issue, for example by designing guidance schemes or noise schedules that continue to scale gracefully with compute, remains an important direction for future work. Moreover, while our framework is fully operational in the latent space, its applicability is

486 currently limited to inpainting, as this is the only observation operator we can reliably lift to the
487 latent domain. Extending the method to accommodate more general forward operators and a broader
488 class of inverse problems, while preserving the same level of efficiency achieved for inpainting, is a
489 challenging yet promising direction for future research.
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 **Reproducibility statement.** We place strong emphasis on reproducibility. To this end, we provide
 541 the full source code of our method along with implementations of all baseline methods used in the
 542 paper. Our repository also includes scripts to reproduce every experiment, as well as configuration
 543 files specifying all hyperparameters and settings for each baseline and experimental setup. Together,
 544 these resources ensure that all results reported in this work can be fully reproduced and easily
 545 extended.

546
 547 **Ethics statement.** While the proposed approach demonstrates clear benefits for applications in
 548 restoration, accessibility, and creative media, it also lies at the borderline of ethical considerations.
 549 Diffusion-based inpainting methods can be misappropriated for producing deceptive or harmful
 550 content, such as manipulated images or synthetic media that obscure authenticity. This dual-use
 551 nature highlights the need for proactive safeguards, including transparent usage guidelines, traceable
 552 model outputs, and continued development of forensic detection tools to ensure responsible integration
 553 of such technologies.

554
 555 **REFERENCES**
 556

557 Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset
 558 and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition*
 559 workshops, pp. 126–135, 2017.

560 Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural
 561 images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 562 pp. 18208–18218, 2022.

563 Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ACM Trans. Graph.*, 42
 564 (4), jul 2023. ISSN 0730-0301. doi: 10.1145/3592450. URL <https://doi.org/10.1145/3592450>.

565 Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas
 566 Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *Proceedings of the*
 567 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 843–852, 2023.

568 Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn,
 569 Jack English, Zion English, Patrick Esser, Sumith Kulal, et al. Flux. 1 kontext: Flow matching for
 570 in-context image generation and editing in latent space. *arXiv e-prints*, pp. arXiv–2506, 2025.

571 Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and*
 572 *Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. ISBN 0387310738.

573 Benjamin Boys, Mark Girolami, Jakiw Pidstrigach, Sebastian Reich, Alan Mosca, and O Deniz Aky-
 574 ildiz. Tweedie moment projected diffusions for inverse problems. *arXiv preprint arXiv:2310.06721*,
 575 2023.

576 Gabriel Cardoso, Yazid Janati El Idrissi, Sylvain Le Corff, and Eric Moulines. Monte Carlo guided
 577 diffusion for Bayesian linear inverse problems. *arXiv preprint arXiv:2308.07983*, 2023.

578 Lucy Chai, Michael Gharbi, Eli Shechtman, Phillip Isola, and Richard Zhang. Any-resolution training
 579 for high-resolution image synthesis. In *European conference on computer vision*, pp. 170–188.
 580 Springer, 2022.

581 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye.
 582 Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International*
 583 *Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=OnD9zGAGT0k>.

584 Hyungjin Chung, Jeongsol Kim, and Jong Chul Ye. Diffusion models for inverse problems. *arXiv*
 585 *preprint arXiv:2508.01975*, 2025.

594 Adrien Corenflos, Zheng Zhao, Thomas B. Schön, Simo Särkkä, and Jens Sjölund. Conditioning
 595 diffusion models by explicit forward-backward bridging. In Yingzhen Li, Stephan Mandt, Shipra
 596 Agrawal, and Emtiyaz Khan (eds.), *Proceedings of The 28th International Conference on Artifi-
 597 cial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp.
 598 3709–3717. PMLR, 03–05 May 2025. URL [https://proceedings.mlr.press/v258/
 599 corenflos25a.html](https://proceedings.mlr.press/v258/corenflos25a.html).

600 Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar,
 601 Alexandros G Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems.
 602 *arXiv preprint arXiv:2410.00083*, 2024.

603

604 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances
 605 in neural information processing systems*, 34:8780–8794, 2021.

606

607 Zehao Dou and Yang Song. Diffusion posterior sampling for linear inverse problem solving: A
 608 filtering perspective. In *The Twelfth International Conference on Learning Representations*, 2024.
 609 URL <https://openreview.net/forum?id=tplXNcHZs1>.

610 Julius Erbach, Dominik Narnhofer, Andreas Dombos, Bernt Schiele, Jan Eric Lenssen, and Konrad
 611 Schindler. Solving inverse problems with flair. *arXiv preprint arXiv:2506.02680*, 2025.

612

613 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 614 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for
 615 high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

616

617 Marc Anton Finzi, Anudhyan Boral, Andrew Gordon Wilson, Fei Sha, and Leonardo Zepeda-Núñez.
 618 User-defined event sampling and uncertainty quantification in diffusion models for physical
 619 dynamical systems. In *International Conference on Machine Learning*, pp. 10136–10152. PMLR,
 620 2023.

621 Ruiqi Gao, Emiel Hoogeboom, Jonathan Heek, Valentin De Bortoli, Kevin P Murphy, and Tim Sal-
 622 mians. Diffusion meets flow matching: Two sides of the same coin. 2024. URL [https://diffusionflow.github.io](https://diffusionflow.

 623 github.io), 2024.

624

625 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 626 Neural Information Processing Systems*, 33:6840–6851, 2020.

627

628 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J
 629 Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646,
 630 2022.

631 Yi Huang, Jiancheng Huang, Yifan Liu, Mingfu Yan, Jiaxi Lv, Jianzhuang Liu, Wei Xiong, He Zhang,
 632 Liangliang Cao, and Shifeng Chen. Diffusion model-based image editing: A survey. *IEEE
 633 Transactions on Pattern Analysis and Machine Intelligence*, 2025.

634

635 Yazid Janati, Badr Moufad, Alain Durmus, Eric Moulines, and Jimmy Olsson. Divide-and-conquer
 636 posterior sampling for denoising diffusion priors. In *Neurips*, 2024.

637

638 Yazid Janati, Badr Moufad, Mehdi Abou El Qassim, Alain Durmus, Eric Moulines, and Jimmy
 639 Olsson. A mixture-based framework for guiding diffusion models. *preprint*, 2025a.

640

641 Yazid Janati, Eric Moulines, Jimmy Olsson, and Alain Oliviero-Durmus. Bridging diffusion posterior
 642 sampling and monte carlo methods: a survey. *Philosophical Transactions A*, 383(2299):20240331,
 643 2025b.

644

645 Xuan Ju, Ailing Zeng, Yuxuan Bian, Shaoteng Liu, and Qiang Xu. Pnp inversion: Boosting
 646 diffusion-based editing with 3 lines of code. In *The Twelfth International Conference on Learning
 647 Representations*, 2024. URL <https://openreview.net/forum?id=FoMZ41jhVw>.

648

649 Zahra Kadkhodaie and Eero P Simoncelli. Solving linear inverse problems using the prior implicit in
 650 a denoiser. *arXiv preprint arXiv:2007.13640*, 2020.

648 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
 649 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
 650 *recognition*, pp. 4401–4410, 2019.

651

652 Bahjat Kawar, Gregory Vaksman, and Michael Elad. Snips: Solving noisy inverse problems stochas-
 653 tically. *Advances in Neural Information Processing Systems*, 34:21757–21769, 2021.

654 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration
 655 models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022.

656

657 Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and
 658 Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the*
 659 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 6007–6017, 2023.

660

661 Jeongsol Kim, Bryan Sangwoo Kim, and Jong Chul Ye. Flowdps: Flow-driven posterior sampling
 662 for inverse problems. *arXiv preprint arXiv:2503.08136*, 2025.

663 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 664 pre-training with frozen image encoders and large language models. In *International conference*
 665 *on machine learning*, pp. 19730–19742. PMLR, 2023.

666

667 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
 668 matching for generative modeling. In *The Eleventh International Conference on Learning Repre-
 669 sentations*, 2023. URL <https://openreview.net/forum?id=PqvMRDCJT9t>.

670 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool.
 671 Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the*
 672 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11461–11471, 2022.

673

674 Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving
 675 inverse problems with diffusion models. In *The Twelfth International Conference on Learning*
 676 *Representations*, 2024. URL <https://openreview.net/forum?id=1YO4EE3SPB>.

677 Ségolène Tiffany Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. Pnp-flow: Plug-and-
 678 play image restoration with flow matching. In *The Thirteenth International Conference on Learning*
 679 *Representations*, 2025. URL <https://openreview.net/forum?id=5AtHrq3B5R>.

680

681 Chenlin Meng, Yang Song, Wenzhe Li, and Stefano Ermon. Estimating high order gradients of
 682 the data distribution by denoising. *Advances in Neural Information Processing Systems*, 34:
 683 25359–25369, 2021.

684

685 Badr Moufad, Yazid Janati, Lisa Bedin, Alain Oliviero Durmus, randal douc, Eric Moulines, and
 686 Jimmy Olsson. Variational diffusion posterior sampling with midpoint guidance. In *The Thirteenth*
 687 *International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=6EUTjXAvmj>.

688

689 Kushagra Pandey, Farrin Marouf Sofian, Felix Draxler, Theofanis Karaletsos, and Stephan Mandt.
 690 Variational control for guidance in diffusion models. In *Forty-second International Conference on*
 691 *Machine Learning*, 2025. URL <https://openreview.net/forum?id=Z0ffRRtOim>.

692

693 Maitreya Patel, Song Wen, Dimitris N. Metaxas, and Yezhou Yang. Steering rectified flow models in
 694 the vector field for controlled image generation. *arXiv preprint arXiv:2412.00100*, 2024.

695

696 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 697 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 698 models from natural language supervision. In *International conference on machine learning*, pp.
 699 8748–8763. PMLR, 2021.

700

701 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 702 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Confer-
 703 ence on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

702 Litu Rout, Yujia Chen, Nataniel Ruiz, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng
 703 Chu. Semantic image inversion and editing using rectified stochastic differential equations. *arXiv*
 704 *preprint arXiv:2410.10792*, 2024a.

705 Litu Rout, Negin Raoof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai.
 706 Solving linear inverse problems provably via posterior sampling with latent diffusion models.
 707 *Advances in Neural Information Processing Systems*, 36, 2024b.

708 François Rozet, Gérôme Andry, François Lanusse, and Gilles Louppe. Learning diffusion priors from
 709 observations by expectation maximization. *Advances in Neural Information Processing Systems*,
 710 37:87647–87682, 2024.

711 Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet,
 712 and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022*
 713 conference proceedings, pp. 1–10, 2022.

714 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
 715 learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*,
 716 pp. 2256–2265. PMLR, 2015.

717 Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse
 718 problems with latent diffusion models via hard data consistency. In *The Twelfth International*
 719 *Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=j8hdRqOUhN>.

720 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International*
 721 *Conference on Learning Representations*, 2021a. URL <https://openreview.net/forum?id=St1giarCHLP>.

722 Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion
 723 models for inverse problems. In *International Conference on Learning Representations*, 2023a.
 724 URL https://openreview.net/forum?id=9_gsMA8MRKQ.

725 Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin
 726 Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation.
 727 In *International Conference on Machine Learning*, pp. 32483–32498. PMLR, 2023b.

728 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.
 729 *Advances in neural information processing systems*, 32, 2019.

