A^{2}PS: SOLVING NOISY INVERSE PROBLEMS VIA ASYMPTOTICALLY ACCURATE POSTERIOR SAMPLING IN DIFFUSION AND FLOW MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Training-free posterior sampling enables pretrained diffusion and flow models to solve inverse problems without task-specific retraining, but approximation errors can accumulate along the generative trajectory and bias the final distribution. We propose A^{2}PS, a posterior sampling framework that progressively corrects these errors. Our key insight is that, near the end of generation, the distribution of the final clean sample conditioned on an intermediate state becomes increasingly well approximated by a Gaussian distribution, yielding an asymptotically accurate posterior score estimate. A^{2}PS uses this estimate with Langevin dynamics to correct errors inherited from earlier steps. We establish the asymptotic accuracy of the posterior score and validate the resulting error-correction mechanism in a controlled setting. Across six noisy linear and nonlinear image inverse problems on ImageNet and FFHQ, with both pixel-space and latent-space models, A^{2}PS achieves the best or near-best restoration quality among evaluated methods while also being the fastest across all tested tasks.

1 INTRODUCTION



Figure 1: Training-free image restoration with A^{2}PS (left side). A single pretrained generative model is used to solve various inverse problems, including deblurring, inpainting, and super-resolution.

Many real-world imaging tasks can be formulated as *noisy inverse problems* (Chalmond, 2003). Given a noisy measurement system, the observation is modeled as

$$\boldsymbol{y} = A(\boldsymbol{x}) + \boldsymbol{\delta}, \qquad \boldsymbol{\delta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_y), \tag{1}$$

where $\boldsymbol{y}$ is the measured result, $\boldsymbol{x}$ is the unknown signal of interest, $A(\cdot)$ denotes the measurement operator, and $\boldsymbol{\Sigma}_y$ characterizes the measurement noise. The goal is to recover $\boldsymbol{x}$ given $\boldsymbol{y}$, $A(\cdot)$, and $\boldsymbol{\Sigma}_y$. Such problems are common in physical measurements and image restoration. For example, different choices of $A(\cdot)$ naturally describe restoration tasks such as inpainting, deblurring, and super-resolution.

A straightforward approach is to train a task-specific restoration model for each operator (Zamir et al., 2022; Potlapalli et al., 2023). However, this requires additional data and optimization for every new task, and the resulting model may no longer be applicable when the operator or noise condition changes. Since these restoration tasks share the same inverse-problem formulation, it is

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

desirable to solve them by simply specifying the measurement process rather than retraining a new model for each setting.

Training-free methods provide such solutions using Bayes' rule, $p(\boldsymbol{x} \mid \boldsymbol{y}) \propto p(\boldsymbol{x})\, p(\boldsymbol{y} \mid \boldsymbol{x})$, where a pretrained generative model provides the image prior $p(\boldsymbol{x})$ while the measurement model specifies the likeliihood $p(\boldsymbol{y} \mid \boldsymbol{x})$. The key idea is then to steer samples generated from the pretrained prior toward regions that are also consistent with the observed measurement, thereby approximating samples from the posterior $p(\boldsymbol{x} \mid \boldsymbol{y})$ without retraining.

Diffusion models (DMs) (Ho et al., 2020; Song et al., 2021) and flow-matching models (FMs) (Lipman et al., 2023; Liu et al., 2023) provide powerful pretrained priors, which are sampled along a trajectory $p_{tt\in[0,1]}$ from a simple distribution $p_0 = \mathcal{N}(\boldsymbol{\mu}0, \sigma_0^2\mathbf{I})$ to the data distribution $p_1 \approx p_{\text{data}}$. To solve an inverse problem, training-free methods incorporate the observation $\boldsymbol{y}$ into this trajectory to steer generation toward the posterior. However, generation only provides intermediate noisy states rather than the clean image $\boldsymbol{x}$, so the likelihood or posterior guidance at these states is generally unavailable in closed form. Existing training-free methods therefore use different approximations to incorporate measurement information along the generative trajectory (Chung et al., 2023; Kim et al., 2025; Patel et al., 2025; Wang et al., 2023; Kawar et al., 2022; Cardoso et al., 2024).

However, this introduces a fundamental difficulty: approximation errors at intermediate steps can propagate and accumulate throughout sampling, gradually driving the generated distribution away from the desired posterior. Recent approaches have recognized this issue and introduced mechanisms to mitigate error propagation (Zhang et al., 2025; Rout et al., 2023). However, these mechanisms do not generally eliminate errors inherited from earlier sampling steps, leaving a residual mismatch with the target posterior. Figure 2 illustrates this limitation through two representative approaches: DPS (Chung et al., 2023), which estimates the posterior score at each step without an explicit error correction mechanism, and DAPS (Zhang et al., 2025), which introduces an error mitigation through decoupled annealing. Both of these methods exhibit noticeable terminal mismatches with the exact posterior.

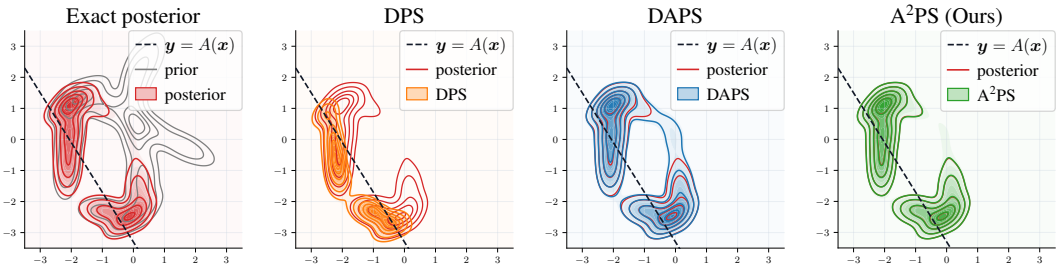


Figure 2: **Terminal distributions for a 2D inverse problem.** All panels use the same coordinate system and measurement $\boldsymbol{y} = A(\boldsymbol{x})$. The leftmost panel shows the exact posterior distribution, serving as the reference. The remaining panels show the terminal distributions produced by DPS, DAPS, and A^{2}PS, respectively. Comparing each panel with the reference (red contour) reveals the residual distributional mismatch at the end of the generative trajectory.

We therefore propose training-free *Asymptotically Accurate Posterior Sampling (A^{2}PS)*, a framework that progressively corrects errors accumulated along the generative trajectory. Our key observation is that the conditional distribution $p(\boldsymbol{x}_1 \mid \boldsymbol{x}_t)$ admits a Gaussian approximation that becomes increasingly accurate as $t \to 1$. This provides an increasingly accurate approximation to the posterior near the generation endpoint.

A^{2}PS leverages this property by using Langevin dynamics to repeatedly re-equilibrate the current samples toward the estimated posterior. Unlike methods that only modify the generative direction, this correction actively reduces errors inherited from earlier sampling steps. As the posterior approximation converges toward the exact posterior near the endpoint, the Langevin correction becomes increasingly accurate, leading to asymptotically accurate posterior sampling. As shown in Figure 2, A^{2}PS substantially reduces the residual mismatch and closely recovers the target posterior.

We evaluate A^{2}PS with both diffusion and flow-matching backbones, as well as pixel-space and latent-space generative models, demonstrating its general applicability. Across diverse noisy linear

and nonlinear inverse problems on FFHQ (Karras et al., 2019) and ImageNet (Deng et al., 2009) datasets, A^{2}PS achieves the best or near-best FID and LPIPS among the evaluated methods while providing up to nearly a $3\times$ runtime speedup over existing approaches.

Our contributions are summarized as follows:

- We propose A^{2}PS, a training-free posterior sampling framework that uses Langevin dynamics to eliminate errors accumulated along the generative trajectory.
- We establish an asymptotically accurate characterization of the posterior near the generative endpoint through a Gaussian approximation of $p(\boldsymbol{x}_1 \mid \boldsymbol{x}_t)$, providing a principled basis for eliminating accumulated sampling errors.
- We demonstrate the generality of A^{2}PS across diffusion and flow-matching backbones, pixel- and latent-space generation, and diverse noisy inverse problems, achieving competitive restoration quality.

2 PRELIMINARIES

Our method builds on two ingredients: a unified formulation of diffusion and flow-matching generative dynamics, and the re-equilibration property of Langevin dynamics. We introduce the notation and properties needed for our derivation of method below.

Diffusion and flow-matching generative models DMs and FMs generate data by gradually evolving a simple distribution $\mathbf{x}_0 \sim p_0 = \mathcal{N}(\boldsymbol{\mu}_0, \sigma_0^2\mathbf{I})$ along a predefined density path $\{p_t\}_{t\in[0,1]}$ toward the complex data distribution $\mathbf{x}_1 \sim p_1 \approx p_{\mathrm{data}}$. Albergo et al. (2025) showed that DMs and FMs should be described under a unified framework. For a Gaussian path, an intermediate sample can be written as $\mathbf{x}_t = a_t\mathbf{x}_0 + b_t\mathbf{x}_1$, where a_t, b_t are the path schedulers, with a_t decreasing from 1 to 0 and b_t increasing from 0 to 1; for example, Rectified Flow (Liu et al., 2023) uses $a_t = 1 - t$ and $b_t = t$. Thus, $\mathbf{x}_t$ gradually transitions from a noisy sample $\mathbf{x}_0$ to a target data sample $\mathbf{x}_1$ as $t \to 1$.

The evolution of p_t is associated with a velocity field $\boldsymbol{v}_t(\boldsymbol{x}) \coloneqq \mathbb{E}[\dot{\mathbf{x}}_t \mid \mathbf{x}_t = \boldsymbol{x}]$, which describes how samples are transported along the density path, and a score field $\boldsymbol{s}_t(\boldsymbol{x}) \coloneqq \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})$, which characterizes the local geometry of the intermediate distribution. To generate samples, particles are initialized from $\mathbf{x}_0 \sim p_0$ and evolved according to the differential equation,

$$\mathrm{d}\mathbf{x}_t = [\boldsymbol{v}_t(\mathbf{x}_t) + \epsilon_t \boldsymbol{s}_t(\mathbf{x}_t)]\,\mathrm{d}t + \sqrt{2\epsilon_t}\,\mathrm{d}\mathbf{w}_t, \quad \text{where } \epsilon_t \geq 0 \text{ controls the stochasticity.} \tag{2}$$

Regardless of ϵ_t, the particle marginals follow the same predefined path p_t and reach p_{data} at $t = 1$. In particular, $\epsilon_t = 0$ yields the deterministic ODE used by FMs, while $\epsilon_t > 0$ yields a family of stochastic SDE samplers used by DMs.

An important property is that the conditional mean of the generative endpoint can be obtained from the velocity field or equivalently from the score field through Tweedie's formula (Robbins, 1992; Efron, 2011),

$$\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) \coloneqq \mathbb{E}[\mathbf{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}] = \frac{a_t\boldsymbol{v}_t(\boldsymbol{x}) - \dot{a}_t\boldsymbol{x}}{a_t\dot{b}_t - \dot{a}_t b_t} = \frac{\boldsymbol{x} - a_t\boldsymbol{\mu}_0 + \sigma_0^2 a_t^2 \boldsymbol{s}_t(\boldsymbol{x})}{b_t}. \tag{3}$$

It implies a bijection between $\boldsymbol{v}_t$ and $\boldsymbol{s}_t$, so either field is sufficient to instantiate Equation (2). In practice, FMs typically predict $\boldsymbol{v}_t$, while DMs predict $\boldsymbol{s}_t$ or an equivalent parameterization.

Langevin dynamics. Langevin dynamics (LD) (Welling & Teh, 2011) describes a stochastic process that relaxes toward a target distribution $f(\boldsymbol{x})$ using only its score $\nabla_{\boldsymbol{x}} \log f(\boldsymbol{x})$. The overdamped LD is defined by

$$\mathrm{d}\mathbf{x}_\tau = \nabla_{\boldsymbol{x}} \log f(\mathbf{x}_\tau)\,\mathrm{d}\tau + \sqrt{2}\,\mathrm{d}\mathbf{w}_\tau, \tag{4}$$

where τ denotes the Langevin evolution time and $\mathbf{w}_\tau$ is a standard Brownian motion.

