

A POSTERIOR-DYNAMICS FRAMEWORK FOR IMAGING INVERSE PROBLEMS WITH PRETRAINED DIFFUSION PRIORS

ZHAOQIANG LIU*, TONGYAO PANG†, RUIBING WANG*, AND YANG ZHENG*

Abstract. Pretrained diffusion models represent image distributions through a continuum of progressively smoothed distributions. This multiscale structure organizes generation from global structure to fine detail and supports high-quality, diverse samples. We exploit the same multiscale diffusion prior for linear imaging inverse problems. Rather than using the pretrained model only as a denoiser in an outer iteration, we define a surrogate likelihood whose center is aligned with the clean-image coordinate and whose covariance accounts for residual diffusion uncertainty. This construction defines an explicit surrogate posterior path, from which we derive continuous posterior dynamics. A tunable Langevin component supports target tracking and allows the amount of posterior exploration to be adapted to the application. We prove endpoint consistency and a finite-horizon tracking bound and, in the exact-score setting, first-order weak accuracy. For computation, we derive the Posterior-Dynamics Implicit–Explicit sampler (PD-IMEX), a stable method using one score evaluation per diffusion scale and an implicit data-consistency update. Experiments on deblurring, super-resolution, and inpainting show strong reconstruction quality at 100 score evaluations, coarse-grid stability, and controllable fidelity–diversity behavior.

Key words. diffusion generative priors; imaging inverse problems; posterior dynamics; likelihood continuation; posterior sampling; implicit-explicit methods

MSC codes. 65J22, 65C30, 62C10, 68T07, 94A08

1. Introduction. Many imaging tasks, including deblurring, super-resolution, tomography, and accelerated magnetic resonance imaging, can be modeled as linear inverse problems. We consider measurements of the form

$$(1.1) \quad y = Ax + n, \quad A \in \mathbb{R}^{m \times d}, \quad n \sim \mathcal{N}(0, \sigma_d^2 I),$$

where $x \in \mathbb{R}^d$ is the unknown clean image. Because A is often ill-conditioned or rank deficient, reconstruction depends critically on prior information. Classical models encode smoothness, sparsity, low-rank structure, or nonlocal self-similarity through explicit regularizers [35, 40, 9, 6, 12, 19, 17]. Task-specific restoration networks learn stronger regularities from paired, noisy, or self-supervised data [46, 28, 45, 25, 3, 24]. Their reconstruction rule, however, is typically tied to the training degradation, data protocol, or task, and can require retraining or adaptation when the sensing operator changes.

Deep generative models provide a distributional alternative. They learn the statistics of plausible images from clean data and can be reused with different sensing operators at inference time. This separates the learned image prior from the explicit physical measurement model and avoids training a new reconstruction network for every degradation.

Diffusion and closely related continuous-time generative models have achieved remarkable success in high-resolution image synthesis, producing realistic images with coherent global structure, fine spatial detail, and broad semantic diversity. Stable Diffusion [34] and the flow-matching model FLUX.1 Kontext [4] provide compelling

*School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China (Zhaoqiang Liu: zqliu12@gmail.com; Ruibing Wang: ruibing-wang0818@163.com; Yang Zheng: 202511081645@std.uestc.edu.cn).

†Corresponding author. Yau Mathematical Sciences Center, Shuangqing Complex, Tsinghua University, Beijing 100084, China (typang@tsinghua.edu.cn).

empirical evidence that these models can represent rich, high-dimensional image distributions. In a diffusion model, generation is organized through a continuum of progressively Gaussian-smoothed image distributions rather than through the clean-image distribution alone [37, 18, 38]. We call this multi-level noise-indexed family the *multiscale diffusion prior*: each *diffusion scale* is a density at a different degree of Gaussian smoothing. At coarse diffusion scales, the model organizes global structure and alternative modes; at fine diffusion scales, it progressively resolves texture and local detail. This coarse-to-fine organization underlies the high fidelity and diversity of diffusion-based generation.

For inverse problems, this structure suggests conditioning the whole multiscale diffusion prior rather than using the pretrained network only as an endpoint denoiser. Coarse diffusion scales can organize global image structure while the measurement remains uncertain, and finer scales can incorporate increasingly detailed data consistency. Retaining a distribution at every diffusion scale also allows multiple data-consistent modes in underdetermined problems [41, 44]. Our objective is to exploit this multiscale representation to obtain an efficient posterior solver without sacrificing the diversity supplied by the generative prior.

Accordingly, we formulate reconstruction as Bayesian posterior inference and ask how the endpoint posterior can be extended across the same diffusion scales:

$$(1.2) \quad p(x | y) \propto p_0(x) \exp\left(-\frac{\|Ax - y\|^2}{2\sigma_d^2}\right),$$

where p_0 is represented implicitly by a diffusion model. A diffusion prior learns the score of progressively smoothed data distributions [37, 38, 18]. For a forward process

$$(1.3) \quad dX_t = f(X_t, t) dt + g(t) dW_t,$$

the reverse-time prior dynamics are

$$(1.4) \quad dX_t = [f(X_t, t) - g(t)^2 \nabla \log p_t(X_t)] dt + g(t) d\bar{W}_t, \quad t : T \rightarrow 0,$$

with the score approximated by a neural network trained by denoising score matching [1, 38].

One principled way to modify the reverse prior dynamics so that their endpoint law matches the posterior is to tilt the reference path law by the endpoint likelihood. Let $\mathbb{P}^{\text{Pr}}$ denote the path law induced by (1.4). The resulting likelihood-tilted path law admits the relative-entropy-regularized variational characterization [16]

$$(1.5) \quad \mathbb{Q}_y^* = \arg \min_{\mathbb{Q} \ll \mathbb{P}^{\text{Pr}}} \{\text{KL}(\mathbb{Q} \| \mathbb{P}^{\text{Pr}}) - \mathbb{E}_{\mathbb{Q}}[\log p(y | X_0)]\}, \quad \frac{d\mathbb{Q}_y^*}{d\mathbb{P}^{\text{Pr}}} = \frac{p(y | X_0)}{p(y)}.$$

By Bayes' rule, the endpoint marginal under $\mathbb{Q}_y^*$ is $p_0(\cdot | y)$. This path-law modification can be realized dynamically by a Doob transform [14] with potential

$$h_t(x) = \mathbb{E}_{\mathbb{P}^{\text{Pr}}}[p(y | X_0) | X_t = x] = p_t(y | x).$$

Under the decreasing-time convention in (1.4), the Doob transform contributes the drift $-g(t)^2 \nabla \log h_t$. The exact posterior dynamics are therefore

$$(1.6) \quad dX_t = [f(X_t, t) - g(t)^2 \nabla \log p_t(X_t | y)] dt + g(t) d\bar{W}_t, \quad t : T \rightarrow 0.$$

Bayes' rule then separates the posterior score in this controlled drift as

$$(1.7) \quad \nabla_{x_t} \log p_t(x_t | y) = \nabla_{x_t} \log p_t(x_t) + \nabla_{x_t} \log p_t(y | x_t).$$

The pretrained model supplies the prior score at every diffusion scale, but the intermediate likelihood score is generally unavailable because

$$(1.8) \quad p_t(y | x_t) = \int p(y | x_0) p(x_0 | x_t) dx_0.$$

Thus the central computational difficulty is to condition the multiscale diffusion prior, not to learn another clean-image prior. The path-space control formulation identifies the exact reference dynamics; the remaining task is to construct a tractable surrogate for its unavailable likelihood term.