730 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 731 Poole. Score-based generative modeling through stochastic differential equations. In *International*
 732 *Conference on Learning Representations*, 2021b.

733 Alessio Spagnoletti, Jean Prost, Andrés Almansa, Nicolas Papadakis, and Marcelo Pereyra. Latino-
 734 pro: Latent consistency inverse solver with prompt optimization. *arXiv preprint arXiv:2503.12615*,
 735 2025.

736 Tristan SW Stevens, Hans van Gorp, Faik C Meral, Junseob Shin, Jason Yu, Jean-Luc Robert,
 737 and Ruud JG van Sloun. Removing structured noise with diffusion models. *arXiv preprint*
 738 *arXiv:2302.05290*, 2023.

739 Brian L. Trippe, Jason Yim, Doug Tischer, David Baker, Tamara Broderick, Regina Barzilay, and
 740 Tommi S. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-
 741 scaffolding problem. In *The Eleventh International Conference on Learning Representations*, 2023.
 742 URL <https://openreview.net/forum?id=6TxBxqNME1Y>.

743 Su Wang, Chitwan Saharia, Ceslee Montgomery, Jordi Pont-Tuset, Shai Noy, Stefano Pellegrini,
 744 Yasumasa Onoe, Sarah Laszlo, David J Fleet, Radu Soricut, et al. Imagen editor and editbench: Ad-
 745 vancing and evaluating text-guided image inpainting. In *Proceedings of the IEEE/CVF conference*
 746 *on computer vision and pattern recognition*, pp. 18359–18369, 2023a.

756 Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion
 757 null-space model. In *The Eleventh International Conference on Learning Representations*, 2023b.
 758 URL <https://openreview.net/forum?id=mRieQgMtNTQ>.

759

760 Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai
 761 Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*,
 762 2025.

763 Luhuan Wu, Brian L. Trippe, Christian A Næsseth, John Patrick Cunningham, and David Blei.
 764 Practical and asymptotically exact conditional sampling in diffusion models. In *Thirty-seventh
 765 Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=eWKqr1zcRv>.

766

767 Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song.
 768 Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of
 769 the Computer Vision and Pattern Recognition Conference*, pp. 20895–20905, 2025.

770

771 Guanhua Zhang, Jiabao Ji, Yang Zhang, Mo Yu, Tommi Jaakkola, and Shiyu Chang. Towards
 772 coherent image inpainting using denoising diffusion implicit models. In *International Conference
 773 on Machine Learning*, pp. 41164–41193. PMLR, 2023.

774

775 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 776 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on
 777 computer vision and pattern recognition*, pp. 586–595, 2018.

778

779 Zheng Zhao. Generative diffusion posterior sampling for informative likelihoods. *arXiv preprint
 arXiv:2506.01083*, 2025.

780

781 Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jie Zhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool.
 782 Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF
 783 Conference on Computer Vision and Pattern Recognition*, pp. 1219–1229, 2023.

784

785 Nicolas Zilberstein, Morteza Mardani, and Santiago Segarra. Repulsive latent score distillation for
 786 solving inverse problems. In *The Thirteenth International Conference on Learning Representations*,
 2025. URL <https://openreview.net/forum?id=bwJxUB0y46>.

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809



Figure 3: Latent-space masking and its correspondence to pixel space using a central square mask. The encoder and decoder of Stable Diffusion 3.5 (medium) were used. The first row shows latent images alongside the encoded mask applied to each, while the second row shows their decoded counterparts. Notice that the masked regions in the latent space translate directly to analogous masked regions in pixel space. For that sake of visualization, since the latent images have 16 channels, we apply PCA and visualize the first 3 components.

A METHOD DETAILS

A.1 DERIVATION OF THE POSTERIOR (3.1)

Recall that given \mathbf{z}_s , the posterior transition of interest is

$$\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y}) \propto \ell_s^\theta(\mathbf{y}|\mathbf{x}_s, \mathbf{z}_s) p_{s|t}^{\eta, \theta}(\mathbf{x}_s|\mathbf{x}_t).$$

Denoting by $\tilde{\mathbf{y}}_s = \alpha_s \mathbf{y} + \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{z}_s, s)[\bar{\mathbf{m}}]$ the effective observation, we have that

$$\ell_s^\theta(\mathbf{y}|\mathbf{x}_s, \mathbf{z}_s) \propto N(\tilde{\mathbf{y}}_s; \mathbf{x}_s[\bar{\mathbf{m}}], \alpha_s^2 \sigma_y^2 \mathbf{I}_{d_y}),$$

and since the reverse transition writes

$$p_{s|t}^{\eta, \theta}(\mathbf{x}_s|\mathbf{x}_t) = N(\mathbf{x}_s[\mathbf{m}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 \mathbf{I}_{d-d_y}) N(\mathbf{x}_s[\bar{\mathbf{m}}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}], \eta_s^2 \mathbf{I}_{d_y}),$$

we obtain

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y}) &= N(\mathbf{x}_s[\mathbf{m}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 \mathbf{I}_{d-d_y}) \\ &\times \frac{N(\tilde{\mathbf{y}}_s; \mathbf{x}_s[\bar{\mathbf{m}}], \alpha_s^2 \sigma_y^2 \mathbf{I}_{d_y}) N(\mathbf{x}_s[\bar{\mathbf{m}}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}], \eta_s^2 \mathbf{I}_{d_y})}{\int N(\tilde{\mathbf{y}}_s; \tilde{\mathbf{x}}_s[\bar{\mathbf{m}}], \alpha_s^2 \sigma_y^2 \mathbf{I}_{d_y}) N(\tilde{\mathbf{x}}_s[\bar{\mathbf{m}}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}], \eta_s^2 \mathbf{I}_{d_y}) d\tilde{\mathbf{x}}_s[\bar{\mathbf{m}}]}. \end{aligned}$$

The formula (3.4) follows by applying (Bishop, 2006, equation 2.116) to the second normalized transition on the right-hand side.

A.2 COMPARISON WITH RELATED WORKS

We start by providing an explicit comparison with the closest works.

Comparison with the transition in Cardoso et al. (2023). Let $\tau \in [0, 1]$ be a timestep such that $\sigma_y = \sigma_\tau / \alpha_\tau$. Such a τ always exists when the linear schedule is used for example. The transition used in the SMC algorithm in Cardoso et al. (2023) for $s > \tau$ is given by

$$\hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) \propto N(\alpha_s \mathbf{y}; \mathbf{x}_s[\bar{\mathbf{m}}], \sigma_{s|\tau}^2 \mathbf{I}_{d_y}) p_{s|t}^{\eta, \theta}(\mathbf{x}_s|\mathbf{x}_t) \quad (\text{A.1})$$

Using the same conjugation formulas as in the previous section, we find that

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) &= N(\mathbf{x}_s[\mathbf{m}], \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 \mathbf{I}_{d-d_y}) \\ &\times N(\mathbf{x}_s[\bar{\mathbf{m}}], (1 - \tilde{\gamma}_{s|\tau}) \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}] + \tilde{\gamma}_{s|\tau} \alpha_s \mathbf{y}, \sigma_{t|\tau}^2 \tilde{\gamma}_{s|\tau} \mathbf{I}_{d-d_y}), \quad (\text{A.2}) \end{aligned}$$

where $\sigma_{t|\tau}^2 := \sigma_t^2 - (\alpha_t / \alpha_\tau)^2 \sigma_\tau^2$ and $\tilde{\gamma}_{s|t} = \eta_s^2 / (\eta_s^2 + \sigma_{t|\tau}^2)$. This is to be contrasted with our update, given a sample \mathbf{z}_s ,

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y}) &= N(\mathbf{x}_s[\mathbf{m}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 \mathbf{I}_{d-d_y}) \\ &\times N(\mathbf{x}_s[\bar{\mathbf{m}}]; (1 - \gamma_{s|\tau}) \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}] + \gamma_{s|\tau} (\alpha_s \mathbf{y} + \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{z}_s, s)[\bar{\mathbf{m}}]), \alpha_s^2 \sigma_y^2 \gamma_{s|\tau} \mathbf{I}_{d_y}), \end{aligned}$$

where $\gamma_{s|\tau} := \eta_s^2 / (\eta_s^2 + \alpha_s^2 \sigma_y^2)$. Hence, MCGDIFF differs from DING on the choice of effective observation, which in this case is $\tilde{\mathbf{y}}_s = \alpha_s \mathbf{y}$, the choice of variance in the transition and the coefficient of the convex combination. From (A.1) it can be seen that MCGDIFF assumes the approximate model $N(\alpha_s \mathbf{y}; \mathbf{x}_s[\bar{\mathbf{m}}], \sigma_{s|\tau}^2 \mathbf{I}_{d_y})$ for the true likelihood $\ell_s(\mathbf{y}|\mathbf{x}_s)$ (2.7).

864 **Comparison with the transition in Zhu et al. (2023) and Martin et al. (2025).** We now write
 865 explicitely the algorithm PNP-FLOW (Martin et al., 2025, Algorithm 3) adapted to the inpainting
 866 problem we consider; see Algorithm 2. We have simply adapted the notations and used $F(\mathbf{x}) =$
 867 $\|\mathbf{y} - \mathbf{x}[\bar{\mathbf{m}}]\|^2 / (2\sigma_y^2)$ in (Martin et al., 2025, Algorithm 3). Thus, the transition used in Algorithm 2 is
 868

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) &\propto N(\mathbf{x}_s[\mathbf{m}], \alpha_s \hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)[\mathbf{m}], \sigma_s^2 I_{d-d_y}) \\ &\times N\left(\mathbf{x}_s[\bar{\mathbf{m}}], \left(1 - \frac{\gamma_s}{\sigma_y^2}\right) \alpha_s \hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)[\bar{\mathbf{m}}] + \frac{\gamma_s}{\sigma_y^2} \alpha_s \mathbf{y}, \sigma_s^2 I_{d-d_y}\right). \end{aligned}$$

873 In the case of the DDIM schedule $\eta_s = \sigma_s$, we have that $\mu_{s|t}^\theta(\mathbf{x}_t) = \alpha_s \hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)$, and the MCGDIFF
 874 transition (Cardoso et al., 2023) in (A.2) writes

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) &= N(\mathbf{x}_s[\mathbf{m}], \alpha_s \hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)[\mathbf{m}], \sigma_s^2 I_{d-d_y}) \\ &\times N(\mathbf{x}_s[\bar{\mathbf{m}}], (1 - \tilde{\gamma}_{s|t}) \alpha_s \hat{\mathbf{x}}_0^\theta(\mathbf{x}_t, t)[\bar{\mathbf{m}}] + \tilde{\gamma}_{s|t} \alpha_s \mathbf{y}, \sigma_{s|t}^2 \tilde{\gamma}_{s|t} I_{d-d_y}). \end{aligned}$$