A key property of Equation (4) is that f is its equilibrium distribution. As τ increases, the distribution of $\mathbf{x}_\tau$ converges toward f, while the influence of its initial distribution vanishes. Consequently, LD progressively forgets the initial distribution and re-equilibrates samples toward the specified target distribution. This initialization-forgetting property of LD is central to A^{2}PS, as it provides a natural mechanism for correcting errors inherited from earlier generative steps by re-equilibrating the current distribution toward a specified target.

3 TRAINING-FREE ASYMPTOTICALLY ACCURATE POSTERIOR SAMPLING

Given a measurement $\boldsymbol{y}$, training-free posterior sampling aims to draw samples from the endpoint posterior $p(\boldsymbol{x}_1 \mid \mathbf{y} = \boldsymbol{y}) \propto p(\boldsymbol{y} \mid \mathbf{x}_1 = \boldsymbol{x}_1)\, p_1(\boldsymbol{x}_1)$, where the endpoint distribution p_1 represented by a pretrained generative model serving as the prior. Rather than training an additional posterior model, we seek to recover the posterior generative dynamics directly from the pretrained model and the known measurement process.

3.1 ORACLE POSTERIOR GENERATIVE PATH

Suppose we have an oracle generative path whose endpoint follows the desired posterior distribution $\mathbf{x}_{1|y} \sim p(\boldsymbol{x}_1 \mid \mathbf{y} = \boldsymbol{y})$. Following the same interpolation schedule as the pretrained generative model, we define $\mathbf{x}_{t|y} = a_t \mathbf{x}_0 + b_t \mathbf{x}_{1|y}$. The corresponding density, score field, and velocity field are

$$p_{t|y}(\boldsymbol{x}) \coloneqq p(\mathbf{x}_{t|y} = \boldsymbol{x}), \qquad \boldsymbol{s}_{t|y}(\boldsymbol{x}) \coloneqq \nabla_{\boldsymbol{x}} \log p_{t|y}(\boldsymbol{x}), \qquad \boldsymbol{v}_{t|y}(\boldsymbol{x}) \coloneqq \mathbb{E}[\dot{\mathbf{x}}_{t|y} \mid \mathbf{x}_{t|y} = \boldsymbol{x}]. \quad (5)$$

If the oracle posterior score $\boldsymbol{s}_{t|y}$ were available, sampling along this path would yield the desired posterior distribution at the generative endpoint. However, this oracle score is generally inaccessible.

Our goal is therefore to directly estimate $\boldsymbol{s}_{t|y}$ rather than training a conditional generative model. In the following subsection, we derive an estimator of $\boldsymbol{s}_{t|y}$ whose accuracy improves asymptotically with t, by establishing a chain of asymptotically accurate approximations.

3.2 ASYMPTOTICALLY ACCURATE POSTERIOR SCORE

(Proofs of Theorem 3.1 and Proposition 3.2 are provided in Appendix D.)

Rather than estimating $\boldsymbol{s}_{t|y}$ directly, we reduce posterior-score estimation to the conditional generative endpoint. By Equation (3), the oracle posterior score can be written as

$$\boldsymbol{s}_{t|y}(\boldsymbol{x}) = \frac{b_t \boldsymbol{\mu}_{1|t,y}(\boldsymbol{x}) - \boldsymbol{x} + a_t \boldsymbol{\mu}_0}{\sigma_0^2 a_t^2}, \qquad \boldsymbol{\mu}_{1|t,y}(\boldsymbol{x}) \coloneqq \mathbb{E}[\mathbf{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}]. \quad (6)$$

This makes estimating the conditional posterior endpoint $\boldsymbol{\mu}_{1|t,y}$ sufficient to recover the oracle posterior score.

By Bayes' rule, $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}) \propto p(\boldsymbol{y} \mid \mathbf{x}_1 = \boldsymbol{x}_1)\, p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$, where $p(\boldsymbol{y} \mid \mathbf{x}_1 = \boldsymbol{x}_1) = \mathcal{N}(\boldsymbol{y}; A(\boldsymbol{x}_1), \boldsymbol{\Sigma}_y)$ under Equation (1). While the likelihood is known, the endpoint conditional $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$ is generally intractable. Our key observation is that this conditional becomes increasingly simple near the generative endpoint.

For convenience, define the effective noise scale $\rho_t \coloneqq \sigma_0 a_t / b_t$, and $\rho_t \to 0$ as $t \to 1$.

Theorem 3.1. *The conditional probability achieves asymptotic Gaussianity with t:*

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}_1; \boldsymbol{\mu}_{1|t}(\boldsymbol{x}), \boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}))\left[1 + \mathcal{O}(\rho_t^3)\right], \textit{where}$$
$$\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) = \frac{\sigma_0^2 a_t^2}{b_t} \frac{\partial \boldsymbol{\mu}_{1|t}(\boldsymbol{x})}{\partial \boldsymbol{x}} = \rho_t^2 \mathbf{I} + \mathcal{O}(\rho_t^4). \quad (7)$$

Here, $\boldsymbol{\mu}_{1|t}$ is the exact conditional mean given by Equation (3)*, and $\boldsymbol{\Sigma}_{1|t}$ gives the exact conditional covariance.*

Hence, near the generative endpoint as $t \to 1$, the intractable conditional $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$ becomes an increasingly concentrated Gaussian centered at $\boldsymbol{\mu}_{1|t}$, with characteristic width of order ρ_t. This motivates the likelihood-conditioned Gaussian approximation

$$\hat{p}(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}) \propto p(\boldsymbol{y} \mid \boldsymbol{x}_1)\, \mathcal{N}(\boldsymbol{x}_1; \boldsymbol{\mu}_{1|t}(\boldsymbol{x}), \boldsymbol{\Sigma}_{1|t}(\boldsymbol{x})). \quad (8)$$

Since $\boldsymbol{\Sigma}_{1|t} = \mathcal{O}(\rho_t^2)$, this posterior also becomes increasingly concentrated, allowing its conditional mean to be recovered asymptotically through its mode.

Proposition 3.2. *Define $\tilde{\boldsymbol{\mu}}_{1|t,y}$ to be the mode of $\hat{p}(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y})$:*

$$\tilde{\boldsymbol{\mu}}_{1|t,y}(\boldsymbol{x}) \coloneqq \arg\min_{\tilde{\boldsymbol{x}}} \left\{ \frac{1}{2}\|A(\tilde{\boldsymbol{x}}) - \boldsymbol{y}\|^2_{\boldsymbol{\Sigma}_y^{-1}} + \frac{1}{2}\|\tilde{\boldsymbol{x}} - \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\|^2_{\boldsymbol{\Sigma}_{1|t}^{-1}(\boldsymbol{x})} \right\}, \tag{9}$$

$\tilde{\boldsymbol{\mu}}_{1|t,y}$ reach endpoint posterior mean $\boldsymbol{\mu}_{1|t,y}$ asymptotically with t, $\tilde{\boldsymbol{\mu}}_{1|t,y} - \boldsymbol{\mu}_{1|t,y} = \mathcal{O}(\rho_t^4)$.

For linear measurement operator A, the approximate posterior in Equation (8) is Gaussian, so Equation (9) recovers its conditional mean exactly. For nonlinear A, the same objective gives its mode, which becomes asymptotically equivalent to the desired conditional mean as the posterior concentrates. Thus, Equation (9) provides a unified estimator for both linear and nonlinear measurement operators, with a closed-form solution in the linear case and gradient-based optimization for differentiable nonlinear operators.

Corollary 3.3. *Under Theorem 3.1 and proposition 3.2, substituting $\tilde{\boldsymbol{\mu}}_{1|t,y}$ into Equation* (6) *gives the asymptotically accurate posterior-score estimator*

$$\tilde{\boldsymbol{s}}_{t|y}(\boldsymbol{x}) \coloneqq \frac{b_t \tilde{\boldsymbol{\mu}}_{1|t,y}(\boldsymbol{x}) - \boldsymbol{x} + a_t \boldsymbol{\mu}_0}{\sigma_0^2 a_t^2}. \tag{10}$$

Since $\tilde{\boldsymbol{s}}_{t|y} - \boldsymbol{s}_{t|y} = \frac{\tilde{\boldsymbol{\mu}}_{1|t} - \boldsymbol{\mu}_{1|t,y}}{b_t \rho_t^2}$, together with Proposition 3.2 we have $\tilde{\boldsymbol{s}}_{t|y} - \boldsymbol{s}_{t|y} = \mathcal{O}(\rho_t^2)$.

Finally, the above results establish the asymptotic chain,

$$\left.\begin{array}{c} p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}_1; \boldsymbol{\mu}_{1|t}, \boldsymbol{\Sigma}_{1|t}) \left[1 + \mathcal{O}(\rho_t^3)\right] \\ \tilde{\boldsymbol{\mu}}_{1|t,y} = \boldsymbol{\mu}_{1|t,y} + \mathcal{O}(\rho_t^4) \end{array}\right\} \implies \tilde{\boldsymbol{s}}_{t|y} = \boldsymbol{s}_{t|y} + \mathcal{O}(\rho_t^2), \tag{11}$$

and $\tilde{\boldsymbol{s}}_{t|y} \to \boldsymbol{s}_{t|y}$ as $t \to 1$. Therefore, although the posterior score is generally inaccessible at arbitrary generative times, it becomes asymptotically recoverable near the endpoint using only the pretrained generative model and the known measurement operator.

3.3 Posterior Re-equilibration and Practical Approach

Posterior re-equilibration. Although Corollary 3.3 shows that $\tilde{\boldsymbol{s}}_{t|y}$ becomes increasingly accurate near the generative endpoint, errors made at earlier times may have already moved the sample away from the desired posterior path. Simply keeping the generative trajectory therefore does not remove these accumulated errors.

To address this, we apply Langevin dynamics at each generative time using $\tilde{\boldsymbol{s}}_{t|y}$ as the target score. As discussed in Equation (4), LD progressively forgets its initialization and re-equilibrates toward the distribution specified by its score. Therefore, errors inherited from earlier generative steps can be corrected rather than only propagated forward. Together with $\tilde{\boldsymbol{s}}_{t|y} \to \boldsymbol{s}_{t|y}$ as $t \to 1$, this turns the asymptotically accurate score estimate into asymptotically accurate posterior sampling.