Existing methods address the unavailable intermediate likelihood through discrete guidance corrections, operator-aware projections, unrolled or plug-and-play reconstruction schemes, and more costly particle or inner Monte Carlo procedures [10, 11, 39, 23, 49, 44, 7, 43, 48]. Recent provable methods obtain posterior correctness through either singular-direction DDIM updates or an outer probability transport equipped with inner Langevin posterior-score estimation [20, 8]. These approaches have different algorithmic structures. Our goal is to preserve the multiscale diffusion prior as the central modeling object and derive an explicit posterior path, continuous dynamics, and numerical solver from one construction.

We work in clean-image coordinates, where scale alignment and likelihood continuation define a normalized surrogate posterior along the Gaussian smoothing path of the diffusion prior. Its transport defect motivates a tunable Langevin component for tracking and exploration. From the same continuous model, we derive the Posterior-Dynamics Implicit–Explicit sampler (PD-IMEX), which uses one score evaluation per scale and implicit linear data consistency.

Our main contribution is a multiscale posterior-dynamics framework that turns a fixed pretrained diffusion model into an efficient and tunable solver for linear imaging inverse problems. Specifically:

- **Multiscale surrogate posterior construction.** We use the complete multiscale diffusion prior rather than treating the pretrained network only as a denoising module in an outer reconstruction procedure. Deterministic scale alignment places the surrogate likelihood center in the clean-image coordinate, while likelihood continuation accounts for the leading projected covariance of residual diffusion uncertainty. Together they define a normalized surrogate posterior at every diffusion scale and recover the desired Bayesian posterior at the endpoint.
- **Posterior dynamics with guarantees and tunable exploration.** We derive continuous dynamics consisting of base multiscale transport and a freely chosen Langevin component. For the moving surrogate posterior, the Langevin intensity supports target tracking and allows the same framework to range from nearly deterministic reconstruction to stronger posterior exploration. We characterize the path mismatch and establish endpoint consistency and a finite-horizon tracking bound.
- **An efficient and stable numerical realization.** We derive PD-IMEX, a one-score-evaluation-per-scale discretization that treats the linear likelihood dynamics implicitly and adds the Langevin innovation after the data-consistency solve. For the exact prior score, we establish first-order weak

accuracy. Experiments show strong reconstruction quality at 100 score evaluations, coarse-grid stability, and control over the fidelity–diversity trade-off.

2. Related work and positioning.

Likelihood guidance and operator-aware solvers. Training-free diffusion inverse solvers commonly modify the unconditional reverse process using an approximation of the unavailable intermediate likelihood score at each diffusion scale. DPS evaluates measurement discrepancy at a Tweedie denoised estimate. Methods based on manifold constraints or pseudoinverses incorporate local geometry or approximate conditional covariance information [10, 11, 39]. TMPD goes beyond a point approximation by using higher-order Tweedie moments to construct a statistically motivated approximation of the conditional distribution [5]. For structured linear problems, SNIPS and DDRM exploit singular-vector decompositions to separate measurement-informed and prior-dominated directions, DDNM decomposes range and null-space components, and DiffPIR alternates diffusion denoising with explicit data consistency [22, 23, 42, 49].

Posterior-oriented and annealed inference. A second line of work places greater emphasis on posterior inference, typically at additional computational cost. Variational approaches such as RED-Diff optimize tractable objectives induced by the diffusion prior, while MGPS uses intermediate-state guidance within a variational posterior-sampling construction [29, 30]. Filtering and sequential Monte Carlo approaches, including auxiliary-variable methods, approximate the conditional distribution through particles or inner Monte Carlo procedures rather than a single reverse-guidance update [15, 7, 43]. PnP-DM treats the pretrained diffusion model as a plug-and-play prior and reduces Bayesian inversion to iterative sampling of Gaussian denoising posteriors [44]. DAPS instead constructs a decoupled noise-annealing process and performs inner sampling at each diffusion scale, allowing larger moves and improved exploration relative to a conventional reverse trajectory [48]. These approaches improve posterior fidelity or exploration through additional inference iterations, whereas our focus is a single-loop dynamics using one score evaluation at each discretized diffusion scale.

Provable diffusion posterior sampling. Recent work has begun to establish rigorous posterior guarantees for diffusion inverse solvers. For noisy linear inverse problems, a provable DDIM-type sampler treats the singular directions of the measurement operator separately: each component follows the diffusion prior while the measurement signal-to-noise ratio is low and switches to a calibrated measurement predictor once the observation becomes sufficiently informative. This gives lightweight coordinate-wise updates with posterior consistency [20]. For more general Bayesian inverse problems, an outer probability transport estimates the time-dependent posterior score through an inner Langevin Monte Carlo procedure with a warm start. Its nonasymptotic analysis controls initialization and sampling errors in this nested procedure, including the effects of problem conditioning [8]. These results show that rigorous posterior sampling can be obtained either by exploiting detailed structure of the observation model or by investing additional computation in posterior-score estimation.

Position of this work. Our distinction lies in using the multiscale diffusion prior as the central object from which posterior modeling and computation are derived. Many practical methods introduce data consistency through discrete guidance, unrolled reconstruction, or plug-and-play iterations, while recent provable methods rely on operator-specific coordinate updates or inner Monte Carlo score estimation.

Our construction is complementary. We define a normalized surrogate posterior at every diffusion scale, quantify its mismatch from the base multiscale transport, and derive continuous transport–Langevin dynamics and the corresponding PD-IMEX discretization from the same posterior path. This gives a single-loop, one-score-per-scale solver while retaining a tunable Langevin intensity for posterior tracking and exploration.

3. Posterior dynamics from a multiscale diffusion prior. The multiscale diffusion prior introduced in Section 1 becomes, in clean-image coordinates, a Gaussian smoothing path. We first condition its endpoint exactly to identify a reference posterior path and a family of transport–Langevin dynamics. We then construct a tractable surrogate posterior path across the same diffusion scales, analyze its tracking error, and derive the PD-IMEX discretization.

3.1. Exact posterior path and transport–Langevin family. We begin with the ideal setting in which the posterior score is available. We identify the unknown signal with the diffusion endpoint $X_0 \in \mathbb{R}^d$ and assume throughout the total-variation analysis that $\sigma_d > 0$. The variance-preserving (VP) forward stochastic differential equation (SDE) is

$$(3.1) \quad dX_t = -\frac{1}{2}\beta(t)X_t dt + \sqrt{\beta(t)} dW_t, \quad X_0 \sim p_0,$$

where $\beta(t) > 0$. Its marginal representation is

$$(3.2) \quad X_t = \sqrt{\alpha_t}X_0 + \sigma_t\varepsilon, \quad \alpha_t = \exp\left(-\int_0^t \beta(s) ds\right), \quad \sigma_t^2 = 1 - \alpha_t.$$

The diffusion state contains both deterministic attenuation and additive Gaussian noise. For the exact posterior path, it is convenient to remove only the known attenuation and use the remaining noise-to-signal ratio as time. We therefore introduce the clean state and log-SNR time

$$(3.3) \quad Z_\lambda := \frac{X_t}{\sqrt{\alpha_t}} = X_0 + \tau_\lambda\varepsilon, \quad \tau_\lambda = e^{-\lambda}, \quad q_\lambda := \tau_\lambda^2 = e^{-2\lambda}, \quad \lambda = \frac{1}{2} \log \frac{\alpha_t}{1 - \alpha_t}.$$

This is an exact change of variables, not yet an approximation of the likelihood. In this coordinate, every marginal is an additive Gaussian perturbation of the same clean variable, with uncertainty q_λ . The log-SNR clock makes the reverse direction monotone and gives $\partial_\lambda q_\lambda = -2q_\lambda$. Thus increasing λ follows the reverse generative direction from coarse to fine diffusion scales. Let P_q denote convolution with $\mathcal{N}(0, qI)$. The unconditional clean-coordinate density is

$$(3.4) \quad \pi_\lambda = P_{q_\lambda} p_0.$$

The family $\{\pi_\lambda\}_\lambda$ is the clean-coordinate representation of the multiscale diffusion prior. Conditioning the initial law on y gives the exact posterior path

$$(3.5) \quad \pi_\lambda^y = P_{q_\lambda} p_0(\cdot | y).$$

Accordingly, $\{\pi_\lambda^y\}_\lambda$ is the exact posterior path across the same diffusion scales. For every finite λ , Gaussian smoothing makes π_λ^y strictly positive and smooth, even when the endpoint posterior is only a Borel probability measure.