875 Hence, the main difference lies in the coefficient of the convex combination and the variance used.
 876

877 **Algorithm 2** PNP-FLOW reinterpreted

878 1: **Input:** Decreasing timesteps $(t_k)_{k=K}^0$ with $t_K = 1$, $t_0 = 0$; adaptive stepsizes $(\gamma_k)_{k=K}^0$.
 879 2: **Initialize:** $\hat{\mathbf{x}}_0 \in \mathbb{R}^d$.
 880 3: **for** $k = K - 1$ **to** 1 **do**
 881 4: $\hat{\mathbf{x}}_0[\bar{\mathbf{m}}] \leftarrow (1 - \frac{\gamma_k}{\sigma_y^2}) \hat{\mathbf{x}}_0[\bar{\mathbf{m}}] + \frac{\gamma_k}{\sigma_y^2} \mathbf{y}$
 882 5: $\mathbf{w} \sim \mathcal{N}(0_d, I_d)$
 883 6: $\mathbf{x} \leftarrow \alpha_{t_k} \hat{\mathbf{x}}_0 + \sigma_{t_k} \mathbf{w}$
 884 7: $\hat{\mathbf{x}}_0 \leftarrow \hat{\mathbf{x}}_0^\theta(\mathbf{x}, t_k)$
 885 8: **end for**
 886 9: **Return:** $\hat{\mathbf{x}}_0$

887 892 **Comparison with the transition in Kim et al. (2025); Patel et al. (2024).** Here we explicitely
 893 write the transition of FLOWDPS for the inpainting case in order to understand the main differences
 894 without our method. For this purpose we rewrite (Kim et al., 2025, Algorithm 1) using our notations.
 895 We note that the algorithm is written for the linear schedule $\alpha_t = 1 - t$, $\sigma_t = t$ and the choice
 896 of DDIM schedule $\eta_t = \sigma_t \sqrt{1 - \sigma_t}$, but we still write it with general notations to streamline the
 897 comparison with Algorithm 1. We also assume for the sake of simplicity that the optimization
 898 problem is solved exactly in (Kim et al., 2025, line 7) (since there is no decoder as we solve the
 899 inverse problem in the latent space). The algorithm is given in Algorithm 3. In the specific setting
 900 where the linear schedule is used, setting $\gamma_k = \sigma_y^2 \sigma_{t_k}$ in Algorithm 2 recovers Algorithm 3 when
 901 $\eta_k = \sigma_{t_k}$. Finally, we note that FLOWDPS can be understood as a noisy version of the FLOWCHEF
 902 algorithm (Patel et al., 2024) and overall, follows the line of work of methods that learn a residual
 903 that is then used to translate the denoiser (Bansal et al., 2023; Zhu et al., 2023).
 904

905 **Comparison with DiffPIR (Zhu et al., 2023) and DDNM (Wang et al., 2023b).** We provide the
 906 DIFFPIR algorithm (Zhu et al., 2023, Algorithm 1) adapted to our inpainting case using our own
 907 notation in Algorithm 4. In Line 6 we write the exact solution to the optimization problem in the
 908 original algorithm. We write the associated transition in a convenient form that allows a seamless
 909 comparison with our algorithm. Define $\gamma_t := \sigma_t^2 / (\sigma_t^2 + \lambda \alpha_t^2 \sigma_y^2)$. Then, the transition

$$\begin{aligned} \hat{\pi}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) &= N(\mathbf{x}_s[\mathbf{m}]; \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\mathbf{m}], \eta_s^2 I_{d-d_y}) \\ &\times N\left(\mathbf{x}_s[\bar{\mathbf{m}}]; (1 - \gamma_t) \mu_{s|t}^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}] + \gamma_t (\alpha_s \mathbf{y} + (\sigma_s^2 - \eta_s^2)^{1/2} \frac{\mathbf{x}_t[\bar{\mathbf{m}}] - \alpha_t \mathbf{y}}{\sigma_t}), \eta_s^2 I_{d_y}\right), \end{aligned}$$

910 914 corresponds to one step of Algorithm 4. We highlight key distinctions:
 911

- 912 • Setting $\eta_s^2 = \sigma_s^2$ recovers the same transition as in PNP-FLOW.
 913 • The main distinction lies in the mean of the Gaussian transition for the unmasked region: it is a
 914 convex combination of $\mu^\theta(\mathbf{x}_t; \eta)[\bar{\mathbf{m}}]$ and an effective observation $\alpha_s \mathbf{y} + (\sigma_s^2 - \eta_s^2)^{1/2} (\mathbf{x}_t[\bar{\mathbf{m}}] -$

918 **Algorithm 3** FLOWDPS reinterpreted

```

919
920 1: Input: decreasing timesteps  $(t_k)_{k=K}^0$  with  $t_K = 1, t_0 = 0$ ; original image  $\mathbf{x}_*$ ; mask  $\mathbf{m}$ ;
921 2:  $\mathbf{y} \leftarrow \mathbf{x}_*[\mathbf{m}]$ 
922 3:  $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)$ 
923 4: for  $k = K - 1$  to 0 do
924 5:  $\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}_0^\theta(\mathbf{x}, t_{k+1})$ 
925 6:  $\hat{\mathbf{x}}_1 \leftarrow (\mathbf{x} - \alpha_{t_{k+1}} \hat{\mathbf{x}}_0) / \sigma_{t_{k+1}}$ 
926 7:  $\hat{\mathbf{x}}_0[\mathbf{m}] \leftarrow \alpha_{t_k} \hat{\mathbf{x}}_0[\mathbf{m}] + \sigma_{t_k} \mathbf{y}$ 
927 8:  $\boldsymbol{\mu} \leftarrow \alpha_{t_k} \hat{\mathbf{x}}_0 + (\sigma_{t_k}^2 - \eta_k^2)^{1/2} \hat{\mathbf{x}}_1$ 
928 9:  $\mathbf{w} \sim \mathcal{N}(0_d, \mathbf{I}_d)$ 
929 10:  $\mathbf{x} \leftarrow \boldsymbol{\mu} + \eta_k \mathbf{w}$ 
930 11: end for
931 12: Return:  $\mathbf{x}$ 

```

933 $\alpha_t \mathbf{y}) / \sigma_t$. In our algorithm, the effective observation instead takes the form $\alpha_s \mathbf{y} + \sigma_s \hat{\mathbf{x}}_1^\theta(\mathbf{x}_s, s)[\mathbf{m}]$.
934 We estimate the residual noise using the pre-trained model at timestep s , whereas DIFFPIR
935 computes it as $(\mathbf{x}_t[\mathbf{m}] - \alpha_t \mathbf{y}) / \sigma_t$.

936

- 937 • This residual noise is scaled differently: by $(\sigma_s^2 - \eta_s^2)^{1/2}$ in DIFFPIR, and by σ_s in our method.
- 938 • The convex combination coefficient in our cases is $\gamma_{s|t} = \eta_s^2 / (\eta_s^2 + \alpha_s^2 \sigma_y^2)$ whereas for DIFFPIR it
939 is set to $\gamma_t = \sigma_t^2 / (\sigma_t^2 + \alpha_t^2 \sigma_y^2)$.
- 940 • Finally, the noise-free ($\sigma_y = 0$) version of DIFFPIR recovers the DDNM algorithm (Zhang et al.,
941 2023).

943 **Algorithm 4** DIFFPIR reinterpreted

```

944
945 1: Input: Decreasing timesteps  $(t_k)_{k=K}^0$  with  $t_K = 1, t_0 = 0$ ; scaling  $\lambda$ ; original image  $\mathbf{x}_*$ ; mask
946 2:  $\mathbf{y} \leftarrow \mathbf{x}_*[\mathbf{m}]$ 
947 3:  $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)$ .
948 4: for  $k = K - 1$  to 1 do
949 5:  $\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}_0^\theta(\mathbf{x}, t_{k+1})$ 
950 6:  $\hat{\mathbf{x}}_0[\mathbf{m}] \leftarrow \frac{\sigma_{t_{k+1}}^2}{\sigma_{t_{k+1}}^2 + \lambda \sigma_y^2 \alpha_{t_{k+1}}^2} \mathbf{y} + \frac{\lambda \sigma_y^2 \alpha_{t_{k+1}}^2}{\sigma_{t_{k+1}}^2 + \lambda \sigma_y^2 \alpha_{t_{k+1}}^2} \hat{\mathbf{x}}_0[\mathbf{m}]$ 
951 7:  $\hat{\mathbf{x}}_1 \leftarrow (\mathbf{x} - \alpha_{t_{k+1}} \hat{\mathbf{x}}_0) / \sigma_{t_{k+1}}$ 
952 8:  $\mathbf{w} \sim \mathcal{N}(0_d, \mathbf{I}_d)$ 
953 9:  $\mathbf{x} \leftarrow \alpha_{t_k} \hat{\mathbf{x}}_0 + (\sigma_{t_k}^2 - \eta_k^2)^{1/2} \hat{\mathbf{x}}_1 + \eta_k \mathbf{w}$ 
954 10: end for
955 11: Return:  $\mathbf{x}$ 

```

956

957 **Further related methods.** Here we continue our discussion of VJP-free methods. The DAPS
958 algorithm (Zhang et al., 2025) proposes sampling, given the previous state $X_{t_{k+1}}$, a clean state \hat{X}_0
959 by performing Langevin Monte Carlo steps on the posterior distribution $\pi_{0|t_{k+1}}(\cdot | X_{t_{k+1}}, \mathbf{y})$. This
960 step is performed approximately by replacing the prior transition $p_{0|t_{k+1}}(\cdot | X_{t_{k+1}})$ with a Gaussian
961 approximation centered at the denoiser $\hat{\mathbf{x}}_0^\theta(X_{t_{k+1}}, t_{k+1})$. Then, given \hat{X}_0 , the next state is drawn
962 from $\mathcal{N}(\alpha_{t_k} \hat{X}_0, \sigma_{t_k}^2 \mathbf{I}_d)$. Then the transition

963

964 One important aspect of our method is that we circumvent differentiation through the denoiser but
965 also the decoder, as the diffusion models we consider operate in the latent space. We do so by
966 downsampling the mask into the latent space. In contrast, the recent work of Spagnoletti et al. (2025)
967 also circumvents differentiation through the denoiser, but does so by lifting the latent states into pixel
968 space and optimizing the likelihood there. The result of the optimization is then projected back into
969 the latent space and then undergoes back-and-forth noise-denoising steps.

970 Finally, several recent works (Mardani et al., 2024; Zilberstein et al., 2025; Erbach et al., 2025) adopt
971 a variational perspective: the target distribution is approximated by a Gaussian distribution whose

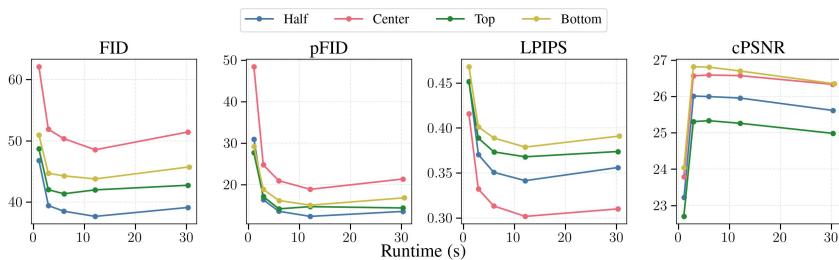
972 parameters are iteratively estimated by minimizing a combination of an observation-fidelity loss and
 973 a score-matching-like loss.