Algorithm 1 A^{2}PS with a flow-matching backbone

Input: measurement $\boldsymbol{y}$, operator A, pretrained FM $\boldsymbol{v}_\theta \approx \boldsymbol{v}_t$, generative schedule $\{t_n\}_{n=0}^N$, Langevin schedule $\{\tau_k\}_{k=0}^K$, guidance strength λ

Sample $\boldsymbol{x} \sim p_0$
for $n = 0, \dots, N-2$ **do**
  $\boldsymbol{x} \leftarrow \texttt{FMStep}(\boldsymbol{x}, \boldsymbol{v}_t, t_n)$ ▷ eq. (2)
  for $k = 0, \dots, K-1$ **do** ▷ Langevin dynamics
    $\boldsymbol{\mu}_{1|t_n}(\boldsymbol{x}),\ \boldsymbol{s}_{t_n}(\boldsymbol{x}) \leftarrow$ eq. (3)
    $\tilde{\boldsymbol{\mu}}_{1|t_n,y}(\boldsymbol{x}) \leftarrow$ eq. (9)
    $\tilde{\boldsymbol{s}}_{t_n|y}(\boldsymbol{x}) \leftarrow$ eq. (10)
    $\tilde{\boldsymbol{s}}^{(\lambda)}_{t_n|y} \leftarrow \tilde{\boldsymbol{s}}_{t_n|y} + \lambda(\tilde{\boldsymbol{s}}_{t_n|y} - \boldsymbol{s}_{t_n})$
    $\boldsymbol{x} \leftarrow \texttt{LDStep}(\boldsymbol{x}, \boldsymbol{s}^{(\lambda)}_{t_n|y}, \tau_k)$ ▷ eq. (4)
  end for
end for
$\boldsymbol{x} \leftarrow \texttt{FMStep}(\boldsymbol{x}, \boldsymbol{v}_t, t_{N-1})$
return $\boldsymbol{x}$

Additional likelihood guidance. The posterior score naturally balances the generative prior and the measurement likelihood, $\boldsymbol{s}_{t|y}(\boldsymbol{x}) = \boldsymbol{s}_t(\boldsymbol{x}) + \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{y} \mid \mathbf{x}_t = \boldsymbol{x})$. This allows us to directly control the strength of the measurement constraint by reweighting the likelihood,

$$p_t(\boldsymbol{x})\, p_t(\boldsymbol{y} \mid \mathbf{x}_t = \boldsymbol{x})^{1+\lambda} \iff \boldsymbol{s}_{t|y}(\boldsymbol{x}) + \lambda(\boldsymbol{s}_{t|y}(\boldsymbol{x}) - \boldsymbol{s}_t(\boldsymbol{x})) =: \boldsymbol{s}^{(\lambda)}_{t|y}(\boldsymbol{x}). \tag{12}$$

The parameter λ therefore changes only the likelihood contribution: $\lambda = 0$ recovers the original posterior, while positive and negative values respectively strengthen and weaken the measurement constraint. We use the likelihood-adjusted posterior score $\boldsymbol{s}_{t|y}^{(\lambda)}$ as the target for LD re-equilibration. In practice, we replace $\boldsymbol{s}_{t|y}$ by its estimate $\tilde{\boldsymbol{s}}_{t|y}$ and denote the resulting score by $\tilde{\boldsymbol{s}}_{t|y}^{(\lambda)}$.

Practical approach. To make the procedure concrete, we instantiate A^{2}PS with FM backbone as an example summarized in Algorithm 1. When computing the full conditional covariance $\boldsymbol{\Sigma}_{1|t}$ is prohibitively expensive in high dimensions, we use the isotropic approximation $\boldsymbol{\Sigma}_{1|t} \approx \hat{\sigma}_t^2 \boldsymbol{I}$, where $\hat{\sigma}_t$ is estimated once for the pretrained model through an offline statistic and reused during sampling (Appendix B). For latent generative models, the decoder $D(\cdot)$ can be absorbed into the measurement operator, i.e., $A \leftarrow A \circ D$, allowing the same formulation to be applied directly in latent space.

Remark 3.4 (Special case.). For zero-noise inpainting, A is a masking operator whose null space is the masked region. Under the isotropic covariance approximation, Equation (9) becomes independent of the conditional covariance, yielding an algorithm mathematically equivalent to the specialized zero-noise inpainting approach proposed by Zheng et al. (2025).

4 Additional Related Work

Training-free posterior sampling. Training-free posterior samplers mainly differ in how they incorporate measurements during generation. Likelihood-guided methods (Yu et al., 2023; Rout et al., 2023; 2024), such as DPS (Chung et al., 2023), approximate the intermediate likelihood using a predicted endpoint. FlowDPS (Kim et al., 2025) extends this idea to FMs, whereas FlowChef (Patel et al., 2025) simplifies the propagation of measurement gradients through the flow dynamics. Another line exploits the structure of linear measurement operators, such as spectral conditioning in DDRM (Kawar et al., 2022) and range-null-space projection in DDNM (Wang et al., 2023). These approaches do not explicitly remove errors inherited from earlier steps. DAPS (Zhang et al., 2025) instead reduces error propagation by decoupling sampling trajectory, but its practical sampler relies on a heuristic Gaussian approximation of the endpoint conditional. In contrast, A^{2}PS derives an asymptotically accurate endpoint conditional and uses posterior re-equilibration to correct accumulated errors. A detailed comparison is provided in Appendix E.

Accelerated Langevin algorithms. The slow convergence of standard Langevin algorithm has motivated various accelerated algorithms based on underdamped dynamics, gradient-adjusted updates, or higher-order methods (Zhang et al., 2026; Zuo et al., 2025; Scott et al., 2026; Zheng et al., 2025). A^{2}PS is compatible with different Langevin solvers; we use the Fast Langevin Dynamics (FLD) solver of Zheng et al. (2025) for large-scale experiments.

5 Experiments and Discussion

We evaluate A^{2}PS around three key questions: (1) *Accuracy:* is its posterior-score estimation closer to the oracle than existing estimations? (2) *Error correction:* can re-equilibration remove accumulated errors, and does it remain robust under practical approximations? (3) *Practicality and generality:* how well does A^{2}PS perform at scale across different model families and representation spaces, and how do its practical components affect quality and efficiency?

We answer the first two questions on a controlled 2D inverse problem with access to oracle quantities, and the third on large-scale image inverse problems using pixel-space diffusion and latent-space flow-matching backbones.

5.1 Controlled 2D Study

Experiment Setup. We consider a controlled 2D inverse problem for which the oracle posterior density $p_{t|y}$ and its score $\boldsymbol{s}_{t|y}$ are analytically available at every generative time. As illustrated in Figure 2, we construct a complex prior such that conditioning on the measurement $\boldsymbol{y} = A(\boldsymbol{x})$ yields a structured multimodal posterior.

This setting allows us to evaluate two sources of error that are inaccessible in large-scale image models: posterior-score approximation and distributional mismatch along the generative trajectory. For the latter, we denote the density produced by an approximation method at time t by q_t and compare it with the oracle $p_{t|y}$ using total variation $\mathrm{TV}(p_{t|y}, q_t)$ and $\mathrm{KL}(p_{t|y}\|q_t)$. All metrics are computed directly from propagated densities rather than finite samples. An oracle-score run using the same numerical density solver provides a numerical error floor.

We compare DPS, DAPS, and A^{2}PS under favorable conditions designed to isolate their error-correction mechanisms. DPS represents the propagation-based guidance approaches based on a likelihood-score estimation but provide no explicit mechanism for removing inherited errors, whereas DAPS corrects them through posterior estimation and decoupled noise annealing. For a stronger comparison, we replace DAPS's original heuristic isotropic Gaussian approximation with the exact conditional mean and full covariance from Equation (7). A^{2}PS is likewise evaluated in an idealized setting with full conditional covariance and fully equilibrated Langevin dynamics. We then evaluate the practical setting with isotropic covariance $\mathrm{Tr}(\mathbf{\Sigma}_{1|t})\mathbf{I}/d$ and finite Langevin evolution.

Score accuracy. At each generative time, we measure the expected squared error between the estimated posterior score and the analytic oracle $\boldsymbol{s}_{t|y}$ under $p_{t|y}$. As shown in Figure 3, A^{2}PS yields substantially lower error than DPS throughout the trajectory, with the discrepancy rapidly decreasing near $t = 1$, consistent with the asymptotic accuracy established in Corollary 3.3.

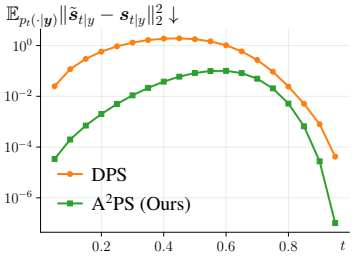


Figure 3: **Posterior-score error.** Expected squared error to the analytic oracle $s_{t|y}$ under $p_{t|y}$. A^{2}PS approaches the oracle near the generative endpoint.

Error correction. Accurate local scores alone cannot remove errors inherited from earlier generative times. We therefore evaluate *A^{2}PS w/o re-equilibration (re-eq)*, which uses the same posterior-score approximation as A^{2}PS but propagates the induced posterior velocity without Langevin correction.

As shown in Figure 4, without an error-correction mechanism, the distributional errors of DPS and A^{2}PS w/o re-eq continue to accumulate. In contrast, DAPS and A^{2}PS can reduce previously accumulated error. DAPS achieves only partial recovery and retains a clear terminal discrepancy, whereas A^{2}PS drives both TV and KL toward the numerical floor. The corresponding terminal distributions are shown in Figure 2, while Figure 7 in Appendix A provides the full density evolution.

These results separate the roles of the two components: the asymptotically accurate score provides an increasingly accurate posterior target, and re-equilibration removes earlier trajectory error.

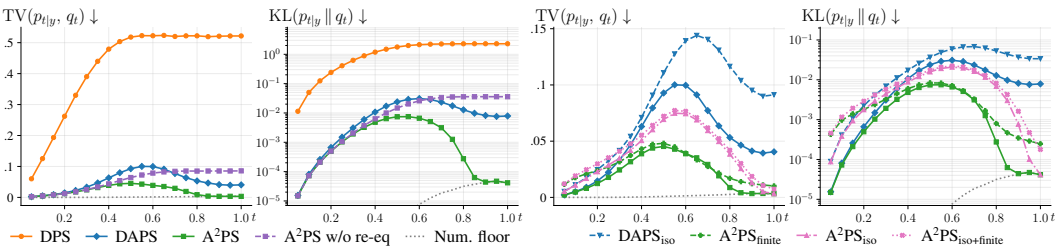


Figure 4: **Distributional error and ablations.** *Left:* TV and $\mathrm{KL}(p_{t|y}\|q_t)$ for DPS, DAPS, A^{2}PS, and A^{2}PS without re-equilibration; the oracle-score run defines the numerical floor. *Right:* A^{2}PS and DAPS under practical approximations, where "iso" uses isotropic conditional covariance and "finite" uses finite Langevin evolution.

Practical approximations. The right panels of Figure 4 replace the full conditional $\mathbf{\Sigma}_{1|t}$ with an isotropic approximation $\mathrm{Tr}(\mathbf{\Sigma}_{1|t})\mathbf{I}/d$ and converged Langevin dynamics with finite evolution ($\tau = 0.8$). Both approximations increase intermediate error, with their combination producing the largest degradation. Nevertheless, TV and KL decrease again near the endpoint, indicating that A^{2}PS remains effective under the approximations used in practice.

Table 1: **Large-scale image restoration.** Each cell reports LPIPS↓ / FID↓ / Runtime per image↓ (s). Runtime is measured on one RTX PRO 6000 GPU. Best results within each backbone and task are shown in **bold**. DDRM and DDNM$^+$ do not support nonlinear operators.

Method	Box Inpainting	Checkerboard Inpainting	Random 70% Inpainting	Super Resolution	Gaussian Deblurring	Non-linear Deblurring
	LPIPS↓ / FID↓ / Runtime per image↓ (s)					
Pixel-Space Diffusion on ImageNet 256×256						
DPS	.415 / 78.91 / 63	.387 / 68.84 / 63	.388 / 68.49 / 63	.493 / 85.52 / 63	.365 / 64.96 / 63	.393 / 70.87 / 66
DDRM	.188 / 55.00 / 35	.082 / 23.13 / 35	**.061 / 15.01** / 35	.257 / 47.28 / 32	.597 / 179.1 / 32	– / – / –
DDNM$^+$	.187 / 55.09 / 62	**.082** / 22.91 / 62	.061 / 15.04 / 62	.252 / 47.75 / 59	.760 / 347.4 / 59	– / – / –
DAPS	.214 / 68.44 / 17	.098 / 32.14 / 17	.073 / 23.70 / 17	.307 / 49.34 / 19	.307 / 50.02 / 26	.187 / 44.74 / 115
A^{2}PS (Ours)	**.163 / 45.35 / 7**	.088 / **20.00 / 7**	.074 / 23.28 / **7**	**.184 / 39.63 / 7**	**.222 / 40.56 / 15**	**.157 / 37.54 / 40**
Latent-Space Flow Matching (SD3.5-M) on FFHQ 768×768						
PnP-Flow	.149 / 64.79 / 67	.081 / 55.05 / 67	.110 / 56.79 / 67	.128 / 70.09 / 67	.220 / 86.33 / 67	.243 / 89.89 / 68
RLSD	.169 / 60.24 / 61	.065 / 62.83 / 61	.114 / 74.82 / 61	.072 / 52.05 / 61	.035 / 31.94 / 62	.106 / 31.23 / 65
FlowChef	.195 / 96.09 / 66	.179 / 78.26 / 66	.157 / 77.65 / 66	.070 / 68.40 / 67	.150 / 85.56 / 67	.077 / 46.75 / 71
FlowDPS	.292 / 96.55 / 93	.245 / 88.14 / 93	.271 / 93.97 / 93	.112 / 63.09 / 95	.193 / 73.08 / 94	.148 / 72.20 / 102
A^{2}PS (Ours)	**.103 / 38.45 / 56**	**.063 / 32.34 / 56**	**.076 / 32.82 / 56**	**.043 / 28.40 / 45**	**.027 / 20.33 / 47**	**.034 / 21.72 / 50**

The corresponding DAPS ablation further separates covariance accuracy from error correction. Replacing the full covariance with an isotropic approximation increases its error, confirming that a more accurate endpoint Gaussian is beneficial. However, even DAPS with the exact full covariance retains a clear terminal discrepancy. Thus, covariance accuracy alone is insufficient to eliminate accumulated trajectory error.