Under the change of variables in (3.3), the exact posterior dynamics (1.6) become

$$dZ_\lambda = 2q_\lambda \nabla \log \pi_\lambda^y(Z_\lambda) d\lambda + \sqrt{2q_\lambda} dW_\lambda.$$

The marginal path (3.5), however, does not uniquely determine the pathwise dynamics. Let $a_\lambda \geq 0$ denote the Langevin intensity, which controls the strength of a matched Langevin score drift and Brownian diffusion. The following proposition gives the resulting family with the same marginals; the dynamics above are recovered when $a_\lambda = q_\lambda$.

PROPOSITION 3.1 (Exact posterior family with arbitrary Langevin intensity). *Let $a_\lambda \geq 0$ be locally bounded and suppose the SDE and its Fokker–Planck equation are well posed on every compact interval $[\lambda_0, \Lambda]$. Consider*

$$(3.6) \quad dZ_\lambda = (q_\lambda + a_\lambda) \nabla \log \pi_\lambda^y(Z_\lambda) d\lambda + \sqrt{2a_\lambda} dW_\lambda, \quad Z_{\lambda_0} \sim \pi_{\lambda_0}^y.$$

Then $Z_\lambda \sim \pi_\lambda^y$ for all $\lambda \geq \lambda_0$. Moreover, $\pi_\lambda^y \Rightarrow p_0(\cdot | y)$ as $\lambda \rightarrow \infty$.

The proof is given in [Appendix A](#). The key identity is

$$(3.7) \quad \partial_\lambda \pi_\lambda^y = -q_\lambda \Delta \pi_\lambda^y.$$

When the candidate density in the Fokker–Planck equation of (3.6) is set equal to π_λ^y , the a_λ -dependent drift and diffusion cancel exactly, leaving (3.7).

Remark 3.2 (Relative Langevin intensity). Choosing $a_\lambda = \zeta q_\lambda$, with $\zeta \geq 0$, parameterizes the Langevin component relative to the multiscale transport coefficient q_λ . It gives

$$dZ_\lambda = (1 + \zeta) q_\lambda \nabla \log \pi_\lambda^y(Z_\lambda) d\lambda + \sqrt{2\zeta q_\lambda} dW_\lambda.$$

Thus $\zeta = 0$ is probability-flow transport and $\zeta = 1$ is the usual reverse SDE in the clean coordinate. Values $\zeta > 1$ are also admissible, and a general a_λ permits a time-dependent Langevin-intensity schedule. In the ideal model, varying ζ changes pathwise exploration without changing the prescribed posterior marginals. For the moving surrogate posterior introduced below, the same Langevin component additionally supplies entropy dissipation and therefore acts as a tracking mechanism.

3.2. Surrogate posterior path across diffusion scales.

Scale-consistent likelihood center. The exact family is not directly implementable because an unconditional diffusion model supplies the prior score but not the intermediate likelihood in (1.8). Write the clean likelihood, up to its normalizing constant, as

$$\ell(x_0) := \exp \left[-\frac{\|Ax_0 - y\|^2}{2\sigma_d^2} \right].$$

The exact quantity is the conditional average

$$p_t(y | x_t) \propto \mathbb{E}[\ell(X_0) | X_t = x_t], \quad X_0 = \frac{x_t}{\sqrt{\alpha_t}} - \frac{\sigma_t}{\sqrt{\alpha_t}} \varepsilon.$$

This decomposition separates two effects of the diffusion scale on the intermediate likelihood. The deterministic attenuation determines where the clean-image forward model should be evaluated, while the additive term represents residual uncertainty about the clean image. Evaluating the latter exactly requires the conditional law of the additive correction, including its state-dependent mean, covariance, and higher

moments. We instead approximate the two effects in stages. First, we retain only the deterministic scaling of the diffusion channel and drop the additive correction:

$$X_0 \approx \frac{x_t}{\sqrt{\alpha_t}}, \quad p_t(y | x_t) \approx \ell\left(\frac{x_t}{\sqrt{\alpha_t}}\right).$$

Here scale consistency refers specifically to the likelihood center: the known attenuation is removed before applying the clean-image forward model. This zeroth-order approximation does not yet account for uncertainty in the discarded additive term. It gives

$$(3.8) \quad L_t^{\text{sc}}(x_t) := \exp\left[-\frac{1}{2\sigma_d^2} \left\| A \frac{x_t}{\sqrt{\alpha_t}} - y \right\|^2\right].$$

The chain rule gives

$$(3.9) \quad \nabla_{x_t} \log L_t^{\text{sc}}(x_t) = -\frac{1}{\sqrt{\alpha_t}\sigma_d^2} A^\top \left(A \frac{x_t}{\sqrt{\alpha_t}} - y \right).$$

The approximation becomes the clean likelihood at the data endpoint, but at a finite diffusion scale it remains only a surrogate for the conditional average.

In the z coordinate, the baseline surrogate posterior is

$$(3.10) \quad \tilde{\pi}_\lambda^{\text{sc}}(z | y) \propto \pi_\lambda(z) \exp\left[-\frac{\|Az - y\|^2}{2\sigma_d^2}\right].$$

Equation (3.10) reveals a simple prior-smoothing structure of the clean-coordinate formulation. The physical likelihood is the same as that at the data endpoint, while all diffusion-scale dependence is confined to

$$\pi_\lambda = P_{q_\lambda} p_0,$$

the Gaussian-smoothed prior density. Related prior-side smoothing ideas appear in proximal optimization and proximal MCMC, where a possibly nonsmooth prior energy is regularized through its Moreau–Yosida envelope and the associated proximity operator [32, 33]. The analogy is structural rather than exact: proximal methods use Moreau smoothing of the prior energy, whereas diffusion applies Gaussian smoothing to the prior density. The z -coordinate makes this separation explicit by removing the deterministic attenuation of the diffusion state and expressing the baseline surrogate as a fixed physical likelihood multiplied by a multiscale, Gaussian-smoothed prior.