974 *VJP-based methods.* A broad class of zero-shot approaches builds on the guidance approximation
 975 (2.8) to estimate $\nabla_{\mathbf{x}_t} \log \ell_t(\mathbf{y}|\mathbf{x}_t)$. [Song et al. \(2023a\)](#) approximate $p_{0|t}$ by a Gaussian with mean
 976 $\hat{\mathbf{x}}_0^\theta(\cdot, t)$ and a tuned covariance. For the inpainting setting in (2.4), plugging this approximation into
 977 (2.7) yields an integral that can be computed in a closed form, providing a proxy for $\ell_t(\mathbf{y}|\cdot)$. Several
 978 works exploit the link between the covariance of $p_{0|t}(\cdot|\mathbf{x}_t)$ and the Jacobian of the denoiser ([Meng
 979 et al., 2021](#)). This observation underpins the methods of [Finzi et al. \(2023\)](#), [Stevens et al. \(2023\)](#),
 980 and [Boys et al. \(2023\)](#), which derive likelihood scores by estimating or inverting the Jacobian. These
 981 approaches require solving large linear systems and backpropagating through the denoiser, both
 982 computationally expensive operations. To reduce cost, these works assume a locally constant Jacobian
 983 around \mathbf{x}_t , but updates still involve either explicit matrix inversion or repeated VJPs. In practice,
 984 diagonal approximations based on row sums are commonly used to approximate the covariance
 985 matrix ([Boys et al., 2023](#)), or conjugate gradient methods are employed to circumvent the need for full
 986 matrix inversion ([Rozet et al., 2024](#)). For general likelihoods $\ell_0(\mathbf{y}|\cdot)$, [Song et al. \(2023b\)](#) combine the
 987 Gaussian posterior model of [Song et al. \(2023a\)](#) with Monte Carlo sampling to approximate $\ell_t(\mathbf{y}|\cdot)$.
 988 In the latent setting, [Rout et al. \(2024b\)](#) apply the DPS approximation jointly with a regularizer that
 989 encourages latent variables to remain near encoder-decoder fixed points. Other methods modify the
 990 sampling dynamics. [Moufad et al. \(2025\)](#) propose a two-stage procedure: the chain is first moved
 991 to an earlier time $\ell \ll t_k$, where the DPS approximation is applied to sample from an approximate
 992 conditional at ℓ , before returning to t_k via additional noising steps. [Janati et al. \(2025a\)](#) incorporate
 993 a related idea into a Gibbs sampling framework. Overall, these methods remain fundamentally
 994 VJP-based and inherit substantial memory and runtime overhead from repeated backpropagation
 995 through the denoiser. By contrast, our decoupled guidance relies exclusively on forward denoiser
 996 evaluations and closed-form Gaussian updates, thereby eliminating VJPs entirely while retaining
 997 competitive performance.

998 For a complete review of zero-shot posterior sampling methods see [Daras et al. \(2024\)](#); [Janati et al.
 999 \(2025b\)](#); [Chung et al. \(2025\)](#).

1000 A.3 BEHAVIOR UNDER INCREASED RUNTIME.

1001 We extend the ablation study in Section 4.2 by examining the behavior of DING when the number
 1002 of NFEs is increased. Specifically, we vary the budget from 20 to 500 NFEs on the DIV2K dataset
 1003 and report results across different masking patterns; see Figure 4. All metrics improve steadily as
 1004 the budget grows, reaching their best values around 200 NFEs (10s runtime). Beyond this point,
 1005 performance saturates and exhibits a slight degradation at 500 NFEs. These results suggest that our
 1006 default DDIM schedule is well suited to low and mid-NFE regimes—which are most relevant for
 1007 practical settings—but may not be fully optimized for larger budgets.



1018 Figure 4: Performance of DING on DIV2K under varying NFE budgets (20 to 500) across different masking
 1019 patterns. Runtimes are measured on a H100 GPU.

1020 A.4 LIMITATION

1021 We observed that the quality of reconstructions is highly sensitive to the specificity of the textual
 1022 prompt. When the prompt is under-specified or lacks sufficient semantic detail, the resulting samples
 1023 may exhibit reduced coherence, particularly in large masked regions where contextual consistency



Figure 5: Effect of prompt precision on inpainting quality

is critical. This issue manifests as mismatched textures or backgrounds, or inconsistent object boundaries, even when the visible area is faithfully preserved. To illustrate this behavior, we compare reconstructions obtained with well-defined prompts against those generated using vague or ambiguous ones. Examples are provided in Figure 5 and 6.

A.5 BIAS IN GAUSSIAN CASE

For the sake of simplicity we assume that $p_0 := \mathcal{N}(0_d, \Sigma)$ where Σ is a covariance matrix. We also write the likelihood as $\ell_0(\mathbf{y}|\mathbf{x}_0) = \mathcal{N}(\mathbf{y}; P_{\bar{\mathbf{m}}}\mathbf{x}_0, \sigma_y^2 \mathbf{I}_{d_y})$ where $P_{\bar{\mathbf{m}}} \in \mathbb{R}^{d_y \times d}$ is the matrix satisfying $P_{\bar{\mathbf{m}}}\mathbf{x} = \mathbf{x}[\bar{\mathbf{m}}]$. Define $D_t := \alpha_t \Sigma (\alpha_t^2 \Sigma + \sigma_t^2 \mathbf{I}_d)^{-1}$. Then, the denoiser and noise predictors are given by

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, t) = D_t \mathbf{x}_t, \quad \hat{\mathbf{x}}_1(\mathbf{x}_t, t) = \sigma_t^{-1} (\mathbf{I}_d - \alpha_t D_t) \mathbf{x}_t.$$

We consider hereafter the DDIM transitions $p_{s|t}^\eta(\mathbf{x}_t|\mathbf{x}_s) := \mathcal{N}(\mathbf{x}_s; \mu_{s|t}(\mathbf{x}_s; \eta), \eta_s^2 \mathbf{I}_d)$ where

$$\mu_{s|t}(\mathbf{x}_s; \eta) := \alpha_s \hat{\mathbf{x}}_0(\mathbf{x}_s, t) + \sqrt{\sigma_s^2 - \eta_s^2} \hat{\mathbf{x}}_1(\mathbf{x}_s, t)$$

In this section we analyze the bias of the DING one-step transition relative to the posterior transition involving the DPS likelihood (2.8); *i.e.* we compare the transition

$$\hat{\pi}_{s|t}^{\text{ding}}(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) := \mathbb{E} \left[\hat{\pi}_{s|t}^{\text{ding}}(\mathbf{x}_s|Z_s, \mathbf{x}_t, \mathbf{y}) \right], \quad (\text{A.3})$$

where $Z_s \sim p_{s|t}^\eta(\cdot|\mathbf{x}_t)$ and

$$\hat{\pi}_{s|t}^{\text{ding}}(\mathbf{x}_s|\mathbf{z}_s, \mathbf{x}_t, \mathbf{y}) \propto \ell_0(\mathbf{y} \mid \frac{\mathbf{x}_s - \sigma_s \hat{\mathbf{x}}_1(\mathbf{z}_s, s)}{\alpha_s}) p_{s|t}^\eta(\mathbf{x}_s|\mathbf{x}_t)$$

against

$$\hat{\pi}_{s|t}^{\text{dps}}(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) \propto \ell_0(\mathbf{y} \mid \hat{\mathbf{x}}_0(\mathbf{x}_s, s)) p_{s|t}^\eta(\mathbf{x}_s|\mathbf{x}_t).$$

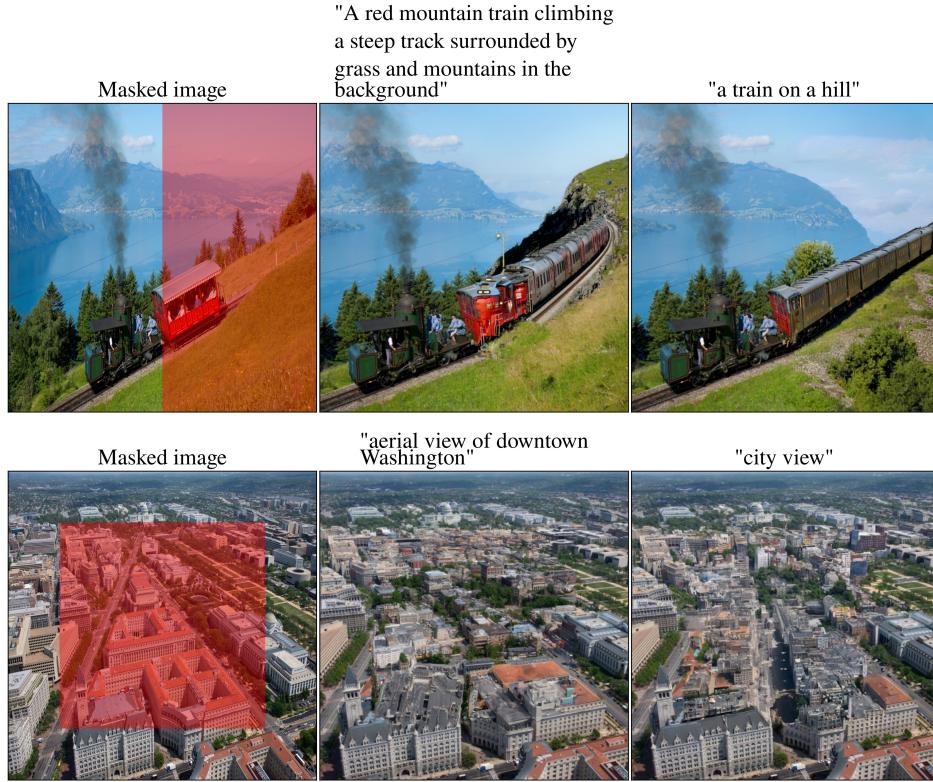


Figure 6: Effect of prompt precision on inpainting quality

We define $M := P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}}$, which is an orthogonal projection matrix since $M^\top = M$, $P_{\bar{\mathbf{m}}} P_{\bar{\mathbf{m}}}^\top = \mathbf{I}_{d_y}$, and thus $M^2 = M$. We also introduce the quantity

$$\varepsilon_s := \|(D_s^\top - \alpha_s^{-1} \mathbf{I}_d) M\|_{\text{op}},$$

which quantifies how far the Jacobian of the denoiser $\hat{\mathbf{x}}_0(\cdot, s)$ deviates from the Jacobian of the DING denoiser approximation *on the observed coordinates*. In the following proposition, we characterize the asymptotic behavior of the DPS and DING posterior transition means and covariances as $\eta_s \rightarrow 0$, and we express the mean bias in terms of ε_s . In Proposition 2, we also provide an explicit upper bound on ε_s in terms of the schedule and the minimum eigenvalue of the prior covariance Σ .

Proposition 1. *Both $\hat{\pi}_{s|t}^{\text{dps}}(\cdot | \mathbf{x}_t, \mathbf{y})$ and $\hat{\pi}_{s|t}^{\text{ding}}(\cdot | \mathbf{x}_t, \mathbf{y})$ are Gaussian distributions with mean and covariance respectively $(\boldsymbol{\mu}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}), \Sigma_{s|t}^{\text{dps}})$ and $(\boldsymbol{\mu}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y}), \Sigma_{s|t}^{\text{ding}})$ satisfying*

$$\|\Sigma_{s|t}^{\text{dps}} - \Sigma_{s|t}^{\text{ding}}\| = \mathcal{O}(\eta_s^4)$$

and

$$\|\boldsymbol{\mu}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}) - \boldsymbol{\mu}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y})\| = \mathcal{O}\left(\eta_s^2(\varepsilon_s(\|\mathbf{y}\| + \|M\boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\|) + \varepsilon_s^2\|\boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\|)\right).$$

as $\eta_s \rightarrow 0$.