5.2 Large-scale Image Restoration

We next examine whether the advantages observed in the controlled setting extend to high-dimensional restoration across model backbones, representation spaces, and degradation operators.

Experiment Setup. We evaluate A^{2}PS with two substantially different generative priors: a pretrained pixel-space DM (Dhariwal & Nichol, 2021) on 1,000 ImageNet validation images at 256×256, and the latent-space FM Stable Diffusion 3.5 Medium (SD3.5-M) (Esser et al., 2024) on 1,000 FFHQ images at 768×768. For each backbone, $\hat{\sigma}_{1|t}$ is estimated once from 10,000 prior trajectories and reused across inverse problems; the resulting $\hat{\sigma}_{1|t}$ are shown in Figure 5. For practical re-equilibration, we use the accelerated Fast Langevin Dynamics (FLD) solver of Zheng et al. (2025).

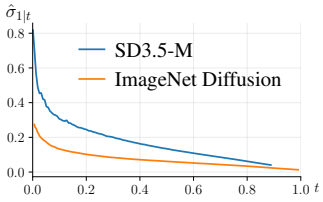


Figure 5: Statistics of $\hat{\sigma}_{1|t}$.

We consider six linear and nonlinear degradations: box, checkerboard, and 70% random inpainting, super-resolution, Gaussian deblurring, and nonlinear deblurring. We report LPIPS (Zhang et al., 2018) for perceptual similarity and FID (Heusel et al., 2017) for distribution-level fidelity. For the diffusion backbone, we compare against DPS (Chung et al., 2023), DDRM (Kawar et al., 2022), DDNM$^+$ (Wang et al., 2023), and DAPS (Zhang et al., 2025); for SD3.5-M, we compare against PnP-Flow (Martin et al., 2025), RLSD (Zilberstein et al., 2025), FlowChef (Patel et al., 2025), and FlowDPS (Kim et al., 2025). Detailed degradation settings and hyperparameters are provided in Appendix C.

Restoration quality and generality. Table 1 shows strong restoration performance across both model backbones. With SD3.5-M, A^{2}PS achieves the best LPIPS and FID on all six degradations. With the pixel-space diffusion backbone, it achieves the best LPIPS on four of six tasks and the best FID on five of six tasks. The remaining exceptions are checkerboard and random inpainting, where the linear-operator-specific DDRM or DDNM$^+$ performs better on some metrics.

Importantly, the same A^{2}PS formulation is applied across all six tasks, including nonlinear deblurring where these specialized linear methods are not applicable. Together, these results show that A^{2}PS generalizes across DM and FM backbones, pixel and latent spaces, and both linear and nonlinear inverse problems.

Efficiency. Despite performing repeated posterior re-equilibration, A^{2}PS is the fastest method across all evaluated tasks and both backbones. It requires only 7–56 seconds per image while maintaining the strongest overall perceptual performance, and provides approximately $1.6\times$–$2.7\times$ speedups over the fastest competing baseline on the diffusion backbone. These results show that posterior re-equilibration remains computationally practical at image scale.

To further examine the practical components of A^{2}PS, we conduct ablations on 100 ImageNet samples with the pixel-space DM. We study likelihood guidance on four inverse problems and re-equilibration efficiency on box inpainting.

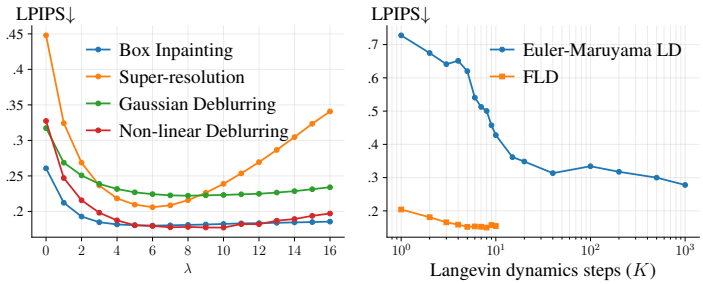


Figure 6: **Likelihood guidance and re-equilibration efficiency.** *Left:* LPIPS versus likelihood-guidance strength λ across four inverse problems. *Right:* LPIPS versus Langevin steps for standard Euler-Maruyama and accelerated FLD on box inpainting.

Likelihood guidance. The left panel of Figure 6 varies the additional likelihood-guidance strength λ across box inpainting, super-resolution, Gaussian deblurring, and nonlinear deblurring. Moderate guidance improves LPIPS across all tasks, whereas overly strong guidance eventually degrades restoration quality. This behavior reflects the role of λ in balancing measurement enforcement against the generative prior, with the preferred strength depending on the inverse problem.

Efficient re-equilibration. Because A^{2}PS repeatedly invokes Langevin dynamics for posterior re-equilibration, the efficiency of the LD solver directly affects the overall sampling cost. Re-equilibration depends on the target distribution rather than a particular Langevin integrator, allowing the inner solver to adopt accelerated algorithms (Zhang et al., 2026; Zuo et al., 2025; Scott et al., 2026; Zheng et al., 2025). As shown in the right panel of Figure 6, FLD reaches the low-LPIPS regime within roughly 10 steps, whereas the standard overdamped Euler-Maruyama discretization does not attain comparable quality within hundreds of steps. Thus, improvements in Langevin sampling can translate directly into a more efficient posterior re-equilibration in A^{2}PS.

6 Conclusion

We presented A^{2}PS, a training-free posterior sampling method for noisy inverse problems with diffusion and flow models. The main idea is to correct the approximation errors accumulated during generation. We showed that $p(\boldsymbol{x}_1 \mid \boldsymbol{x}_t)$ admits an asymptotically accurate Gaussian approximation as $t \to 1$, which yields a posterior-score estimate whose error vanishes near the generative endpoint. A^{2}PS uses this increasingly accurate estimate to re-equilibrate the current distribution with Langevin dynamics and thereby reduce errors inherited from earlier sampling steps.

Our controlled experiments isolate this effect and show that re-equilibration can recover from substantial intermediate distributional mismatch. Large-scale experiments further demonstrate that A^{2}PS scales to high-dimensional image restoration with both diffusion and flow-matching models, in pixel and latent spaces, while remaining computationally efficient across linear and nonlinear inverse problems. Taken together, these results suggest that posterior guidance need not be uniformly accurate throughout the entire trajectory: errors made earlier can still be corrected when more accurate posterior information becomes available near the endpoint. Extending this principle to convergence guarantees for finite-time re-equilibration and imperfect learned generative fields remains a direction for future work.

AI USE STATEMENT

In this work, we used generative AI tools for proposing and refining hypotheses, providing critical ingredients for proving mathematical claims, assisting in the writing of proofs, and assisting with translation. We have not used generative AI tools for developing theoretical models or conceptual frameworks, formulating mathematical claims, designing or providing feedback on research methodology or experiments, implementing methods, supporting qualitative and thematic data analysis, or interpreting results. The generation of synthetic data sets and cleaning or reformatting datasets were not applicable to this work.

Additionally, we used generative AI tools for discovering research topics and identifying research gaps, identifying relevant literature, summarizing or analyzing existing literature, searching for information, creating and modifying scientific figures or images, and editing the manuscript to improve readability. Generative AI tools were also used to create and edit software code for the 2D experiments. For the large-scale experiments, generative AI tools were used to assist with troubleshooting implementation issues, such as identifying possible causes and solutions for software bugs. The authors independently evaluated these suggestions, modified the code themselves, and verified the resulting implementation through testing.

We have reviewed all AI-assisted work. We checked LLM-generated research ideas for potential plagiarism through a manual literature survey, and we verified and tested the LLM-assisted code used for the 2D experiments. All mathematical proofs were independently checked by two authors. All LLM-generated edits and suggestions to the manuscript were reviewed and discussed by the authors before being incorporated into the final manuscript. We take responsibility for the final content of this work, including text, claims, code, and other artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Complete details of A^2PS are provided in Section 3, including its mathematical formulation, underlying principles, and practical realization. Given the flexibility of A^2PS, Algorithm 1 provides a concrete instantiation of A^2PS with a flow-matching backbone. Complete proofs of all theoretical results are provided in Appendix D. Experimental setups are described in the first part of each experimental section, while additional experimental details and hyperparameters for the large-scale experiments are provided in Appendix C. The source code is provided in the supplementary materials.

REFERENCES

Michael Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic Interpolants: A Unifying Framework for Flows and Diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025.

Gabriel Cardoso, Yazid Janati el idrissi, Sylvain Le Corff, and Eric Moulines. Monte Carlo guided Denoising Diffusion models for Bayesian linear inverse problems. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *Proceedings of the International Conference on Learning Representations (ICLR)*, volume 2024, pp. 44001–44037, 2024.

Bernard Chalmond. *Modeling and Inverse Problems in Imaging Analysis*. Number 155 in Applied Mathematical Sciences. Springer, New York, NY, 2003.

Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255, 2009.

Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794. Curran Associates, Inc., 2021.

Bradley Efron. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4396–4405. IEEE, 2019.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising Diffusion Restoration Models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 23593–23606. Curran Associates, Inc., 2022.

Jeongsol Kim, Bryan Sangwoo Kim, and Jong Chul Ye. Flowdps: Flow-driven posterior sampling for inverse problems. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12328–12337. IEEE, 2025.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow Matching for Generative Modeling. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

Ségolène Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. PnP-Flow: Plug-and-Play Image Restoration with Flow Matching. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *Proceedings of the International Conference on Learning Representations (ICLR)*, volume 2025, pp. 45466–45492, 2025.

Maitreya Patel, Song Wen, Dimitris N. Metaxas, and Yezhou Yang. FlowChef: Steering of Rectified Flow Models for Controlled Generations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15308–15318, 2025.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. PromptIR: Prompting for All-in-One Image Restoration. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 71275–71293. Curran Associates, Inc., 2023.

Herbert E. Robbins. An Empirical Bayes Approach to Statistics. In Samuel Kotz and Norman L. Johnson (eds.), *Breakthroughs in Statistics: Foundations and Basic Theory*, pp. 388–394. Springer New York, New York, NY, 1992.

Litu Rout, Negin Raoof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai. Solving Linear Inverse Problems Provably via Posterior Sampling with Latent Diffusion Models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 49960–49990. Curran Associates, Inc., 2023.

Litu Rout, Yujia Chen, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu. Beyond First-Order Tweedie: Solving Inverse Problems using Latent Diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9472–9481, 2024.

Maximilian Scott, Dáire O'Kane, Andraž Jelinčič, and James Foster. Underdamped Langevin MCMC with third order convergence. *Journal of Machine Learning Research*, 27(126):1–63, 2026.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

Phong Tran, Anh Tuan Tran, Quynh Phung, and Minh Hoai. Explore Image Deblurring via Encoded Blur Kernel Space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11956–11965, 2021.

Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 681–688, 2011.

Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. Freedom: Training-free energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 23117–23127. IEEE, 2023.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient Transformer for High-Resolution Image Restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5728–5739, 2022.

Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving Diffusion Inverse Problem Solving with Decoupled Noise Annealing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20895–20905, 2025.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595. IEEE, 2018.

Shiyuan Zhang, Qiwei Di, Xuheng Li, and Quanquan Gu. Dimension-Independent Convergence of Underdamped Langevin Monte Carlo in KL Divergence. In *ICLR 2026 2nd Workshop on Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*, 2026.

Candi Zheng, Yuan Lan, and Yang Wang. LanPaint: Training-Free Diffusion Inpainting with Asymptotically Exact and Fast Conditional Sampling. *Transactions on Machine Learning Research*, 2025.

Nicolas Zilberstein, Morteza Mardani, and Santiago Segarra. Repulsive Latent Score Distillation for Solving Inverse Problems. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025.

Xinzhe Zuo, Stanley Osher, and Wuchen Li. Gradient-Adjusted Underdamped Langevin Dynamics for Sampling. *SIAM/ASA Journal on Uncertainty Quantification*, 13(4):1735–1765, 2025.

A Additional Results

Figure 7 visualizes the posterior density evolution. DPS accumulates errors along the trajectory as the oracle posterior becomes increasingly complex near the endpoint. In contrast, DAPS and A^{2}PS reduce previously accumulated mismatch, with A^{2}PS most closely aligning with the posterior as $t \to 1$. This provides trajectory-level evidence that A^{2}PS improves posterior approximation through progressive correction along the sampling path.

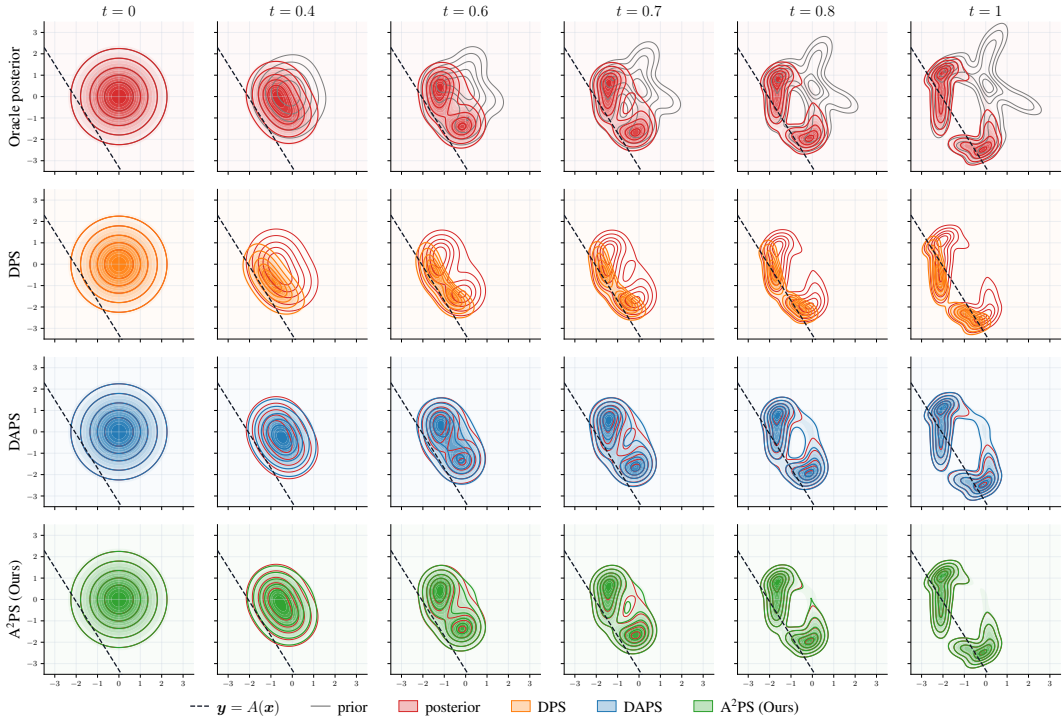


Figure 7: **Posterior evolution.** Columns denote generative time; rows show the oracle posterior path and the densities produced by DPS, DAPS, and A^{2}PS. Unlike propagation-only DPS, DAPS and A^{2}PS can reverse previously accumulated mismatch, with A^{2}PS increasingly aligning with the oracle path.

B Estimation of $\hat{\sigma}_{1|t}$

A^{2}PS approximates $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$ by $\mathcal{N}(\boldsymbol{x}_1; \boldsymbol{\mu}_{1|t}(\boldsymbol{x}), \boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}))$. Rather than estimating a state-dependent covariance online, we model $\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) \approx \hat{\sigma}^2_{1|t}\boldsymbol{I}$ as a function of the generative time.

We estimate the uncertainty profile from samples generated by prior model. For n trajectories in dimension d, we compute the root-mean-square distance between the conditional endpoint mean and the final generated sample,

$$\hat{\sigma}_{1|t} = \sqrt{\frac{1}{nd}\sum_{i=1}^{n}\left\|\boldsymbol{\mu}_{1|t}(\boldsymbol{x}_{i,t}) - \boldsymbol{x}_{i,1}\right\|^2}. \tag{13}$$

The resulting lookup table is interpolated at inference time. A scalar multiplier η rescales the profile to account for differences among backbones and degradation operators. This procedure does not estimate the full covariance in Equation (7); it provides an inexpensive average uncertainty scale suitable for the optimization in Equation (9).

C Experimental Details for Large-Scare Experiments

C.1 Details of the Dataset and Prior model Setup

The measurement-noise standard deviation is $\mathbf{\Sigma}_y = 0.05$ for ImageNet and $\mathbf{\Sigma}_y = 0.01$ for FFHQ. For all FFHQ tasks with SD3.5, we use the fixed text prompt "a photo of a close face", and classifier-free guidance scale of 2.

C.2 Description of Each Inverse Problems

The evaluation includes six linear and non-linear degradations:

1. **Box inpainting:** a central 128×128 region is removed from ImageNet images and a 384×384 region from FFHQ images.
2. **Checkerboard inpainting:** pixels are removed in a checkerboard pattern with 16×16 cells.
3. **Random inpainting:** 70% of pixels are masked uniformly at random.
4. **Super-resolution:** average pooling downsamples ImageNet by a factor of 4 and FFHQ by a factor of 8.
5. **Gaussian deblurring:** images are convolved with a 61×61 Gaussian kernel with standard deviation 3.
6. **Non-linear deblurring:** a neural degradation model produces non-linear motion blur (Tran et al., 2021), following the protocol of DPS.

C.3 Details for Baselines and A^{2}PS for Each Task

We tuned baseline hyperparameters according to the authors' recommendations and comparable runtime budgets. When a paper did not report settings for a task considered here, we selected the configuration by grid search.

C.3.1 ImageNet Task with Pixel-level Diffusion Model

DPS We use the default hyperparameters for all tasks.

DDRM We use the default parameters $\eta_B = 1$ and $\eta = 0.85$ with 500 diffusion steps.

DDNM$^+$ We use the default parameters $\eta = 0.85$, $l = 1$, and $r = 1$ with 250 diffusion steps.

DAPS We use 100 Langevin steps, $N_A = 150$, and $N_{\mathrm{ODE}} = 3$.

A^{2}PS-Diffusion We use 20 diffusion steps and Fast Langevin Dynamics (FLD). Here N_{LD} is the number of FLD steps, η_{LD} is the FLD step size, Γ is the FLD friction parameter, η_{GD} is the gradient-descent learning rate for Equation (9), N_{GD} is the number of gradient steps, and η rescales $\hat{\sigma}_{1|t}$.

- Gaussian deblurring: $N_{\mathrm{LD}} = 5$, $\lambda = 6$, $\eta_{\mathrm{LD}} = 0.15$, $\Gamma = 15$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 50$, $\eta = 1$.
- Super-resolution: $N_{\mathrm{LD}} = 5$, $\lambda = 6$, $\eta_{\mathrm{LD}} = 0.15$, $\Gamma = 15$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 50$, $\eta = 1$.
- Inpainting: $N_{\mathrm{LD}} = 5$, $\lambda = 3$, $\eta_{\mathrm{LD}} = 0.35$, $\Gamma = 25$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 50$, $\eta = 1.5$.
- Non-linear deblurring: $N_{\mathrm{LD}} = 5$, $\lambda = 6$, $\eta_{\mathrm{LD}} = 0.15$, $\Gamma = 15$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 50$, $\eta = 1$.

C.3.2 FFHQ task with Latent-Based Flow-matching Model

PnP-Flow We adapt PnP-Flow to Stable Diffusion 3.5 Medium and perform a linear hyperparameter search for each task. We use 50 ODE steps and 10 samples per step.

- Gaussian deblurring: $\gamma = 1$, $\alpha = 1$.
- Super-resolution: $\gamma = 90$, $\alpha = 0.5$.

- Inpainting: $\gamma = 0.6$, $\alpha = 0.06$.
- Non-linear deblurring: $\gamma = 1$, $\alpha = 1$.

FlowChef We use the default $s = 15$ and $\mathrm{NFE} = 400$. For inpainting, we set $s = 1$.

FlowDPS We use the default $s = 15$, three gradient-descent steps, and $\mathrm{NFE} = 250$. For inpainting, we set $s = 6$.

RLSD We adapt RLSD to Stable Diffusion 3.5 Medium and use 250 ODE steps.

- Gaussian deblurring: $l_{r_x} = 0.2$, $l_{r_z} = 0.2$, $\tilde{\rho} = 0.2$, $\lambda_t = 0.1$.
- Super-resolution: $l_{r_x} = 0.2$, $l_{r_z} = 0.2$, $\tilde{\rho} = 0.2$, $\lambda_t = 0.1$.
- Inpainting: $l_{r_x} = 0.2$, $l_{r_z} = 0.3$, $\tilde{\rho} = 0.1$, $\lambda_t = 0.1$.
- Non-linear deblurring: $l_{r_x} = 0.2$, $l_{r_z} = 0.2$, $\tilde{\rho} = 0.2$, $\lambda_t = 0.1$.

A^{2}PS-Flow We use 20 flow-ODE steps with Stable Diffusion 3.5 Medium.

- Gaussian deblurring: $N_{\mathrm{LD}} = 1$, $\lambda = 9$, $\eta_{\mathrm{LD}} = 0.15$, $\Gamma = 15$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 10$, $\eta = 8$.
- Super-resolution: $N_{\mathrm{LD}} = 1$, $\lambda = 9$, $\eta_{\mathrm{LD}} = 0.15$, $\Gamma = 15$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 10$, $\eta = 8$.
- Inpainting: $N_{\mathrm{LD}} = 2$, $\lambda = 2.5$, $\eta_{\mathrm{LD}} = 0.35$, $\Gamma = 25$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 6$, $\eta = 8$.
- Non-linear deblurring: $N_{\mathrm{LD}} = 1$, $\lambda = 9$, $\eta_{\mathrm{LD}} = 0.15$, $\Gamma = 15$, $\eta_{\mathrm{GD}} = 0.1$, $N_{\mathrm{GD}} = 10$, $\eta = 8$.