Likelihood continuation for residual uncertainty. The baseline surrogate in (3.10) therefore already contains a prior-side smoothing path. The likelihood continuation introduced next serves a different purpose: it accounts for the residual uncertainty of the clean image conditional on the current diffusion state. Away from the endpoint,

$$Z_\lambda = X_0 + \tau_\lambda \varepsilon$$

remains uncertain around X_0 , so evaluating the endpoint likelihood directly at Z_λ can be overconfident. We therefore retain a leading covariance approximation of the discarded diffusion perturbation. This motivates the effective covariance

$$(3.11) \quad \Gamma_\lambda := \sigma_d^2 I + \eta_\lambda A A^\top, \quad \eta_\lambda \geq 0, \quad \eta_\lambda \rightarrow 0 \quad (\lambda \rightarrow \infty).$$

We call η_λ the likelihood-continuation schedule. The nominal choice $\eta_\lambda = q_\lambda$ matches the covariance of the projected isotropic perturbation. This covariance correction remains a controlled surrogate rather than the exact state-dependent conditional likelihood. More generally, the likelihood continuation accounts for the diffusion-scale uncertainty remaining in Z_λ and introduces measurement information progressively as the prior moves from coarse to fine diffusion scales. The original likelihood is recovered as $\eta_\lambda \rightarrow 0$. Define

$$(3.12) \quad L_\lambda(z) := \exp \left[-\frac{1}{2} (Az - y)^\top \Gamma_\lambda^{-1} (Az - y) \right]$$

and the normalized surrogate posterior at diffusion scale λ

$$(3.13) \quad \bar{\pi}_\lambda(z | y) := \frac{1}{Z_\lambda} \pi_\lambda(z) L_\lambda(z), \quad Z_\lambda := \int \pi_\lambda(z) L_\lambda(z) dz.$$

Its score is

$$(3.14) \quad \bar{s}_\lambda(z; y) = s_\lambda^{\text{pr}}(z) + s_\lambda^{\text{lik}}(z) = \nabla \log \pi_\lambda(z) - H_\lambda z + b_\lambda,$$

where

$$(3.15) \quad H_\lambda := A^\top \Gamma_\lambda^{-1} A, \quad b_\lambda := A^\top \Gamma_\lambda^{-1} y.$$

Let $t = t(\lambda)$ denote the inverse log-SNR parametrization and write $\alpha_\lambda := \alpha_{t(\lambda)}$. We use $\widehat{\varepsilon}_\theta(\cdot, \lambda)$ as shorthand for the pretrained network evaluated at the corresponding diffusion-time label $t(\lambda)$. In practice,

$$(3.16) \quad s_{\theta, \lambda}^{\text{pr}}(z) := -e^\lambda \widehat{\varepsilon}_\theta(\sqrt{\alpha_\lambda} z, \lambda)$$

approximates $\nabla \log \pi_\lambda(z)$.

The likelihood continuation is useful only if it returns to the desired Bayesian model as the diffusion noise vanishes. The next proposition verifies this endpoint property separately for densities and scores. The distinction is important: density convergence identifies the limiting law, whereas score convergence alone does not guarantee convergence of normalized probability measures.

PROPOSITION 3.3 (Endpoint and score consistency). *Assume that p_0 is absolutely continuous with an $L^1(\mathbb{R}^d)$ probability density and $\sigma_d > 0$. If $\eta_\lambda \rightarrow 0$, then*

$$(3.17) \quad \|\bar{\pi}_\lambda(\cdot | y) - p(\cdot | y)\|_{\text{TV}} \rightarrow 0.$$

If, in addition, p_0 is strictly positive and C^1 , and

$$\nabla \log(P_{q_\lambda} p_0) \rightarrow \nabla \log p_0 \quad \text{locally uniformly,}$$

then $\bar{s}_\lambda(\cdot; y)$ converges locally uniformly to the clean-posterior score

$$\nabla \log p_0(z) - \sigma_d^{-2} A^\top (Az - y).$$

The L^1 density assumption is essential for total-variation convergence. If the endpoint prior is singular, Gaussian mollifications still converge weakly but need not converge in total variation. The proof of [Proposition 3.3](#) is in [Appendix B](#). Thus the surrogate posterior path has the correct endpoint, but this statement alone says nothing about whether a dynamics can track it.

Transport defect and posterior dynamics. The remaining issue comes from the way the surrogate was constructed. The exact posterior path in (3.5) is a Gaussian smoothing of the joint endpoint density and therefore satisfies the backward heat equation. By contrast, (3.13) multiplies a smoothed prior by a separately continued likelihood. This product has the correct endpoint but need not follow the same transport equation. Consequently, simply substituting $\bar{s}_\lambda$ into the probability-flow drift does not in general produce marginals $\bar{\pi}_\lambda$. We quantify this continuous-time mismatch before deciding how much correction is needed.

DEFINITION 3.4 (Transport defect). *For a positive smooth surrogate posterior path $\{\bar{\pi}_\lambda\}$, define*

$$(3.18) \quad r_\lambda(z) := \frac{\partial_\lambda \bar{\pi}_\lambda(z) + q_\lambda \Delta \bar{\pi}_\lambda(z)}{\bar{\pi}_\lambda(z)} = \partial_\lambda \log \bar{\pi}_\lambda(z) + q_\lambda \left(\Delta \log \bar{\pi}_\lambda(z) + \|\bar{s}_\lambda(z)\|^2 \right).$$

The transport defect measures the mismatch between the surrogate posterior path and the base multiscale transport inherited from the diffusion prior. Indeed, the base transport velocity is

$$v_\lambda(z) = q_\lambda \nabla \log \bar{\pi}_\lambda(z),$$

and (3.18) is equivalently

$$(3.19) \quad r_\lambda \bar{\pi}_\lambda = \partial_\lambda \bar{\pi}_\lambda + \nabla \cdot (\bar{\pi}_\lambda v_\lambda).$$

Thus r_λ is a normalized Fokker–Planck or continuity-equation residual. It is a signed local source term rather than a nonnegative norm. Under sufficient decay,

$$(3.20) \quad \int r_\lambda d\bar{\pi}_\lambda = 0.$$

Positive and negative values describe where the prescribed surrogate posterior gains or loses mass faster than the base multiscale transport can provide.

For an observable F , define its centered version by $\mathcal{C}_{\bar{\pi}_\lambda}[F] := F - \mathbb{E}_{\bar{\pi}_\lambda}[F]$. The following identity applies directly to both the nominal and practical continuation schedules.

PROPOSITION 3.5 (Transport defect for a differentiable likelihood continuation). *Assume that η_λ is differentiable and that the required differentiation and integration-by-parts operations are justified. Then*

$$(3.21) \quad r_\lambda = 2q_\lambda \mathcal{C}_{\bar{\pi}_\lambda} [s_\lambda^{\text{pr}} \cdot s_\lambda^{\text{lik}}] + \left(q_\lambda + \frac{1}{2} \dot{\eta}_\lambda \right) \mathcal{C}_{\bar{\pi}_\lambda} [\|s_\lambda^{\text{lik}}\|^2].$$

In particular, $\eta_\lambda = q_\lambda$ implies $\dot{\eta}_\lambda = -2q_\lambda$, so the second term vanishes.

The proof is in [Appendix C](#). The first term is the prior–likelihood interaction created because the heat flow does not preserve products. The second records the mismatch between the likelihood continuation clock and the nominal variance q_λ . For the practical schedule $\eta_\lambda = q_\lambda^2/(q_\lambda + \tau_c^2)$,

$$(3.22) \quad q_\lambda + \frac{1}{2} \dot{\eta}_\lambda = \frac{q_\lambda \tau_c^4}{(q_\lambda + \tau_c^2)^2},$$

so continuation adds an explicit likelihood self-interaction to the nominal prior–likelihood defect.

The defect explains why the base multiscale transport alone can be fragile for the surrogate posterior path. We therefore retain this transport and add a Langevin component at every λ . When the diffusion scale is frozen, this component preserves $\bar{\pi}_\lambda$; while the posterior path moves, it supplies the dissipative mechanism in the tracking estimate and permits exploration of posterior ambiguity. Whether this dissipation dominates the transport defect depends on the functional-inequality and defect-control conditions below. Let $a_\lambda \geq 0$ denote the Langevin intensity. The continuous surrogate model is

$$(3.23) \quad \boxed{dZ_\lambda = (q_\lambda + a_\lambda) \bar{s}_\lambda(Z_\lambda; y) d\lambda + \sqrt{2a_\lambda} dW_\lambda.}$$

It may be decomposed algebraically as

$$(3.24) \quad dZ_\lambda = \underbrace{q_\lambda \bar{s}_\lambda(Z_\lambda; y) d\lambda}_{\text{base multiscale transport}} + \underbrace{\left[a_\lambda \bar{s}_\lambda(Z_\lambda; y) d\lambda + \sqrt{2a_\lambda} dW_\lambda \right]}_{\text{Langevin component}}.$$

The second bracket has $\bar{\pi}_\lambda$ as its invariant law when λ is frozen. Thus the base multiscale transport advances the diffusion scale, while the Langevin component supports tracking and exploration. Increasing a_λ strengthens both effects, while the definition of the surrogate target $\bar{\pi}_\lambda$ itself remains unchanged.