Proof. Using the standard Gaussian conjugation formula (Bishop, 2006, equation 2.116), we have that $\hat{\pi}_{s|t}^{\text{dps}}(\mathbf{x}_s | \mathbf{x}_t, \mathbf{y}) = \mathcal{N}(\mathbf{x}; \mathbf{m}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}), \Sigma_s^{\text{dps}})$ with

$$\begin{aligned} \mathbf{m}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}) &:= \Sigma_{s|t}^{\text{dps}}(\eta_s^{-2}\boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + \sigma_{\mathbf{y}}^{-2}D_s^\top P_{\bar{\mathbf{m}}}^\top \mathbf{y}), \\ \Sigma_{s|t}^{\text{dps}} &:= (\eta_s^{-2}\mathbf{I}_d + \sigma_{\mathbf{y}}^{-2}(P_{\bar{\mathbf{m}}} D_s)^\top P_{\bar{\mathbf{m}}} D_s)^{-1}. \end{aligned}$$

1134 Next, for the DING transition, first set $b_s(Z_s) := -(\sigma_s/\alpha_s)P_{\bar{\mathbf{m}}}\hat{\mathbf{x}}_1(Z_s, s)$. Gaussian conjugacy with
 1135 $p_{s|t}^\eta(\mathbf{x}_s|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_s; \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta), \eta_s^2 \mathbf{I}_d)$ gives
 1136

$$1137 \hat{\pi}_{s|t}^{\text{ding}}(\mathbf{x}_s|Z_s, \mathbf{x}_t, \mathbf{y}) = \mathcal{N}\left(\mathbf{x}_s; \tilde{\Sigma}_{s|t}^{\text{ding}}\left(\eta_s^{-2} \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + \sigma_y^{-2} \alpha_s^{-1} P_{\bar{\mathbf{m}}}^\top(\mathbf{y} - b_s(Z_s))\right), \tilde{\Sigma}_{s|t}^{\text{ding}}\right),$$

1138 and $\tilde{\Sigma}_{s|t}^{\text{ding}} := (\eta_s^{-2} \mathbf{I}_d + \alpha_s^{-2} \sigma_y^{-2} P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}})^{-1}$. Since the mean of this conditional distribution is clearly
 1139 affine in Z_s , we integrate it out, yielding that $\hat{\pi}_{s|t}^{\text{ding}}(\mathbf{x}_s|\mathbf{x}_t, \mathbf{y}) = \mathcal{N}(\mathbf{x}_s; \mathbf{m}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y}), \Sigma_{s|t}^{\text{ding}})$ where
 1140

$$1141 \mathbf{m}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y}) := \tilde{\Sigma}_{s|t}^{\text{ding}}\left(\eta_s^{-2} \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + \sigma_y^{-2} \alpha_s^{-1} P_{\bar{\mathbf{m}}}^\top \mathbf{y}\right. \\ 1142 \quad \left. + (\sigma_y^{-2} \alpha_s^{-2}) P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}} (\mathbf{I}_d - \alpha_s D_s) \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\right), \\ 1143 \Sigma_{s|t}^{\text{ding}} := \tilde{\Sigma}_{s|t}^{\text{ding}} + \frac{\eta_s^2}{\sigma_y^4 \alpha_s^4} \tilde{\Sigma}_{s|t}^{\text{ding}} P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}} (\mathbf{I}_d - \alpha_s D_s) (\tilde{\Sigma}_{s|t}^{\text{ding}} P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}} (\mathbf{I}_d - \alpha_s D_s))^\top. \\ 1144$$

1145

1146 **Small-noise regime.** We now study the behavior of both transitions when the DDIM kernel variance
 1147 η_s^2 tends to zero. For simplicity we define
 1148

$$1149 K_{\text{dps}} := \sigma_y^{-2} D_s^\top P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}} D_s, \quad K_{\text{ding}} := \alpha_s^{-2} \sigma_y^2 M$$

1150 and $R_s = \mathbf{I}_d - \alpha_s D_s$, $M = P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}}$. Then,
 1151

$$1152 \Sigma_{s|t}^{\text{dps}} = (\eta_s^{-2} \mathbf{I}_d + K_{\text{dps}})^{-1}, \\ 1153 \Sigma_{s|t}^{\text{ding}} = (\eta_s^{-2} \mathbf{I}_d + K_{\text{ding}})^{-1} + \frac{\eta_s^2}{\alpha_s^4 \sigma_y^4} (\eta_s^{-2} \mathbf{I}_d + K_{\text{ding}})^{-1} M R_s R_s^\top M (\eta_s^{-2} \mathbf{I}_d + K_{\text{ding}})^{-1}. \\ 1154$$

1155 We use throughout that for any fixed matrix K , we have that when $\eta_s^2 \|K\|_{\text{op}} < 1$,
 1156

$$1157 (\eta_s^{-2} \mathbf{I}_d + K)^{-1} = \eta_s^2 (\mathbf{I}_d - \eta_s^2 K) + R_2(\eta_s), \quad \|R_2(\eta_s)\| \leq \eta_s^6 \frac{\|K\|_{\text{op}}^2}{1 - \eta_s^2 \|K\|_{\text{op}}}. \quad (\text{A.4})$$

1158 This follows from the standard Neumann (geometric) series expansion. Applying (A.4) with $\eta_s^2 \leq \min(1/\|K_{\text{dps}}\|_{\text{op}}, 1/\|K_{\text{ding}}\|_{\text{op}})$, we get
 1159

$$1160 \Sigma_{s|t}^{\text{dps}} = \eta_s^2 (\mathbf{I}_d - \eta_s^2 K_{\text{dps}}) + \mathcal{O}(\eta_s^6), \\ 1161 \Sigma_{s|t}^{\text{ding}} = \eta_s^2 (\mathbf{I}_d - \eta_s^2 K_{\text{ding}}) + \frac{\eta_s^6}{\alpha_s^4 \sigma_y^4} M R_s R_s^\top M + \mathcal{O}(\eta_s^6).$$

1162 and thus
 1163

$$\Sigma_{s|t}^{\text{dps}} - \Sigma_{s|t}^{\text{ding}} = \mathcal{O}(\eta_s^4).$$

1164 Plugging these expansions in the mean terms, we find that
 1165

$$1166 \mathbf{m}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}) = \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + \eta_s^2 (\sigma_y^{-2} D_s^\top P_{\bar{\mathbf{m}}}^\top \mathbf{y} - K_{\text{dps}} \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)) + \mathcal{O}(\eta_s^4), \\ 1167 \mathbf{m}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y}) = \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + \eta_s^2 (\alpha_s^{-1} \sigma_y^{-2} P_{\bar{\mathbf{m}}}^\top \mathbf{y} + \alpha_s^{-2} \sigma_y^{-2} M R_s \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) - K_{\text{ding}} \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)) + \mathcal{O}(\eta_s^4).$$

1168 This yields
 1169

$$1170 \mathbf{m}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}) - \mathbf{m}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y}) = \eta_s^2 \sigma_y^{-2} \left[(\alpha_s^{-1} \mathbf{I}_d - D_s^\top) P_{\bar{\mathbf{m}}}^\top \mathbf{y} \right. \\ 1171 \quad \left. + \alpha_s^{-2} M R_s \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) - (\alpha_s^{-2} M - D_s^\top M D_s) \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) \right] + \mathcal{O}(\eta_s^4), \\ 1172$$

1173 We now proceed to further upper bound the leading term. Define $E_s := D_s - \alpha_s^{-1} \mathbf{I}_d$. Then
 1174 $R_s = -\alpha_s E_s$ and we have that
 1175

$$1176 \mathbf{m}_{s|t}^{\text{dps}}(\mathbf{x}_t, \mathbf{y}) - \mathbf{m}_{s|t}^{\text{ding}}(\mathbf{x}_t, \mathbf{y}) = \\ 1177 \eta_s^2 \sigma_y^{-2} \left(-E_s^\top P_{\bar{\mathbf{m}}}^\top \mathbf{y} + \alpha_s^{-1} E_s^\top M \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + E_s^\top M E_s \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) \right) + \mathcal{O}(\eta_s^4).$$

1178 with $M = P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}}$, which is an orthogonal projection matrix since $M^\top = M$ and $P_{\bar{\mathbf{m}}}^\top P_{\bar{\mathbf{m}}}^\top = \mathbf{I}_{d_y}$ and
 1179 thus $M^2 = M$. We proceed by bounding each term of
 1180

$$1181 -E_s^\top P_{\bar{\mathbf{m}}}^\top \mathbf{y} + \alpha_s^{-1} E_s^\top M \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta) + E_s^\top M E_s \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)$$

1188 separately. Define $\varepsilon_s := \|ME_s\|_{\text{op}}$. Then, since $\mathbf{v} := P_{\overline{\mathbf{m}}}^\top \mathbf{y} \in \text{range}(M)$, we have $M\mathbf{v} = \mathbf{v}$. Hence
 1189

$$1190 \quad E_s^\top P_{\overline{\mathbf{m}}}^\top \mathbf{y} = E_s^\top M P_{\overline{\mathbf{m}}}^\top \mathbf{y} = (ME_s)^\top (MP_{\overline{\mathbf{m}}}^\top \mathbf{y})$$

1191 where we have used that $M^\top M = M$. By the operator norm inequality, and the fact that $\|P_{\overline{\mathbf{m}}}^\top \mathbf{y}\| = \|y\|$, we get
 1192

$$1193 \quad \|E_s^\top P_{\overline{\mathbf{m}}}^\top \mathbf{y}\| \leq \|ME_s\|_{\text{op}} \|MP_{\overline{\mathbf{m}}}^\top \mathbf{y}\| = \varepsilon_s \|P_{\overline{\mathbf{m}}}^\top \mathbf{y}\| = \varepsilon_s \|\mathbf{y}\|.$$

1194 Next, using the same operator norm inequality we get that
 1195

$$1196 \quad \|E_s^\top M \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\| \leq \varepsilon_s \|M \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\|, \quad \|E_s^\top ME_s \boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\| \leq \varepsilon_s^2 \|\boldsymbol{\mu}_{s|t}(\mathbf{x}_t; \eta)\|.$$

1197 which yields the desired bound. \square
 1198

1199 **Proposition 2** (Upperbound on ε_s). *We have that*

$$1200 \quad \varepsilon_s \leq \frac{\sigma_s^2}{\alpha_s} \frac{1}{\alpha_s^2 \lambda_{\min}(\Sigma) + \sigma_s^2}$$

1201 where $\lambda_{\min}(\Sigma)$ is the smallest eigenvalue of Σ .
 1202

1203 *Proof.* By noting that $(\alpha_s^2 \Sigma + \sigma_s^2 \mathbf{I})E_s = -\alpha_s^{-1} \sigma_s^2 \mathbf{I}_d$, we get the alternative expression
 1204

$$1205 \quad E_s = -\frac{\sigma_s^2}{\alpha_s} (\alpha_s^2 \Sigma + \sigma_s^2 \mathbf{I}_d)^{-1}.$$