D Asymptotic Gaussianity for Conditional Endpoint

D.1 Definitions and Assumptions of the Density Path

We consider the conditional Gaussian density path

$$p_t(\boldsymbol{x} \mid \mathbf{x}_1 = \boldsymbol{x}_1) = \mathcal{N}(\boldsymbol{x};\, a_t \boldsymbol{\mu}_0 + b_t \boldsymbol{x}_1,\, a_t^2 \sigma_0^2 \mathbf{I}), \tag{14}$$

where $\mathbf{x}_0 \sim \mathcal{N}(\boldsymbol{\mu}_0, \sigma_0^2 \mathbf{I})$ is independent of $\mathbf{x}_1$. We assume that $a_t > 0$ and $b_t > 0$ along the density path, with

$$a_t \to 0, \qquad b_t \to 1, \qquad \text{as } t \to 1. \tag{15}$$

Consequently, the marginal density at an intermediate time is

$$p_t(\boldsymbol{x}) = \int_{\mathbb{R}^d} p_1(\boldsymbol{x}_1) \mathcal{N}(\boldsymbol{x};\, a_t \boldsymbol{\mu}_0 + b_t \boldsymbol{x}_1,\, \sigma_0^2 a_t^2 \mathbf{I})\, \mathrm{d}\boldsymbol{x}_1. \tag{16}$$

Throughout the analysis, we define the effective noise scale

$$\rho_t \coloneqq \frac{\sigma_0 a_t}{b_t}, \qquad \rho_t \to 0 \quad \text{as } t \to 1. \tag{17}$$

We use the following notation for the conditional mean and covariance of the generative endpoint:

$$\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) \coloneqq \mathbb{E}[\mathbf{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}], \qquad \boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) \coloneqq \mathrm{Cov}(\mathbf{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}). \tag{18}$$

Regularity assumptions. We assume that the endpoint density p_1 is strictly positive and smooth in an open neighborhood U containing the endpoint region of interest:

$$p_1(\boldsymbol{u}) > 0, \qquad p_1 \in C^\infty(U). \tag{19}$$

It follows immediately that $\log p_1 \in C^\infty(U)$. Hence, on every compact subset $K \Subset U$, all finite-order derivatives of $\log p_1$ are uniformly bounded. Accordingly, no separate local boundedness assumptions on the endpoint score or its higher-order derivatives are required.

All asymptotic statements below are understood locally on such compact subsets of U. More precisely, we consider intermediate states $\boldsymbol{x}$ for which the corresponding endpoint coordinate

$$\frac{\boldsymbol{x} - a_t \boldsymbol{\mu}_0}{b_t} \tag{20}$$

remains in a fixed compact subset $K \Subset U$ as $t \to 1$.

This assumption also provides the required local regularity of the intermediate density. To see this, let

$$\phi_\rho(\boldsymbol{u}) \coloneqq \mathcal{N}(\boldsymbol{u};\, \mathbf{0},\, \rho^2\mathbf{I}). \tag{21}$$

Using $\sigma_0 a_t = b_t \rho_t$, Equation (16) can be written as

$$p_t(\boldsymbol{x}) = \frac{1}{b_t^d}\,(p_1 * \phi_{\rho_t})\left(\frac{\boldsymbol{x} - a_t\boldsymbol{\mu}_0}{b_t}\right). \tag{22}$$

Therefore, by the standard local approximation properties of Gaussian mollification, the smoothness and strict positivity of p_1 imply that, on every fixed compact subset $K \Subset U$, the derivatives of $\log(p_1 * \phi_{\rho_t})$ of any fixed finite order remain uniformly bounded as $\rho_t \to 0$. In particular,

$$\nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x}) = \frac{1}{b_t^2}\nabla_{\boldsymbol{u}}^2 \log\left(p_1 * \phi_{\rho_t}\right)(\boldsymbol{u})\Big|_{\boldsymbol{u}=(\boldsymbol{x}-a_t\boldsymbol{\mu}_0)/b_t} = \mathcal{O}(1) \quad \text{as } t \to 1. \tag{23}$$

For a fixed observation y, we further assume that the likelihood is strictly positive and sufficiently smooth in the same neighborhood:

$$p(y \mid \boldsymbol{u}) > 0, \qquad \log p(y \mid \boldsymbol{u}) \in C^3(U). \tag{24}$$

As above, its derivatives up to third order are automatically uniformly bounded on every compact subset $K \Subset U$.

The following lemma characterizes the conditional covariance of the generative endpoint and establishes its small-noise asymptotics.

Lemma D.1. *Under the density path in Equation* (14) *and the regularity condition in Equation* (19)*, the conditional covariance satisfies the exact identity*

$$\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) = b_t\rho_t^2\frac{\partial}{\partial \boldsymbol{x}}\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) = \rho_t^2\mathbf{I} + b_t^2\rho_t^4\nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x}). \tag{25}$$

Consequently, locally on the endpoint region specified above,

$$\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) = \rho_t^2\mathbf{I} + \mathcal{O}(\rho_t^4) \quad \textit{as} \quad t \to 1. \tag{26}$$

Proof. From Equation (14), we have

$$\mathbf{x}_t \mid \mathbf{x}_1 = \boldsymbol{x}_1 \sim \mathcal{N}\big(a_t\boldsymbol{\mu}_0 + b_t\boldsymbol{x}_1,\, \sigma_0^2 a_t^2\mathbf{I}\big)\,. \tag{27}$$

Applying Tweedie's formula to this Gaussian corruption gives

$$\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) = \frac{\boldsymbol{x} - a_t\boldsymbol{\mu}_0}{b_t} + \frac{\sigma_0^2 a_t^2}{b_t}\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}). \tag{28}$$

Since

$$\sigma_0^2 a_t^2 = b_t^2\rho_t^2, \tag{29}$$

differentiating Equation (28) with respect to $\boldsymbol{x}$ yields

$$\frac{\partial}{\partial \boldsymbol{x}}\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) = \frac{1}{b_t}\mathbf{I} + b_t\rho_t^2\nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x}). \tag{30}$$

We next relate the Jacobian of the conditional mean to the conditional covariance. By Bayes' rule,

$$p(\boldsymbol{x}_1 \mid \boldsymbol{x}) \propto p_1(\boldsymbol{x}_1)\exp\!\left(-\frac{\|\boldsymbol{x} - a_t\boldsymbol{\mu}_0 - b_t\boldsymbol{x}_1\|^2}{2\sigma_0^2 a_t^2}\right). \tag{31}$$

Differentiating its log-density with respect to $\boldsymbol{x}$ gives

$$\nabla_{\boldsymbol{x}} \log p(\boldsymbol{x}_1 \mid \boldsymbol{x}) = \frac{b_t}{\sigma_0^2 a_t^2}\left(\boldsymbol{x}_1 - \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\right). \tag{32}$$

Therefore, differentiating the conditional expectation under the integral,

$$\begin{aligned}\frac{\partial}{\partial \boldsymbol{x}}\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) &= \int \boldsymbol{x}_1 \left[\nabla_{\boldsymbol{x}} p(\boldsymbol{x}_1 \mid \boldsymbol{x})\right]^\top \mathrm{d}\boldsymbol{x}_1 \\ &= \frac{b_t}{\sigma_0^2 a_t^2}\int \boldsymbol{x}_1 \left(\boldsymbol{x}_1 - \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\right)^\top p(\boldsymbol{x}_1 \mid \boldsymbol{x})\,\mathrm{d}\boldsymbol{x}_1 \\ &= \frac{b_t}{\sigma_0^2 a_t^2}\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}). \end{aligned} \tag{33}$$

Using $\sigma_0^2 a_t^2 = b_t^2 \rho_t^2$, this is equivalently

$$\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) = b_t \rho_t^2 \frac{\partial}{\partial \boldsymbol{x}}\boldsymbol{\mu}_{1|t}(\boldsymbol{x}). \tag{34}$$

Substituting Equation (30) into Equation (34) gives

$$\begin{aligned}\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) &= b_t \rho_t^2 \left(\frac{1}{b_t}\mathbf{I} + b_t \rho_t^2 \nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x})\right) \\ &= \rho_t^2 \mathbf{I} + b_t^2 \rho_t^4 \nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x}). \end{aligned} \tag{35}$$

Finally, Equation (23) gives

$$\nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x}) = \mathcal{O}(1) \qquad \text{as } t \to 1 \tag{36}$$

uniformly on the local region under consideration. Since $b_t \to 1$,

$$b_t^2 \rho_t^4 \nabla_{\boldsymbol{x}}^2 \log p_t(\boldsymbol{x}) = \mathcal{O}(\rho_t^4), \tag{37}$$

which proves Equation (26). □

D.2 Proof of Theorem D.2

Theorem D.2 (Asymptotic Gaussianity of the Endpoint Conditional)**.** *As $t \to 1$ (and thus $\rho_t \to 0$), the conditional endpoint density satisfies*

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) = \mathcal{N}\big(\boldsymbol{x}_1;\, \boldsymbol{\mu}_{1|t}(\boldsymbol{x}),\, \boldsymbol{\Sigma}_{1|t}(\boldsymbol{x})\big)\left[1 + \mathcal{O}(\rho_t^3)\right] \tag{38}$$

uniformly on every standardized compact neighborhood. Moreover, the exact conditional covariance satisfies

$$\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x}) = \frac{\sigma_0^2 a_t^2}{b_t}\frac{\partial \boldsymbol{\mu}_{1|t}(\boldsymbol{x})}{\partial \boldsymbol{x}} = \rho_t^2 \mathbf{I} + \mathcal{O}(\rho_t^4). \tag{39}$$

Proof. For notational brevity throughout the proof, we drop the explicit dependence on $\boldsymbol{x}$ for the conditional moments, writing $\boldsymbol{\mu}_{1|t}$ and $\boldsymbol{\Sigma}_{1|t}$. Let us define the unconditional rescaled mean

$$\boldsymbol{\mu}_t \coloneqq \frac{\boldsymbol{x} - a_t \boldsymbol{\mu}_0}{b_t}. \tag{40}$$

By Equation (14), the exact posterior density of $\mathbf{x}_1$ given $\mathbf{x}_t = \boldsymbol{x}$ is obtained via Bayes' rule:

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) = \frac{p_1(\boldsymbol{x}_1)\exp\!\left(-\frac{\|\boldsymbol{x}_1 - \boldsymbol{\mu}_t\|^2}{2\rho_t^2}\right)}{\displaystyle\int_{\mathbb{R}^d} p_1(\boldsymbol{u})\exp\!\left(-\frac{\|\boldsymbol{u} - \boldsymbol{\mu}_t\|^2}{2\rho_t^2}\right)\mathrm{d}\boldsymbol{u}}. \tag{41}$$

This shows that the conditional distribution is a Gaussian measure with covariance $\rho_t^2 \mathbf{I}$ tilted by the prior p_1.

From Lemma D.1, we have already established that the true conditional mean satisfies

$$\boldsymbol{\mu}_{1|t} = \boldsymbol{\mu}_t + b_t \rho_t^2 \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}). \tag{42}$$

Because the Gaussian mollification ensures $\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}) = \mathcal{O}(1)$ locally, we have a precise order for the mean shift:

$$\boldsymbol{\mu}_{1|t} - \boldsymbol{\mu}_t = \mathcal{O}(\rho_t^2). \tag{43}$$

Let us recenter the analysis around the exact conditional mean by defining the local variable

$$\boldsymbol{u} \coloneqq \boldsymbol{x}_1 - \boldsymbol{\mu}_{1|t}. \tag{44}$$

The distance in the Gaussian exponent becomes

$$\boldsymbol{x}_1 - \boldsymbol{\mu}_t = \boldsymbol{u} + \big(\boldsymbol{\mu}_{1|t} - \boldsymbol{\mu}_t\big). \tag{45}$$

To tightly control the linear terms in the log-density expansion, we use the exact score identity. Differentiating the logarithm of Equation (41) with respect to $\boldsymbol{x}_1$ and taking the conditional expectation yields:

$$0 = -\frac{\boldsymbol{\mu}_{1|t} - \boldsymbol{\mu}_t}{\rho_t^2} + \mathbb{E}[\nabla_{\boldsymbol{x}_1} \log p_1(\mathbf{x}_1) \,|\, \mathbf{x}_t = \boldsymbol{x}]\,. \tag{46}$$

By the local C^∞ (and thus C^2) property of p_1 (Equation (19)) and the fact that $\mathbb{E}[\boldsymbol{u} \mid \mathbf{x}_t = \boldsymbol{x}] = 0$, expanding the score around $\boldsymbol{\mu}_{1|t}$ gives:

$$\mathbb{E}[\nabla \log p_1(\mathbf{x}_1) \,|\, \mathbf{x}_t = \boldsymbol{x}] = \nabla \log p_1(\boldsymbol{\mu}_{1|t}) + \mathcal{O}(\rho_t^2). \tag{47}$$