The continuous model is now fully specified. Its convergence requires two distinct facts. First, the surrogate posterior must approach the clean posterior, which was established in [Proposition 3.3](#). Second, the sampler law must track the moving surrogate posterior despite the defect r_λ . To address the latter, let μ_λ denote the law of (3.23) and define

$$(3.25) \quad H(\lambda) := \text{KL}(\mu_\lambda \| \bar{\pi}_\lambda), \quad I(\mu_\lambda \| \bar{\pi}_\lambda) := \int \|\nabla \log(\mu_\lambda / \bar{\pi}_\lambda)\|^2 d\mu_\lambda.$$

The following theorem shows that the Langevin component produces relative-Fisher-information dissipation, while the transport defect is the forcing term generated by motion of the surrogate posterior path. Combining this tracking estimate with end-point consistency yields convergence to the posterior under the stated conditions.

THEOREM 3.6 (Entropy identity and finite-horizon posterior tracking). *Assume that the coefficients, surrogate posterior path, and sampler densities are sufficiently regular for the SDE, Fokker–Planck equation, differentiation under the integral, and the integrations by parts below. Then*

$$(3.26) \quad H'(\lambda) = -a_\lambda I(\mu_\lambda \| \bar{\pi}_\lambda) - \int r_\lambda d\mu_\lambda.$$

Suppose further that, for each λ , $\bar{\pi}_\lambda$ satisfies the log-Sobolev inequality

$$(3.27) \quad \text{Ent}_{\bar{\pi}_\lambda}(f^2) \leq \frac{2}{\rho_\lambda} \int \|\nabla f\|^2 d\bar{\pi}_\lambda, \quad \rho_\lambda > 0,$$

and that for every law ν encountered in the evolution with finite relative entropy,

$$(3.28) \quad - \int r_\lambda d\nu \leq c_\lambda \text{KL}(\nu \| \bar{\pi}_\lambda) + \varepsilon_\lambda, \quad c_\lambda, \varepsilon_\lambda \geq 0.$$

On a fixed interval $[\lambda_0, \Lambda]$, set

$$(3.29) \quad k_\lambda := 2a_\lambda \rho_\lambda - c_\lambda.$$

If $H(\lambda_0) < \infty$, then for every $\lambda \leq \Lambda$,

$$(3.30) \quad H(\lambda) \leq e^{-\int_{\lambda_0}^{\lambda} k_u du} H(\lambda_0) + \int_{\lambda_0}^{\lambda} e^{-\int_s^{\lambda} k_u du} \varepsilon_s ds.$$

The proof is in [Appendix D](#). The first term in (3.30) propagates the mismatch between the practical high-noise initialization and the initial surrogate posterior; the second accumulates transport-defect forcing after Langevin dissipation. The log-Sobolev and entropy-variational tools are standard [2, 16]. In particular, the entropy variational inequality verifies (3.28) whenever, for some $\theta_\lambda > 0$,

$$(3.31) \quad c_\lambda = \theta_\lambda^{-1}, \quad \varepsilon_\lambda = \theta_\lambda^{-1} \log \int e^{-\theta_\lambda r_\lambda} d\bar{\pi}_\lambda < \infty.$$

Combined with [Proposition 3.3](#), the finite-horizon estimate yields $\|\mu_\lambda - p(\cdot | y)\|_{\text{TV}} \rightarrow 0$ whenever $\int_{\lambda_0}^{\infty} k_u du = \infty$ and the weighted defect term in (3.30) vanishes.

3.3. PD-IMEX discretization. PD-IMEX discretizes the surrogate posterior dynamics over a grid of diffusion scales. At each diffusion scale, the nonlinear prior prediction is evaluated once, the linear likelihood drift is treated implicitly, and the Langevin innovation is added after the data-consistency solve. Using the score decomposition (3.14), the continuous dynamics (3.23) can be written as

$$(3.32) \quad dZ_\lambda = (q_\lambda + a_\lambda)[s_\lambda^{\text{pr}}(Z_\lambda) - H_\lambda Z_\lambda + b_\lambda] d\lambda + \sqrt{2a_\lambda} dW_\lambda.$$

The prior score is nonlinear and supplied by the network, while the likelihood drift is linear with coefficients H_λ, b_λ determined by the likelihood-continuation covariance Γ_λ . A fully implicit treatment would require repeated network evaluations, whereas an explicit likelihood step can be unstable when H_λ is stiff. This structure motivates an implicit-explicit (IMEX) step that evaluates the network once and treats the linear likelihood implicitly.

Let

$$(3.33) \quad \lambda_0 < \lambda_1 < \dots < \lambda_K, \quad h_k := \lambda_{k+1} - \lambda_k,$$

and write $q_k := q_{\lambda_k}$ and $\tau_k := \tau_{\lambda_k}$. By Tweedie's identity, $\mathbb{E}[X_0 | Z_\lambda = z] = z + q_\lambda s_\lambda^{\text{pr}}(z)$. Thus the noise-prediction network also gives the clean prediction

$$(3.34) \quad \hat{\varepsilon}_k := \hat{\varepsilon}_\theta(\sqrt{\alpha_{\lambda_k}} z_k, \lambda_k), \quad \hat{x}_{0,k} := z_k - \tau_k \hat{\varepsilon}_k.$$

Although score, noise, and clean predictions are algebraically equivalent at λ_k , their frozen schemes have different numerical behavior. Direct clean-data prediction remains effective in high-dimensional pixel spaces where prediction of noised quantities degrades [27]. Moreover, recent error analysis shows that posterior-mean freezing admits a discretization bound governed by the metric-entropy dimension, whereas score freezing retains ambient-dimension dependence [31]. We therefore freeze the clean prediction, leaving the known Gaussian state dependence unfrozen:

$$(3.35) \quad s_\lambda^{\text{pr}}(z) \approx \frac{\hat{x}_{0,k} - z}{q_\lambda}.$$

Set $\chi_\lambda := q_\lambda + a_\lambda$ and $\kappa_\lambda := \chi_\lambda / q_\lambda$. On $[\lambda_k, \lambda_{k+1}]$, substituting (3.35) into the dynamics (3.32) gives

$$(3.36) \quad dZ_\lambda = [\kappa_\lambda(\hat{x}_{0,k} - Z_\lambda) - \chi_\lambda H_\lambda Z_\lambda + \chi_\lambda b_\lambda] d\lambda + \sqrt{2a_\lambda} dW_\lambda.$$

Equation (3.36) keeps the continuation dependence of H_λ, b_λ explicit through (3.11)–(3.15).