1206 By the submultiplicativity of the operator norm and the fact that M is a non-trivial orthogonal
 1207 projection matrix, we have that
 1208

$$1209 \quad \|E^\top M\|_{\text{op}} \leq \|E\|_{\text{op}} = \frac{\sigma_s^2}{\alpha_s} \frac{1}{\lambda_{\min}(\alpha_s^2 \Sigma + \sigma_s^2 \mathbf{I}_d)} \leq \frac{\sigma_s^2}{\alpha_s} \frac{1}{\alpha_s^2 \lambda_{\min}(\Sigma) + \sigma_s^2}.$$

1210 \square
 1211

1212 B DETAILS ABOUT THE EXPERIMENTS

1213 B.1 MODELS

1214 We use both the SD 3 and SD 3.5 (medium) (Esser et al., 2024) models with the linear schedule
 1215 $\alpha_t = 1 - t$ and $\sigma_t = t$. In all the experiments we run the zero-shot methods with a guidance scale
 1216 of 2. The fine-tuned baseline, which we refer to as SD3 Inpaint, is based on the publicly available
 1217 model² trained for inpainting with a ControlNet-augmented version of Stable Diffusion 3. It has
 1218 been finetuned on a large dataset of approximately 12 million 1024×1024 image–mask pairs to
 1219 directly predict high-quality inpainted completions conditioned on the masked image and the mask
 1220 itself. We have found the model to perform well also on lower resolutions, despite not undergoing
 1221 multi-resolution training. Examples of image editing of lower resolution images are presented in
 1222 the the HuggingFace page of the smae project. We run this baseline using a guidance scale of 7 for
 1223 optimal results.
 1224

1225 Finally, all experiments use `bfloat16` for model forward passes (and backward passes for baselines
 1226 that require it), with other computations performed in `float32`.
 1227

1228 B.2 MASK DOWNSAMPLING

1229 To construct the mask in the latent space, we start from the original binary mask defined in pixel
 1230 space. Since the encoder reduces spatial resolution by a fixed factor (here, 8), we downsample the
 1231 pixel-space mask to match the resolution of the latent representation. This is done by applying
 1232 bilinear interpolation with antialiasing. The resulting low-resolution mask captures the proportion of
 1233 masked pixels within each latent receptive field. Finally, we threshold this downsampled mask at
 1234 0.95 to obtain a binary latent mask, slightly overestimating the masked region to prevent boundary
 1235 artifacts during sampling.
 1236

1237 ²<https://huggingface.co/alimama-creative/SD3-Controlnet-Inpainting>

1242 B.3 IMPLEMENTATION OF THE BASELINES
1243

1244 Here, we give implementation details of the baselines. We stress that *each baseline is run in the latent*
1245 *space*, and thus no method computes the gradient w.r.t. the input of the decoder. We also manually
1246 tuned each baseline for the considered tasks. We provide the used hyperparameters in Table 8.
1247

1248 **BLENDED-DIFF.** We implemented Avrahami et al. (2023, Algorithm 1) following their official
1249 code³. The codebase includes an additional hyperparameter, `blending_percentage`, which
1250 determines at what fraction of the inference steps blending begins. We set it to zero, as applying
1251 blending across all steps produced the best results. A key detail is the original implementation is
1252 that the observed region (background) is re-noised to the noise level defined by the current timestep;
1253 see Avrahami et al. (2023, step 1-2 within the for loop in Algo 1), yet the reconstructed region
1254 (foreground) has less noise as it comes from applying a DDIM transition. This causes the background
1255 and foreground to follow different noise levels, and hence, introduces minor artifacts in the final
1256 reconstructions. We fixed this issue in our implementation by matching the two noise levels.
1257

1258 **DAPS.** We adapt Zhang et al. (2025, Algorithm 1) based on the released code⁴ to the flow matching
1259 formulation. We found that using Langevin as MCMC sampler for enforcing data consistency works
1260 the best for low NFE regime.
1261

1262 **DIFFPIR.** We make Zhu et al. (2023, Algorithm 1) compatible with the flow matching formulation
1263 with step 4 being implemented in the case of mask operator. We found in practice that the hyper-
1264 parameter λ has little impact on the quality of reconstructions and hence we use the recommended
1265 values $\lambda = 1^5$. On the other hand for the second hyperparameter ζ , we find that using $\zeta = 0.3$ yielded
1266 the best reconstructions.
1267

1268 **DDNM.** We adapt the implementation in the released code⁶ to the flow matching formulation with
1269 the step 4 in Wang et al. (2023b, Algorithm 3) being implemented for a mask operator. The official
1270 implementation uses a DDIM transition in step 5 of Algorithm 3 whose stochasticity is controlled by
the hyperparameters η . As recommended, we set the latter to $\eta = 0.85$.
1271

1272 **FLOWCHEF & FLOWDPS.** For both algorithms, we adapt the implementations available in the
1273 released codes FLOWCHEF⁷⁸ to our codebase. We observe that the two algorithms are quite similar,
1274 with FLOWDPS being distinct by adding stochasticity between iterations.
1275

1276 **PNP-FLOW.** We reimplement Martin et al. (2025, Algorithm 3) while taking as a reference the
1277 released code⁹. For the stepsizes on data fidelity term, we find that a constant scheduler with higher
1278 stepsize enables the algorithm to fit the observation, mitigate the smooth and blurring effects in the
1279 reconstruction and hence yield better reconstructions.
1280

1281 **PSLD.** We implement the PSLD algorithm provided in Rout et al. (2024b, Algorithm 2). We find
1282 that PSLD algorithm requires several diffusion steps, e.g. at least 150 diffusion steps, to yield good
1283 results. Unfortunately, we were not able to make it work well for the low NFE setup.
1284

1285 **REDDIFF.** We implement Mardani et al. (2024, Algorithm 1) based on the official code¹⁰ and adapt
1286 it to the flow matching formulation. We initialize the algorithm with a sample for a standard Gaussian.
1287 For low NFE setups, we find that using a constant weight schedule yields better results, namely in
1288 terms fitting the observation and providing consistent reconstructions.
1289

³<https://github.com/omriav/blended-latent-diffusion>

⁴<https://github.com/zhangbingliang2019/DAPS>

⁵<https://github.com/yuanzhi-zhu/DiffPIR>

⁶<https://github.com/wyhuai/DDNM>

⁷<https://github.com/FlowDPS-Inverse/FlowDPS>

⁸<https://github.com/FlowChef/flowchef>

⁹<https://github.com/annegnx/PnP-Flow>

¹⁰<https://github.com/NVlabs/RED-diff>

1296 RESAMPLE. We reimplemented [Song et al. \(2024, Algorithm 1\)](#) based on the provided implementation
 1297 details in [Song et al. \(2024, Appendix\)](#) and the reference code¹¹. As noted in [Janati et al. \(2025a\)](#),
 1298 we set the tolerance ε for optimizing the data consistency to the noise level σ_y . Since we are working
 1299 with low NEFs, we set the frequency at which hard data consistency is applied (skip step size) to 5.
 1300 That aside, we found that the algorithm requires several diffusion steps (200) in order to output good
 1301 enough reconstructions. We note that removing the DPS step in the data consistency steps reduces
 1302 the quality of the reconstructions.

1303

1304

1305 Table 8: Hyperparameters for each algorithm (using the same notations as in their paper) and task variations.
 1306 “—” indicates identical across tasks.

Algorithm	n_{steps}	Base hyperparameters	Latent tasks				
			Half	Top	Bottom	Center	Strip
BLEND-DIFF	50	<code>blending_percentage=0</code>	—	—	—	—	—
DAPS	50	$N_{\text{node}} = 2$ $\text{MCMC_steps} = 20$ $\beta_y = 10^{-2}$ $\text{Min_ratio} = 0.43$ $\text{MCMC_sampler} = \text{Langevin}$ $\rho = 1$	$\eta_0 = 2 \times 10^{-5}$	$\eta_0 = 3 \times 10^{-5}$	$\eta_0 = 2 \times 10^{-5}$	$\eta_0 = 9 \times 10^{-6}$	$\eta_0 = 2 \times 10^{-5}$
DDNM	50	$\eta = 0.85$	—	—	—	—	—
DIFFPIR	50	$\lambda = 1$ $\zeta = 0.3$	—	—	—	—	—
FLOWCHEF	50	$\text{step_size} = 0.9$ $\text{grad_descent_steps} = 10$	—	—	—	—	—
FLOWDPS	50	$\text{grad_descent_steps} = 3$	$\text{step_size} = 20$	$\text{step_size} = 10$	$\text{step_size} = 10$	$\text{step_size} = 10$	$\text{step_size} = 10$
PNP-FLOW	50	$\alpha = 1.0$ $\text{lr_style} = \text{constant}$	$\gamma_n = 0.8$	$\gamma_n = 1.3$	$\gamma_n = 1.4$	$\gamma_n = 0.8$	$\gamma_n = 0.8$
PSLD	50	$\text{DDIM_param} = 1.0$	$\gamma = 0.01$ $\eta = 0.01$	$\gamma = 0.01$ $\eta = 0.01$	$\gamma = 0.01$ $\eta = 0.01$	$\gamma = 0.05$ $\eta = 0.1$	$\gamma = 0.1$ $\eta = 0.5$
REDDIFF	50	$\text{lr} = 0.2$ $\text{grad_term_weight} = 0.25$ $\text{obs_weight} = 1.0$	—	—	—	—	—
RESAMPLE	50	$C = 5$ $\text{grad_descent_steps} = 200$ $\gamma_{\text{scale}} = 40.0$ $\text{lr}_{\text{pixel}} = 10^{-2}$ $\text{lr}_{\text{latent}} = 5 \times 10^{-3}$	—	—	—	—	—
DING (ours)	25	$\eta = \sigma_s(1 - \alpha_s)$	—	—	—	—	—