Combining Equation (46) and Equation (47) precisely characterizes the mean shift:

$$\frac{\boldsymbol{\mu}_{1|t} - \boldsymbol{\mu}_t}{\rho_t^2} = \nabla \log p_1(\boldsymbol{\mu}_{1|t}) + \mathcal{O}(\rho_t^2). \tag{48}$$

We now execute the local Taylor expansion. Let

$$\boldsymbol{H}_t \coloneqq \nabla^2 \log p_1(\boldsymbol{\mu}_{1|t}). \tag{49}$$

Since $p_1 \in C^2(U)$, we can expand the log-prior around the conditional mean:

$$\log p_1(\boldsymbol{\mu}_{1|t} + \boldsymbol{u}) = \log p_1(\boldsymbol{\mu}_{1|t}) + \nabla \log p_1(\boldsymbol{\mu}_{1|t})^\top \boldsymbol{u} + \frac{1}{2}\boldsymbol{u}^\top \boldsymbol{H}_t \boldsymbol{u} + \mathcal{O}(\|\boldsymbol{u}\|^3). \tag{50}$$

Substituting Equation (45) and Equation (50) into the logarithm of the numerator in Equation (41), and applying the refined mean shift Equation (48), we obtain:

$$\log p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) = C_t - \frac{\|\boldsymbol{u}\|^2}{2\rho_t^2} + \frac{1}{2}\boldsymbol{u}^\top \boldsymbol{H}_t \boldsymbol{u} + \mathcal{O}(\rho_t^3), \tag{51}$$

where C_t absorbs terms independent of $\boldsymbol{u}$. The linear term in $\boldsymbol{u}$ rigorously cancels out up to $\mathcal{O}(\rho_t^3)$ because

$$\left(\nabla \log p_1(\boldsymbol{\mu}_{1|t}) - \frac{\boldsymbol{\mu}_{1|t} - \boldsymbol{\mu}_t}{\rho_t^2}\right)^\top \boldsymbol{u} = \mathcal{O}(\rho_t^2) \cdot \mathcal{O}(\rho_t) = \mathcal{O}(\rho_t^3) \tag{52}$$

for any standard neighborhood $\|\boldsymbol{u}\| = \mathcal{O}(\rho_t)$.

To handle the normalization integral (denominator of Equation (41)), we note that p_1 is a probability measure ($\int p_1 = 1$). For any $\delta > 0$ such that the δ-neighborhood of $\boldsymbol{\mu}_{1|t}$ is in U, the Gaussian tail strictly bounds the outer integral:

$$\int_{\{\|\boldsymbol{x}_1 - \boldsymbol{\mu}_{1|t}\| > \delta\}} p_1(\boldsymbol{x}_1) \exp\left(-\frac{\|\boldsymbol{x}_1 - \boldsymbol{\mu}_t\|^2}{2\rho_t^2}\right) \mathrm{d}\boldsymbol{x}_1 \le \exp\left(-\frac{c}{\rho_t^2}\right) \int_{\mathbb{R}^d} p_1(\boldsymbol{x}_1)\, \mathrm{d}\boldsymbol{x}_1 = o(\rho_t^k) \tag{53}$$

for any $k > 0$. Thus, global assumptions are unnecessary; integrating the local Gaussian expansion in Equation (51) uniquely determines the normalization and covariance. The same local expansion yields the exact covariance derived in Lemma D.1:

$$\boldsymbol{\Sigma}_{1|t} = \rho_t^2 \mathbf{I} + \rho_t^4 \boldsymbol{H}_t + \mathcal{O}(\rho_t^6). \tag{54}$$

Using the matrix expansion $\boldsymbol{\Sigma}_{1|t}^{-1} = \rho_t^{-2}\mathbf{I} - \boldsymbol{H}_t + \mathcal{O}(\rho_t^2)$, the logarithm of the exact Gaussian distribution $\mathcal{N}(\boldsymbol{x}_1; \boldsymbol{\mu}_{1|t}, \boldsymbol{\Sigma}_{1|t})$ expands as:

$$\log \mathcal{N}\big(\boldsymbol{x}_1;\, \boldsymbol{\mu}_{1|t},\, \boldsymbol{\Sigma}_{1|t}\big) = \widetilde{C}_t - \frac{\|\boldsymbol{u}\|^2}{2\rho_t^2} + \frac{1}{2}\boldsymbol{u}^\top \boldsymbol{H}_t \boldsymbol{u} + \mathcal{O}(\rho_t^4). \tag{55}$$

Subtracting Equation (55) from Equation (51) proves that the log-densities match up to $\mathcal{O}(\rho_t^3)$:

$$\log \frac{p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})}{\mathcal{N}\big(\boldsymbol{x}_1;\, \boldsymbol{\mu}_{1|t},\, \boldsymbol{\Sigma}_{1|t}\big)} = \mathcal{O}(\rho_t^3). \tag{56}$$

Exponentiating this ratio via $\exp(\mathcal{O}(\rho_t^3)) = 1 + \mathcal{O}(\rho_t^3)$ yields Equation (38).

Finally, Equation (39) is directly imported from Lemma D.1, completing the proof. □

D.3 Proof of Proposition 3.2

Proposition 3.2. *Define $\tilde{\boldsymbol{\mu}}_{1|t,y}$ to be the mode of $\hat{p}(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y})$:*

$$\tilde{\boldsymbol{\mu}}_{1|t,y}(\boldsymbol{x}) \coloneqq \arg\min_{\tilde{\boldsymbol{x}}} \left\{ \frac{1}{2}\|A(\tilde{\boldsymbol{x}}) - \boldsymbol{y}\|^2_{\boldsymbol{\Sigma}_y^{-1}} + \frac{1}{2}\|\tilde{\boldsymbol{x}} - \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\|^2_{\boldsymbol{\Sigma}_{1|t}^{-1}(\boldsymbol{x})} \right\}, \tag{9}$$

$\tilde{\boldsymbol{\mu}}_{1|t,y}$ reach endpoint posterior mean $\boldsymbol{\mu}_{1|t,y}$ asymptotically with t, $\tilde{\boldsymbol{\mu}}_{1|t,y} - \boldsymbol{\mu}_{1|t,y} = \mathcal{O}(\rho_t^4)$.

Recall the measurement model

$$\boldsymbol{y} = A(\boldsymbol{x}_1) + \boldsymbol{\delta}, \qquad \boldsymbol{\delta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_y), \qquad \boldsymbol{\Sigma}_y \succ 0. \tag{57}$$

For a fixed observation $\boldsymbol{y}$, the log-likelihood is

$$\log p(\boldsymbol{y} \mid \boldsymbol{x}_1) = -\frac{1}{2}\|A(\boldsymbol{x}_1) - \boldsymbol{y}\|^2_{\boldsymbol{\Sigma}_y^{-1}} + C_y, \tag{58}$$

where C_y is independent of $\boldsymbol{x}_1$. The approximate likelihood-conditioned posterior introduced in Equation (8) is defined by tilting the exact conditional Gaussian prior:

$$\widehat{p}(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}) \propto \exp\left(-\frac{1}{2}\|A(\boldsymbol{x}_1) - \boldsymbol{y}\|^2_{\boldsymbol{\Sigma}_y^{-1}} - \frac{1}{2}\|\boldsymbol{x}_1 - \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\|^2_{\boldsymbol{\Sigma}_{1|t}(\boldsymbol{x})^{-1}} \right). \tag{59}$$

Thus, the surrogate posterior mode $\tilde{\boldsymbol{\mu}}_{1|t,y}(\boldsymbol{x})$ is the global minimizer of the negative log-density in Equation (59).

For notational brevity, we drop the explicit dependence on $\boldsymbol{x}$ and write $\boldsymbol{\mu}_{1|t}$ and $\boldsymbol{\Sigma}_{1|t}$. We first establish that this minimizer is highly localized. By optimality, the objective value at $\tilde{\boldsymbol{\mu}}_{1|t,y}$ is no larger than its value at the prior mean $\boldsymbol{\mu}_{1|t}$, yielding

$$\frac{1}{2}\left\|\tilde{\boldsymbol{\mu}}_{1|t,y} - \boldsymbol{\mu}_{1|t}\right\|^2_{\boldsymbol{\Sigma}_{1|t}^{-1}} \le \frac{1}{2}\left\|A(\boldsymbol{\mu}_{1|t}) - \boldsymbol{y}\right\|^2_{\boldsymbol{\Sigma}_y^{-1}}. \tag{60}$$

Because $\boldsymbol{\mu}_{1|t}$ remains in a compact subset of U and A is continuous, the right-hand side is uniformly bounded as $t \to 1$. By Lemma D.1, $\boldsymbol{\Sigma}_{1|t} = \rho_t^2\mathbf{I} + \mathcal{O}(\rho_t^4)$, which implies the largest eigenvalue satisfies $\lambda_{\max}(\boldsymbol{\Sigma}_{1|t}) = \mathcal{O}(\rho_t^2)$. It immediately follows from Equation (60) that the mode collapses toward the prior mean at a rate of

$$\left\|\tilde{\boldsymbol{\mu}}_{1|t,y} - \boldsymbol{\mu}_{1|t}\right\| = \mathcal{O}(\rho_t). \tag{61}$$

Consequently, $\tilde{\boldsymbol{\mu}}_{1|t,y} \in U$ for all t sufficiently close to 1.

Let the displacement be

$$\Delta\boldsymbol{\mu}_t \coloneqq \tilde{\boldsymbol{\mu}}_{1|t,y} - \boldsymbol{\mu}_{1|t} = \mathcal{O}(\rho_t). \tag{62}$$

The first-order optimality condition of the minimizer gives

$$\Delta\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_{1|t}\nabla \log p\left(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t} + \Delta\boldsymbol{\mu}_t\right). \tag{63}$$

Since $\boldsymbol{\Sigma}_{1|t} = \mathcal{O}(\rho_t^2)$ and the likelihood score is uniformly bounded on U, we can bootstrap Equation (61) to a sharper bound:

$$\Delta\boldsymbol{\mu}_t = \mathcal{O}(\rho_t^2). \tag{64}$$

Applying Taylor's theorem to the likelihood score around $\boldsymbol{\mu}_{1|t}$ yields

$$\nabla \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t} + \Delta\boldsymbol{\mu}_t) = \nabla \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}) + \nabla^2 \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t})\Delta\boldsymbol{\mu}_t + \mathcal{O}(\|\Delta\boldsymbol{\mu}_t\|^2). \tag{65}$$

Substituting this back into Equation (63) and noting that $\|\Delta\boldsymbol{\mu}_t\|^2 = \mathcal{O}(\rho_t^4)$, we obtain

$$\Delta\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_{1|t}\nabla \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}) + \boldsymbol{\Sigma}_{1|t}\nabla^2 \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t})\Delta\boldsymbol{\mu}_t + \mathcal{O}(\rho_t^6). \tag{66}$$

Because the Hessian is bounded and $\|\boldsymbol{\Sigma}_{1|t}\| = \mathcal{O}(\rho_t^2)$, the second term is $\mathcal{O}(\rho_t^4)$. Therefore, the mode shift is explicitly characterized by

$$\tilde{\boldsymbol{\mu}}_{1|t,y} = \boldsymbol{\mu}_{1|t} + \boldsymbol{\Sigma}_{1|t}\nabla \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}) + \mathcal{O}(\rho_t^4). \tag{67}$$