We integrate the mean-reverting prior part of (3.36) exactly, freeze the likelihood coefficients at the right endpoint, and apply backward Euler only to that linear likelihood term. The required step coefficients are

$$(3.37) \quad \varrho_k := \exp\left(-\int_{\lambda_k}^{\lambda_{k+1}} \kappa_s \, ds\right),$$

$$(3.38) \quad \Omega_k := \int_{\lambda_k}^{\lambda_{k+1}} \chi_s \, ds, \quad V_k := 2 \int_{\lambda_k}^{\lambda_{k+1}} a_s \exp\left(-2 \int_s^{\lambda_{k+1}} \kappa_u \, du\right) \, ds.$$

Here ϱ_k is the exact mean contraction of the frozen mean-reverting component, Ω_k is the accumulated likelihood weight, and V_k is the conditional variance generated by the Langevin diffusion. With $H_{k+1} := H_{\lambda_{k+1}}$ and $b_{k+1} := b_{\lambda_{k+1}}$, the PD-IMEX step is

$$(3.39) \quad \boxed{\begin{aligned} (I + \Omega_k H_{k+1}) z_{k+1}^{\det} &= \widehat{x}_{0,k} + \varrho_k (z_k - \widehat{x}_{0,k}) + \Omega_k b_{k+1}, \\ z_{k+1} &= z_{k+1}^{\det} + \sqrt{V_k} \xi_k, \quad \xi_k \sim \mathcal{N}(0, I). \end{aligned}}$$

Algorithm 3.1 PD-IMEX posterior sampler

Require: Measurement y , operator A , measurement-noise level σ_d , score network $\widehat{\varepsilon}_\theta$, grid $\{\lambda_k\}_{k=0}^K$, likelihood-continuation schedule η_λ , Langevin-intensity schedule a_λ

Ensure: Approximate posterior sample z_K

- 1: Draw $x_{\lambda_0} \sim \mathcal{N}(0, I)$ and convert it to the clean coordinate: $z_0 \leftarrow x_{\lambda_0} / \sqrt{\alpha_{\lambda_0}}$
 - 2: **for** $k = 0, \dots, K-1$ **do**
 - 3: Form $\Gamma_{\lambda_{k+1}}$, H_{k+1} , and b_{k+1} from (3.11) and (3.15)
 - 4: Evaluate $\widehat{\varepsilon}_k$ and $\widehat{x}_{0,k}$ from (3.34)
 - 5: Compute ϱ_k , Ω_k , and V_k from (3.37)–(3.38)
 - 6: Solve the deterministic IMEX equation in (3.39) for $z_{k+1}^{\det}$
 - 7: Draw $\xi_k \sim \mathcal{N}(0, I)$ and set $z_{k+1} \leftarrow z_{k+1}^{\det} + \sqrt{V_k} \xi_k$
 - 8: **end for**
 - 9: Return z_K
-

The implicit solve in (3.39) stabilizes the deterministic likelihood drift. The Gaussian term $\sqrt{V_k} \xi_k$ is the Langevin innovation after integration through the frozen prior mean-reverting subflow. Adding it after the likelihood solve prevents additional filtering by the likelihood resolvent. Classical analyses of stiff stochastic integrators identify the same loss of fluctuations when an under-resolved relaxation is handled by an L-stable implicit update [26]. If the innovation were instead placed inside the backward-Euler solve, a likelihood eigenmode of curvature ν would retain only $(1 + \Omega_k \nu)^{-2}$ of its innovation variance. This post-solve ordering preserves the innovation covariance prescribed by the update, while the following proposition establishes first-order consistency with the continuous posterior dynamics. It is not claimed to integrate the full stiff frozen covariance exactly.

PROPOSITION 3.7 (Post-solve Langevin innovation and local consistency). *Assume H_{k+1} is symmetric positive semidefinite and let $R_k = (I + \Omega_k H_{k+1})^{-1}$. Conditional on z_k , the post-solve update (3.39) has covariance $V_k I$. If the same Langevin*

innovation is placed inside the implicit likelihood solve, its covariance is $V_k R_k^2$; hence post-solve placement avoids artificial damping by the likelihood resolvent.

If a_λ , q_λ , and the coefficients are C^1 on a fixed finite interval, then as $h_k \rightarrow 0$,

$$(3.40) \quad \varrho_k = 1 - \kappa_{\lambda_k} h_k + O(h_k^2), \quad \Omega_k = \chi_{\lambda_k} h_k + O(h_k^2), \quad V_k = 2a_{\lambda_k} h_k + O(h_k^2).$$

Consequently, when $\hat{x}_{0,k} = z_k + q_k s_{\lambda_k}^{\text{pr}}(z_k)$, the conditional mean and covariance of (3.39) match those of (3.23) to first order.

The proof is in [Appendix E](#). This statement concerns the innovation covariance of the chosen update decomposition; it does not claim exact integration of the full stiff frozen covariance. The first-order expansion fixes the finite interval and its coefficients, so it is not uniform in a simultaneous stiff limit where $\Omega_k \|H_{k+1}\|$ remains order one under refinement.

Having established retention of the post-solve Langevin innovation and local consistency, we next quantify the global weak error relative to the continuous sampler in the exact-score setting. Let Q_k denote the exact transition of (3.23) from λ_k to λ_{k+1} , and let $\hat{Q}_k$ be the one-step kernel of (3.39) using the surrogate posterior score in (3.14).

THEOREM 3.8 (First-order weak error of PD-IMEX). *Fix $\Lambda > \lambda_0$ and let $\lambda_K = \Lambda$. Suppose PD-IMEX uses the exact prior score, so that $\hat{x}_{0,k} = z_k + q_k s_{\lambda_k}^{\text{pr}}(z_k)$, and write $B_\lambda(z) = (q_\lambda + a_\lambda) \bar{s}_\lambda(z; y)$. Assume on $[\lambda_0, \Lambda]$ that B_λ , a_λ , H_λ , and b_λ are C^1 in λ , and that the spatial derivatives of B_λ through order four have polynomial growth uniformly in λ . Assume also that the exact process and the numerical chains are nonexplosive and satisfy the polynomial moment bounds used below, uniformly on the finite interval and over sufficiently fine grids. For a test function $\phi \in C_b^4(\mathbb{R}^d)$, suppose that the associated backward Kolmogorov solution has spatial derivatives through order four with polynomial-growth bounds uniformly on $[\lambda_0, \Lambda]$. Initialize both processes from the same law and let*

$$h := \max_{0 \leq k < K} h_k.$$

Then there is a constant $C_{\phi, \Lambda}$, independent of the grid, such that

$$(3.41) \quad \left| \mathbb{E} \phi(z_K^{\text{alg}}) - \int \phi d\mu_\Lambda \right| \leq C_{\phi, \Lambda} h.$$

These are finite-horizon weak-approximation assumptions. Stating the required backward regularity explicitly avoids imposing stronger global conditions on the endpoint prior than are needed for the argument.

The proof in [Appendix E](#) establishes the local expansion

$$(3.42) \quad \hat{Q}_k f(z) = f(z) + h_k \mathcal{L}_{\lambda_k} f(z) + O((1 + \|z\|^m) h_k^2),$$

where

$$(3.43) \quad \mathcal{L}_\lambda f = (q_\lambda + a_\lambda) \bar{s}_\lambda(\cdot; y) \cdot \nabla f + a_\lambda \Delta f.$$

The exact transition has the same first-order expansion. A telescoping argument with backward Kolmogorov regularity then gives the global result.

Combining the finite-horizon weak error with continuous tracking and endpoint consistency yields the following decomposition for the oracle-score chain.

COROLLARY 3.9 (Finite-horizon weak-error decomposition for oracle-score PD-IMEX). *Under the assumptions of [Theorems 3.6](#) and [3.8](#), for every $\phi \in C_b^4(\mathbb{R}^d)$,*

$$(3.44) \quad \left| \mathbb{E} \phi(z_K^{\text{alg}}) - \int \phi \, dp(\cdot \mid y) \right| \leq C_{\phi, \lambda_K} h + 2 \|\phi\|_{\infty} \sqrt{\frac{1}{2} H(\lambda_K)} \\ + 2 \|\phi\|_{\infty} \|\bar{\pi}_{\lambda_K} - p(\cdot \mid y)\|_{\text{TV}}.$$

The three terms in [\(3.44\)](#) are, respectively, time-discretization error, sampler-to-surrogate tracking error, and surrogate endpoint error.