1325

1326

1327 C EXAMPLES OF RECONSTRUCTIONS

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

¹¹<https://github.com/soominkwon/resample>

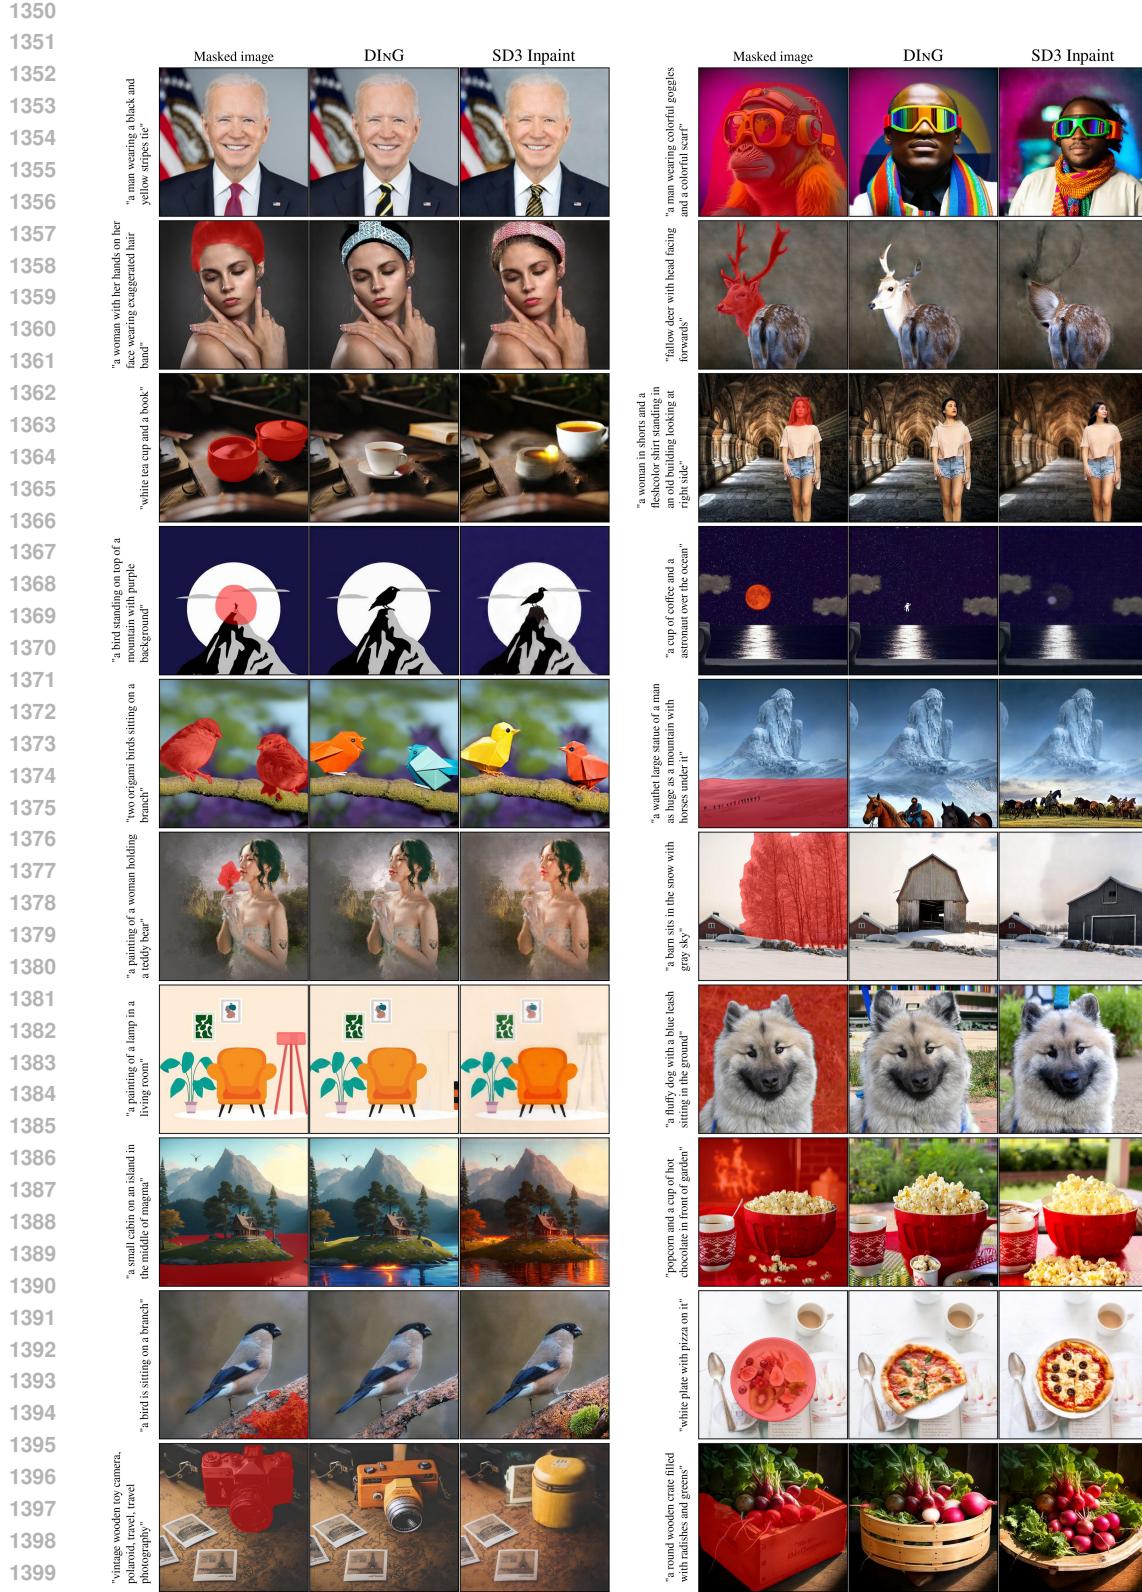


Figure 7: Comparison of DING and finetuned SD3 on PIE-Bench. Both methods have the same runtime of 2.2s.

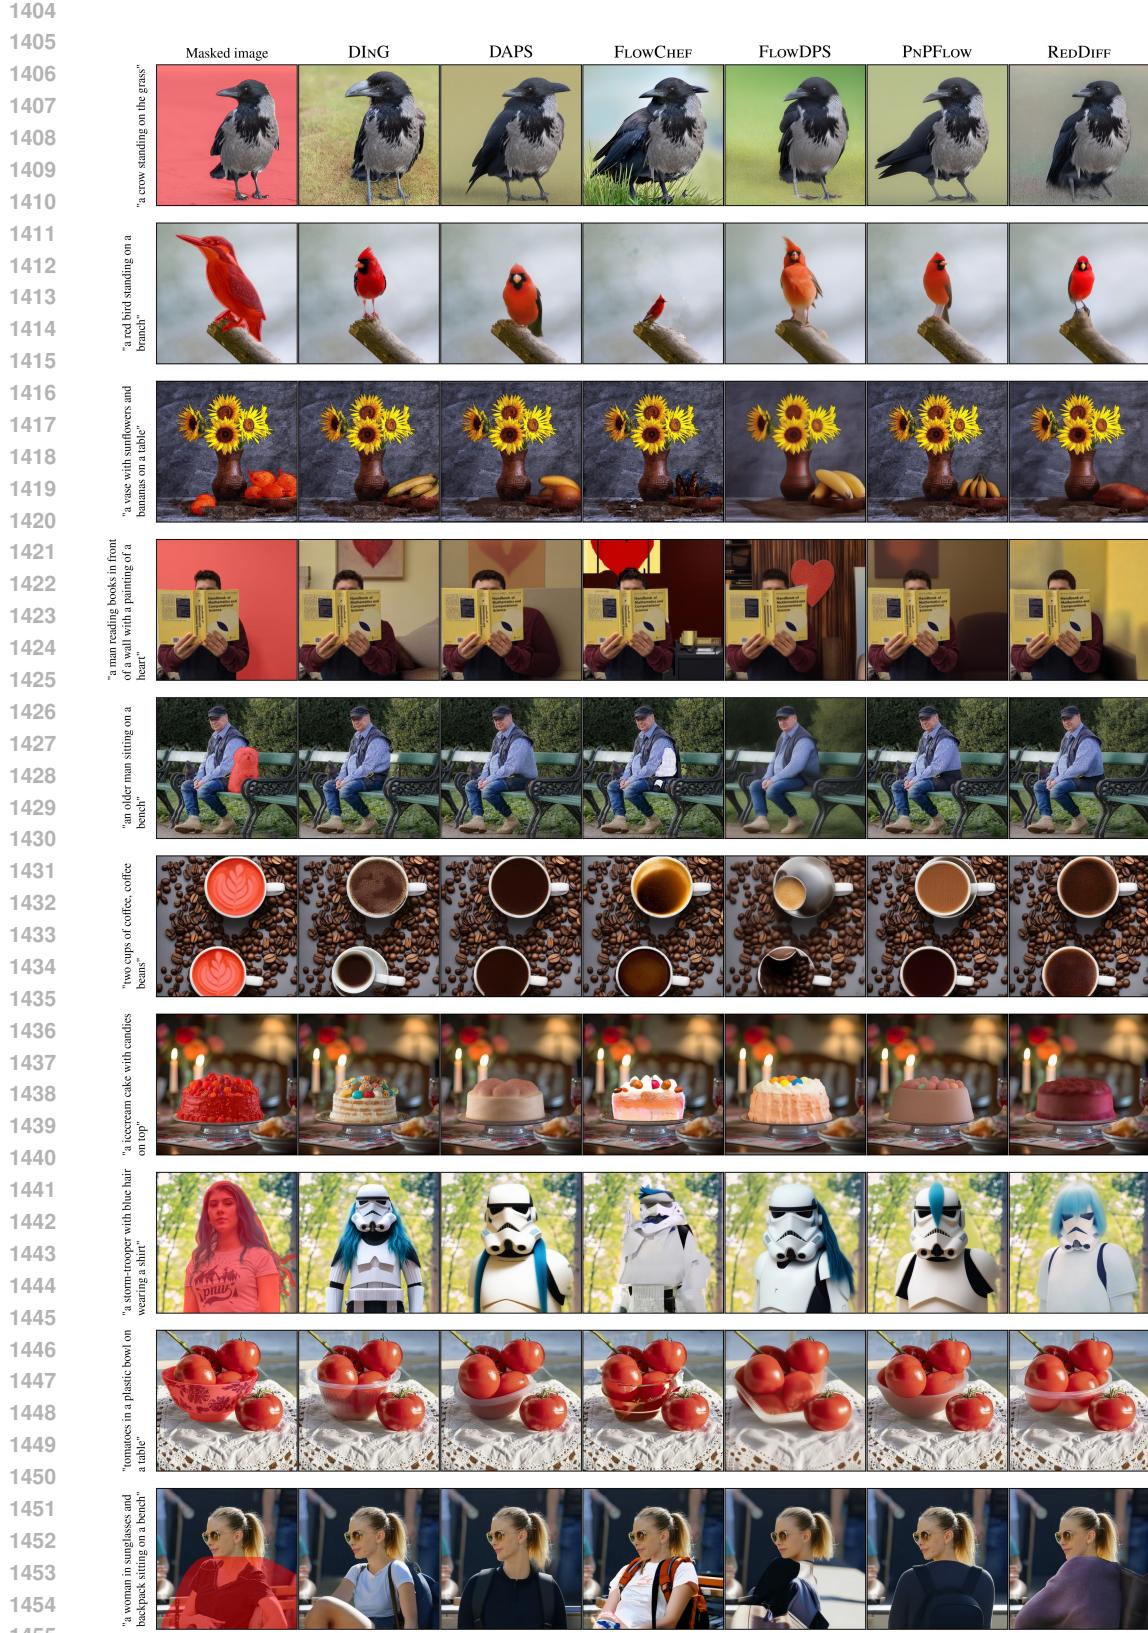


Figure 8: Comparison of DING and zero-shot baselines on PIE-Bench. All methods use 50 NFEs.