Next, we establish that the mean of this surrogate posterior, denoted $\widehat{\boldsymbol{\mu}}_{1|t,y}$, is exceptionally close to its mode. Introducing the standardized coordinate

$$\boldsymbol{u} = \boldsymbol{\Sigma}_{1|t}^{-1/2}\big(\boldsymbol{x}_1 - \tilde{\boldsymbol{\mu}}_{1|t,y}\big), \tag{68}$$

we expand the negative log-density of $\widehat{p}$ around the mode. The linear terms cancel identically by the optimality condition. The quadratic term is dominated by $-\frac{1}{2}\|\boldsymbol{u}\|^2$, perturbed by $\mathcal{O}(\rho_t^2\|\boldsymbol{u}\|^2)$ from the likelihood Hessian. The first asymmetric term capable of shifting the mean is cubic, arising from the third derivative of the likelihood, and is scaled by $\boldsymbol{\Sigma}_{1|t}^{3/2} = \mathcal{O}(\rho_t^3)$. Hence, the standardized surrogate density behaves as

$$\log \widehat{p}_u(\boldsymbol{u}) = \widetilde{C}_t - \frac{1}{2}\|\boldsymbol{u}\|^2 + \mathcal{O}\big(\rho_t^2\|\boldsymbol{u}\|^2 + \rho_t^3\|\boldsymbol{u}\|^3\big). \tag{69}$$

Standard Gaussian integration over this slightly skewed density yields a mean shift in the standardized space of $\mathbb{E}_{\widehat{p}}[\boldsymbol{u}] = \mathcal{O}(\rho_t^3)$. Mapping back to the original space multiplies this by $\boldsymbol{\Sigma}_{1|t}^{1/2} = \mathcal{O}(\rho_t)$, giving

$$\widehat{\boldsymbol{\mu}}_{1|t,y} - \tilde{\boldsymbol{\mu}}_{1|t,y} = \mathcal{O}(\rho_t^4). \tag{70}$$

Finally, we compare the mean of the exact likelihood-conditioned posterior, $\boldsymbol{\mu}_{1|t,y}$, to the surrogate mean $\widehat{\boldsymbol{\mu}}_{1|t,y}$. By Theorem D.2, the exact prior locally matches the surrogate prior with a relative error of $\mathcal{O}(\rho_t^3)$. Multiplying both by the identical observation likelihood $p(\boldsymbol{y} \mid \boldsymbol{x}_1)$ and re-normalizing preserves this local error:

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}) = \widehat{p}(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y})\left[1 + \mathcal{O}(\rho_t^3)\right] \tag{71}$$

within any standardized neighborhood. Crucially, because p_1 integrates to 1, the Gaussian exponent in the conditional prior heavily penalizes deviations. The exact and surrogate posteriors thus assign exponentially vanishing mass to the exterior of this neighborhood, effectively truncating any global integration errors to $o(\rho_t^k)$. The difference in their means is therefore entirely determined by the local error $\mathcal{O}(\rho_t^3)$ acting on the characteristic spatial scale of the posterior, which is $\mathcal{O}(\rho_t)$. This yields

$$\boldsymbol{\mu}_{1|t,y} - \widehat{\boldsymbol{\mu}}_{1|t,y} = \mathcal{O}(\rho_t^4). \tag{72}$$

Combining Equation (70) and Equation (72), the difference between the true posterior mean and the surrogate mode is

$$\tilde{\boldsymbol{\mu}}_{1|t,y}(\boldsymbol{x}) = \boldsymbol{\mu}_{1|t,y}(\boldsymbol{x}) + \mathcal{O}(\rho_t^4), \tag{73}$$

completing the proof.

E Approximation Mechanisms of Existing Posterior Samplers

We provide a detailed comparison of how existing training-free posterior samplers make the otherwise intractable conditioning operation computationally tractable. For consistency with the main text, we follow the convention where $\mathbf{x}_1$ denotes the clean endpoint (which several diffusion-based works denote by $\boldsymbol{x}_0$).

For a measurement $\mathbf{y} = \boldsymbol{y}$, the exact intermediate likelihood is given by

$$p_t(\boldsymbol{y} \mid \boldsymbol{x}) = \int p(\boldsymbol{y} \mid \boldsymbol{x}_1) p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})\, \mathrm{d}\boldsymbol{x}_1. \tag{74}$$

The primary challenge for likelihood-guided posterior sampling lies in evaluating the intractable endpoint conditional $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$ or quantities derived from it.

DPS. DPS (Chung et al., 2023) replaces the expectation in Equation (74) by evaluating the likelihood directly at the conditional endpoint mean:

$$p_t(\boldsymbol{y} \mid \boldsymbol{x}) = \mathbb{E}[p(\boldsymbol{y} \mid \mathbf{x}_1) \mid \mathbf{x}_t = \boldsymbol{x}] \approx p\left(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\right), \tag{75}$$

where $\boldsymbol{\mu}_{1|t}(\boldsymbol{x}) = \mathbb{E}[\mathbf{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}]$ is obtained from the pretrained generative model via Tweedie's formula. The resulting likelihood-score approximation is

$$\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{y} \mid \boldsymbol{x}) \approx \nabla_{\boldsymbol{x}} \log p\left(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\right). \tag{76}$$

Thus, DPS replaces the integration over $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$ by a single likelihood evaluation at its conditional mean.

FlowDPS. FlowDPS (Kim et al., 2025) extends the endpoint substitution to flow-based models:

$$\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{y} \mid \boldsymbol{x}) \approx \nabla_{\boldsymbol{x}} \log p\left(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}(\boldsymbol{x})\right). \tag{77}$$

To avoid backpropagating through the full Jacobian of the endpoint predictor,

$$\nabla_{\boldsymbol{x}} \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}) = \left(\frac{\partial \boldsymbol{\mu}_{1|t}}{\partial \boldsymbol{x}}\right)^{\top} \nabla_{\boldsymbol{\mu}_{1|t}} \log p(\boldsymbol{y} \mid \boldsymbol{\mu}_{1|t}), \tag{78}$$

FlowDPS assumes that the clean-data manifold is locally an affine subspace. Under this assumption, the Jacobian simplifies to a schedule-dependent scalar multiplied by a projection matrix onto the tangent space. By further assuming that the likelihood gradient lies in the data-manifold subspace, this projection is omitted. FlowDPS therefore relies on both the clean-endpoint point estimate and an additional geometric proxy to bypass full Jacobian computation.

FlowChef. FlowChef (Patel et al., 2025) also guides a flow trajectory through a predicted clean endpoint, but its approximation differs from that of DPS. Under its original time convention, FlowChef estimates the clean endpoint as

$$\hat{\boldsymbol{x}}_{\text{clean}} = \boldsymbol{x}_t + t\, \boldsymbol{u}_\theta(\boldsymbol{x}_t, t), \tag{79}$$

and defines a measurement loss $\mathcal{L}(A(\hat{\boldsymbol{x}}_{\text{clean}}), \boldsymbol{y})$. By the chain rule,

$$\nabla_{\boldsymbol{x}_t} \mathcal{L} = (\boldsymbol{I} + t\, \boldsymbol{J}_{u_\theta})^{\top} \nabla_{\hat{\boldsymbol{x}}_{\text{clean}}} \mathcal{L}, \tag{80}$$

where $\boldsymbol{J}_{u_\theta} = \frac{\partial \boldsymbol{u}_\theta}{\partial \boldsymbol{x}_t}$ is the spatial Jacobian of the velocity field. FlowChef then skips the Jacobian-dependent factor and directly uses the endpoint-space gradient $\nabla_{\hat{\boldsymbol{x}}_{\text{clean}}} \mathcal{L}$ for guidance. Thus, unlike DPS, FlowChef does not explicitly approximate the intermediate likelihood in Equation (74); instead, it approximates how measurement guidance is backpropagated through the flow dynamics.

DDRM and DDNM. DDRM (Kawar et al., 2022) and DDNM (Wang et al., 2023) circumvent general likelihood approximations by restricting the measurement operator to be linear. DDRM factorizes

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\top} \tag{81}$$

and performs conditional sampling in the corresponding spectral coordinates. The linear-Gaussian measurement problem then decomposes across singular directions, allowing the measurement and diffusion uncertainties to be combined through tractable 1D Gaussian transitions. Its tractability therefore stems from the spectral decomposition of the linear operator rather than from a heuristic point estimate of the clean endpoint.

DDNM instead uses the range-null decomposition. For a noise-free linear measurement $\boldsymbol{y} = A\boldsymbol{x}_1$, it replaces the unconditional endpoint prediction by

$$\tilde{\boldsymbol{\mu}}^{\text{DDNM}}_{1|t,y}(\boldsymbol{x}) = \boldsymbol{A}^{\dagger}\boldsymbol{y} + \left(\boldsymbol{I} - \boldsymbol{A}^{\dagger}\boldsymbol{A}\right) \boldsymbol{\mu}_{1|t}(\boldsymbol{x}), \tag{82}$$

where $\boldsymbol{A}^{\dagger}$ is the pseudoinverse of $\boldsymbol{A}$. This construction exactly satisfies measurement consistency,

$$\boldsymbol{A}\tilde{\boldsymbol{\mu}}^{\text{DDNM}}_{1|t,y} = \boldsymbol{y}, \tag{83}$$

while retaining the model prediction in the null space. However, measurement consistency alone does not imply

$$\tilde{\boldsymbol{\mu}}^{\text{DDNM}}_{1|t,y} = \mathbb{E}[\mathbf{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}]. \tag{84}$$

Thus, DDNM obtains a measurement-consistent endpoint estimate through range-null projection, without explicitly recovering the posterior conditional mean.

DAPS. DAPS (Zhang et al., 2025) differs from the above trajectory-guidance methods by explicitly decoupling consecutive sampling states. Its ideal clean-space update requires sampling from

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}, \mathbf{y} = \boldsymbol{y}) \propto p(\boldsymbol{y} \mid \boldsymbol{x}_1) p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}). \tag{85}$$

Since $p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x})$ remains intractable, practical DAPS adopts the isotropic Gaussian approximation

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) \approx \mathcal{N}\left(\boldsymbol{x}_1; \hat{\boldsymbol{\mu}}_{1|t}(\boldsymbol{x}), r_t^2 \boldsymbol{I}\right), \tag{86}$$

where $\hat{\boldsymbol{\mu}}_{1|t}$ is obtained by solving the unconditional probability-flow ODE and r_t^2 is chosen heuristically. It consequently samples with Langevin dynamics from

$$\tilde{p}_{\mathrm{DAPS}}(\boldsymbol{x}_1 \mid \boldsymbol{x}, \boldsymbol{y}) \propto p(\boldsymbol{y} \mid \boldsymbol{x}_1)\mathcal{N}\left(\boldsymbol{x}_1; \hat{\boldsymbol{\mu}}_{1|t}(\boldsymbol{x}), r_t^2 \boldsymbol{I}\right). \tag{87}$$

While DAPS also discusses a more expensive diffusion-score approximation of the clean conditional, it adopts this Gaussian proxy in practice for efficiency. Although this decoupled scheme reduces error propagation across timesteps, the target distribution of its inner Langevin sampler still depends on the heuristic Gaussian approximation in Equation (86).

Comparison with A^{2}PS. The above approaches achieve tractability through a point estimate, structural assumptions on the measurement operator or flow geometry, or a prescribed heuristic approximation to the clean conditional. In contrast, A^{2}PS directly characterizes the asymptotic behavior of the true conditional density in Equation (74):

$$p(\boldsymbol{x}_1 \mid \mathbf{x}_t = \boldsymbol{x}) = \mathcal{N}\left(\boldsymbol{x}_1; \boldsymbol{\mu}_{1|t}(\boldsymbol{x}), \boldsymbol{\Sigma}_{1|t}(\boldsymbol{x})\right)\left[1 + \mathcal{O}(\rho_t^3)\right], \tag{88}$$

with

$$\boldsymbol{\Sigma}_{1|t} = \frac{a_t^2 \sigma_0^2}{b_t} \frac{\partial \boldsymbol{\mu}_{1|t}(\boldsymbol{x})}{\partial \boldsymbol{x}} = \rho_t^2 \boldsymbol{I} + \mathcal{O}(\rho_t^4). \tag{89}$$

Rather than discarding conditional uncertainty or fixing an isotropic heuristic covariance, A^{2}PS retains an asymptotically rigorous conditional distribution, yielding the asymptotically accurate endpoint posterior mean and posterior score derived in Proposition 3.2 and corollary 3.3.