4. Experiments. The experiments evaluate the three layers of the framework: construction of the surrogate posterior path from the multiscale diffusion prior, the role of the Langevin intensity in tracking and exploration, and the numerical realization of the dynamics.

4.1. Experimental protocol. We evaluate on 100 images from Flickr-Faces-HQ (FFHQ) 256×256 [\[21\]](#) and 100 images from ImageNet 256×256 [\[36\]](#). For FFHQ we use the pretrained prior from DPS [\[10\]](#); for ImageNet we use the unconditional prior of Dhariwal and Nichol [\[13\]](#). We consider bicubic $\times 4$ super-resolution, Gaussian deblurring with a 61×61 kernel of standard deviation 3.0, and motion deblurring with a 61×61 kernel of intensity 0.5. Independent Gaussian measurement noise with $\sigma_d = 0.05$ is added in every main task. We report peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), and learned perceptual image patch similarity (LPIPS) [\[47\]](#).

We count each score-network call as one network function evaluation (NFE). The reported 100-NFE outputs use 99 PD-IMEX transitions followed by one terminal denoising evaluation, which serves as a finite-noise reconstruction readout. The endpoints are the coarsest and finest trained diffusion scales, $\lambda_0 = -5.05884$ and $\lambda_K = 4.60512$, and the grid is uniform in log-SNR up to rounding to the nearest distinct training index. We parameterize the Langevin intensity as $a_{\lambda} = \zeta q_{\lambda}$, where ζ is the relative Langevin intensity. The likelihood covariance is

$$(4.1) \quad \Gamma(\tau) = \sigma_d^2 I + \eta(\tau) A A^{\top}, \quad \eta(\tau) = \frac{\tau^4}{\tau^2 + \tau_c^2}, \quad \tau_c = 6.$$

This choice follows the uncertainty interpretation in [\(3.11\)](#). Since $z = X_0 + \tau \varepsilon$, the perturbation observed through A has covariance $\tau^2 A A^{\top}$, giving the nominal likelihood covariance $\sigma_d^2 I + \tau^2 A A^{\top}$. We temper it by the smooth gate $\eta(\tau)/\tau^2 = \tau^2/(\tau^2 + \tau_c^2)$. Thus $\eta(\tau) \sim \tau^2$ at coarse diffusion scales, where projected diffusion uncertainty should be retained, whereas $\eta(\tau) \sim \tau^4/\tau_c^2$ near the endpoint, where the continuation is removed more rapidly to expose the original likelihood. The crossover satisfies $\eta(\tau_c) = \tau_c^2/2$. We fix $\tau_c = 6$ for all main tasks; the resulting schedule is positive, smooth, and endpoint-consistent. We use $\zeta = 0.1$ for FFHQ and $\zeta = 0.01$ for ImageNet. MGPS and DAPS use multiple score calls within each outer level, so both levels and actual NFEs are reported.

Every likelihood solve is noniterative. Deblurring is diagonalized by the two-dimensional FFT, bicubic super-resolution uses the explicit singular-vector basis of the separable degradation, and inpainting uses a diagonal closed form. The controlled ablations share the same images, measurement-noise realizations, operator seed, and posterior-sampling seeds.

At 100 score-network evaluations, PD-IMEX provides the strongest overall distortion performance in the comparison and remains competitive in perceptual quality,

TABLE 1

Reconstruction results with $\sigma_d = 0.05$, averaged over 100 images per dataset. Computational cost is reported as outer levels/actual score-network evaluations. Higher PSNR and SSIM and lower LPIPS are better.

Method	Levels/ NFEs↓	Gaussian deblur			Motion deblur			Super-resolution (×4)		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
<i>FFHQ</i>										
DPS [10]	1000/1000	25.41	.708	<u>.144</u>	23.69	.658	.178	24.21	.673	.194
DiffPIR [49]	100/100	28.24	.774	.184	27.09	.704	.194	27.67	.772	<u>.118</u>
DDRM [23]	100/100	<u>29.12</u>	<u>.837</u>	.132	<u>29.23</u>	<u>.846</u>	.089	<u>29.30</u>	<u>.843</u>	.141
RED-Diff [29]	100/100	28.46	<u>.705</u>	.242	26.62	.768	.235	27.72	.715	.361
MGPS [30]	100/814	27.88	.792	.149	26.82	.769	.147	27.55	.783	.115
DAPS [48]	100/500	28.71	.792	.155	28.39	.795	.142	28.66	.764	.142
Ours (PD-IMEX)	100/100	29.52	.847	.208	30.57	.869	<u>.139</u>	29.49	.847	.201
<i>ImageNet</i>										
DPS [10]	1000/1000	21.81	.526	.314	21.15	.503	.352	21.58	.507	.404
DiffPIR [49]	100/100	24.18	.597	.493	23.16	.514	.422	23.71	.577	.338
DDRM [23]	100/100	<u>25.59</u>	<u>.714</u>	.336	<u>26.61</u>	<u>.761</u>	.189	25.64	.720	.321
RED-Diff [29]	100/100	24.52	.661	.491	23.76	.635	.429	24.46	.609	.537
MGPS [30]	100/674	24.68	.657	.397	24.14	.642	.352	24.77	.667	.305
DAPS [48]	100/500	24.98	.661	<u>.349</u>	25.33	.685	.303	25.08	.669	<u>.319</u>
Ours (PD-IMEX)	100/100	25.62	.716	.399	27.59	.779	<u>.197</u>	<u>25.35</u>	<u>.712</u>	.390




FIG. 1. Representative FFHQ reconstructions for the methods in Table 1. Columns share the same measurement within each row.

with the expected fidelity–perception trade-off on several tasks. It ranks first in ten of the twelve PSNR and SSIM comparisons and second in the remaining two, including against methods using five to ten times as many score evaluations. DDRM and MGPS retain lower LPIPS on several tasks, whereas PD-IMEX is second in motion-deblurring LPIPS on both datasets.

4.2. Ablation of multiscale posterior construction. We first test how the measurement model is coupled to the multiscale diffusion prior. The naive baseline applies time-independent likelihood guidance, the second configuration corrects the likelihood center through scale consistency, and the full configuration additionally continues its covariance. All other stochastic and numerical settings are fixed. The naive multiplier is tuned separately for each task, giving $\gamma^* = 1$ on FFHQ and $\gamma^* = 0.3$ on ImageNet.

Table 2 separates the two components of the surrogate likelihood. Scale consistency corrects its center by undoing the deterministic attenuation of the diffusion

TABLE 2

Ablation of the multiscale posterior construction over 100 images at a fixed 100-NFE budget. The naive baseline uses its best time-independent scalar multiplier. The second configuration aligns the likelihood center with the clean-image coordinate, and the third additionally continues its covariance to account for residual diffusion uncertainty. All Langevin and numerical settings are fixed.

Configuration	FFHQ Gaussian deblur			ImageNet motion deblur		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Naive guidance (best scalar)	19.228	.67088	.33503	23.799	<u>.70211</u>	.34586
Scale-consistent likelihood	<u>29.502</u>	.84730	.20597	<u>24.729</u>	.56936	<u>.20600</u>
Scale-consistent + continuation	29.523	<u>.84723</u>	<u>.20753</u>	27.585	.77864	.19684




FIG. 2. *Effect of the relative Langevin intensity ζ on reconstruction fidelity and perceptual quality for two noisy deblurring tasks. Each panel reports PSNR on the left axis and LPIPS on the right axis over 100 images. Only ζ is varied; the continuation, numerical grid, score network, measurements, and random seeds are fixed. Dashed vertical lines indicate the values used in Table 1.*

state. A scalar guidance weight can compensate for part of a center mismatch on some tasks but cannot replace this scale-dependent alignment uniformly. Likelihood continuation then changes the effective covariance, accounting for residual diffusion uncertainty at coarse scales. Its effect is small on FFHQ Gaussian deblurring but substantial on ImageNet motion deblurring.