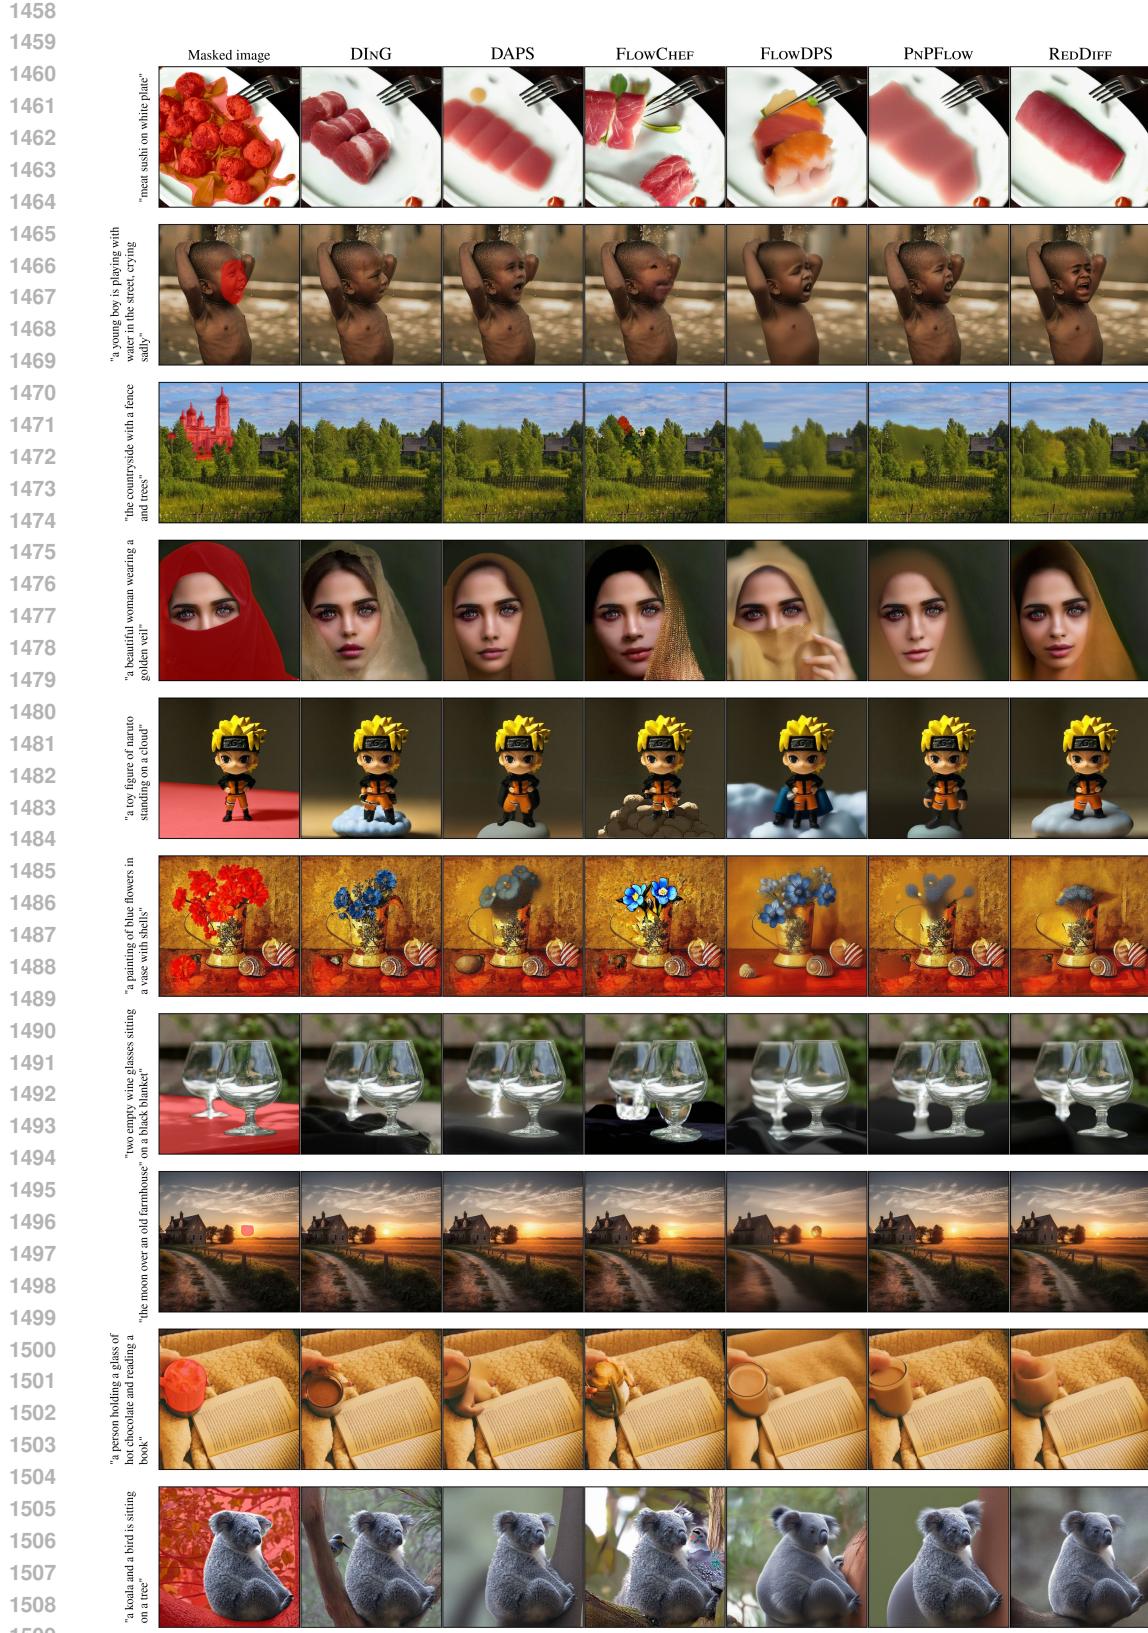


Figure 9: Comparison of DING and zero-shot baselines on PIE-Bench. All methods use 50 NFEs.

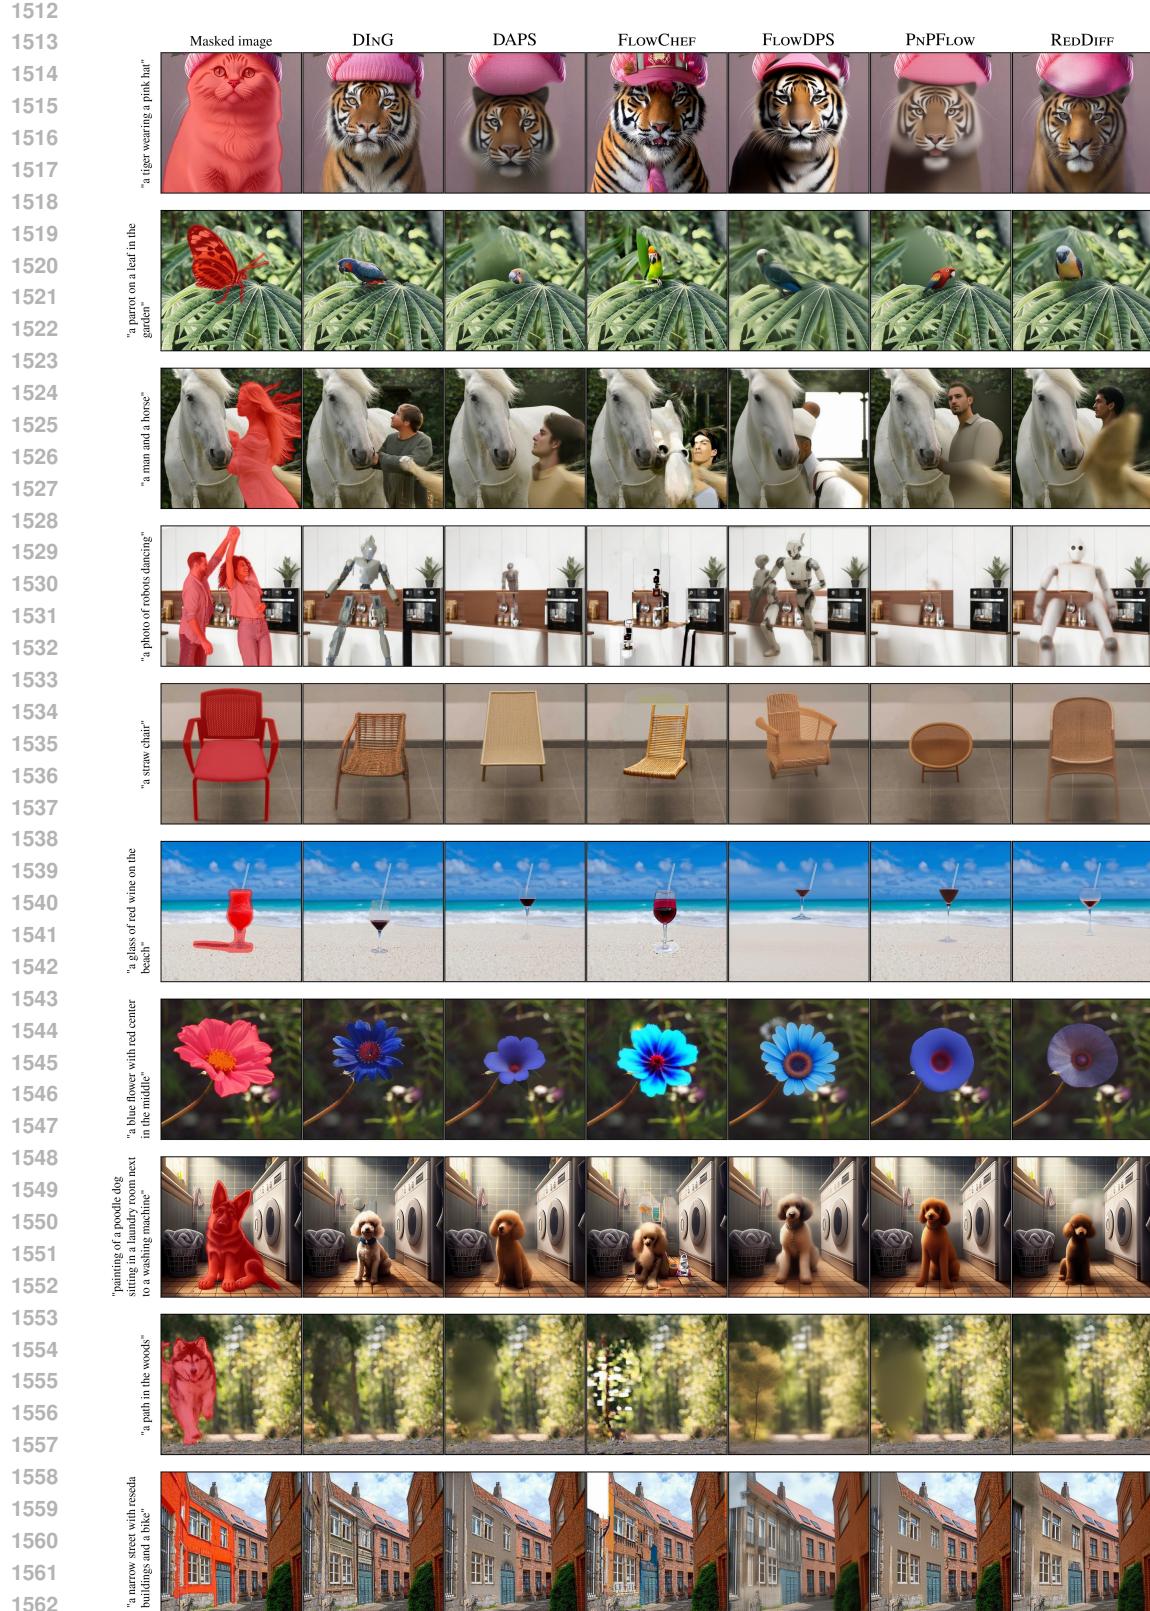


Figure 10: Comparison of DING and zero-shot baselines on PIE-Bench. All methods use 50 NFES.