4.3. Langevin intensity: tracking and exploration. Figure 2 shows that the practical effect of the relative Langevin intensity is nonmonotone and task dependent. On FFHQ Gaussian deblurring, a moderate value improves distortion accuracy, while a larger value can improve LPIPS at the expense of PSNR. ImageNet motion deblurring favors a smaller value, and an excessively large value degrades both metrics. Thus ζ controls a finite-step tracking–exploration trade-off rather than serving as a universally beneficial choice of Langevin intensity.

Figure 3 complements the aggregate metrics with a multimodal inpainting example. At $\zeta = 0$, the finite-step deterministic implementation contracts different initializations toward similar, relatively smooth completions. In contrast, the selected larger value produces coherent but visibly different facial completions, illustrating how the relative Langevin intensity controls sample variability in an underdetermined problem.

4.4. Numerical realization: coarse-grid stability and innovation placement.




FIG. 3. Effect of the relative Langevin intensity on sample diversity for centered-box inpainting. The two rows use $\zeta = 0$ and $\zeta = 24$, respectively, and the five central columns use independent sampling seeds while keeping the measurement and operator fixed.

TABLE 3

NFE refinement of the explicit and implicit likelihood updates. Entries are mean PSNR, SSIM, and LPIPS over 100 images. Within each NFE, the better result between the two updates is bold; fail denotes numerical failure on all 100 images. All configurations use the same surrogate posterior path, Langevin-intensity schedule, score network, images, measurements, and random seeds; only the log-SNR grid and the treatment of the linear likelihood drift differ.

Task	NFE	Explicit			Implicit		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
FFHQ Gaussian deblur	100		<i>fail</i>		29.523	.8472	.2075
	200	28.698	.7891	.2311	29.505	.8470	.2066
	500	29.499	.8472	.2058	29.489	.8469	.2064
	1000	29.494	.8470	.2063	29.486	.8468	.2066
ImageNet motion deblur	100	5.434	.1344	.9360	27.585	.7786	.1968
	200	18.342	.2843	.5728	27.643	.7841	.2217
	500	27.661	.7847	.2284	27.638	.7838	.2357
	1000	27.637	.7838	.2375	27.627	.7833	.2403

Coarse-grid stability. We next examine the two numerical choices in PD-IMEX by refining the log-SNR grid. The explicit and implicit variants have the same surrogate posterior path and Langevin-intensity schedule; they differ only in the treatment of the linear likelihood drift.

The explicit update is unusable on the 100-NFE grid: it fails numerically on all FFHQ images and attains only 5.43 dB PSNR, .1344 SSIM, and .9360 LPIPS on ImageNet. Refinement restores the same limiting behavior as the implicit update, but at an operator-dependent rate. On FFHQ the explicit method remains less accurate at 200 NFEs and agrees with the implicit method across all three metrics by 500 NFEs; on ImageNet it remains under-resolved at 200 NFEs and recovers only at 500 NFEs. Over the same refinement, the maximum stability indicator $\max_k \Omega_k \lambda_{\max}(H_{k+1})$ decreases from 7.30 to 1.06 at 500 NFEs and .50 at 1000 NFEs on FFHQ; the corresponding values on ImageNet are 6.70, .97, and .46. The implicit update is stable throughout. Thus implicit treatment follows the same posterior dynamics on a substantially coarser grid and hence requires fewer score evaluations.

Innovation placement. Noiseless centered-box inpainting is treated separately as a qualitative stress test. Removing a 128×128 square leaves a large null space, so many semantically different faces can agree exactly with the observed pixels. Pixelwise PSNR and SSIM against one hidden ground truth would penalize plausible posterior samples merely for choosing a different eye shape, expression, or facial detail. We therefore use this task only to inspect visual coherence and sample variability. All methods use the same FFHQ prior, centered mask, and noiseless observed pixels. For PD-IMEX, we use a quadratic grid, initialize the observed region from the appropriately noised measurement, and apply no hard projection. The 100-NFE configuration uses $\tau_c = 24$ and $\zeta = 24$.

This task also makes the placement of the Langevin innovation especially relevant. To display the distinction explicitly, write

$$d_k := \hat{x}_{0,k} + \varrho_k(z_k - \hat{x}_{0,k}) + \Omega_k b_{k+1}, \quad R_k := (I + \Omega_k H_{k+1})^{-1}.$$

Using the same Gaussian draw ξ_k , the two variants are

$$(4.2) \quad \begin{aligned} z_{k+1}^{\text{inside}} &= R_k(d_k + \sqrt{V_k} \xi_k), \\ z_{k+1}^{\text{post}} &= R_k d_k + \sqrt{V_k} \xi_k, \quad \xi_k \sim \mathcal{N}(0, I). \end{aligned}$$

Thus, inside-solve placement filters the innovation through the likelihood resolvent, whereas post-solve placement applies the resolvent only to the deterministic update. Along a likelihood eigenmode of curvature ν , the local relaxation time is approximately $(\kappa_\lambda + \chi_\lambda \nu)^{-1}$. Near the endpoint, the product $\Omega_k \nu$ can remain order one or larger, so this fast scale is not resolved by the grid. The inside-solve variant then retains only the fraction $(1 + \Omega_k \nu)^{-2}$ of the innovation variance. The post-solve variant in (4.2) instead stabilizes the deterministic likelihood drift without filtering the Langevin innovation.

To isolate this placement effect at finite NFE, the amount of residual refresh must also remain comparable. Over step k , the fresh-innovation fraction is $1 - \exp(-2\zeta \Delta \lambda_k)$. Thus, holding ζ fixed while refining the grid weakens the refresh per step and confounds innovation placement with a different practical stochastic regime. We use $\zeta = 24, 48, 120, 240$ at 100, 200, 500, 1000 NFEs, respectively, which approximately preserves the typical per-step refresh strength of the 100-NFE configuration.

Within each NFE, the inside- and post-solve variants share the initialization and every Gaussian draw; only the placement of the innovation changes. Figure 4 shows that post-solve innovation gives a coherent completion throughout the tested NFE range, whereas filtering the same innovation through the inside-solve resolvent leaves the missing region unresolved. The experiment therefore supports post-solve placement when the finite-step refresh strength is matched.

Figure 5 places this controlled study in a broader inverse-problem comparison. Since many images are compatible with the same masked observation, we do not rank the methods by pixelwise proximity to the single hidden image. Instead, the figure compares preservation of the observed region and the generation of coherent, semantically plausible content in the missing region.




FIG. 4. *Langevin innovation-placement study for noiseless centered-box inpainting. From 100 to 1000 NFEs, we use $\zeta = 24, 48, 120, 240$, respectively, to keep the typical per-step residual refresh comparable. Within each column, the inside- and post-solve variants use the same initialization, deterministic implicit update, and Gaussian draws; only the innovation placement changes.*




FIG. 5. *Qualitative comparison for noiseless centered-box inpainting. From top to bottom, the rows show images 38, 15, and 80. All methods use the same FFHQ model, centered mask, and noiseless observed pixels. PD-IMEX uses $\zeta = 24$, post-solve Langevin innovations, and 100 NFEs. Because the conditional problem is highly nonunique, the ground truth is shown as one compatible latent image rather than as the unique desired completion.*

5. Conclusion. We developed posterior dynamics that use the full multiscale diffusion prior for linear imaging inverse problems. The key idea is to condition the diffusion prior across diffusion scales, rather than using the pretrained network only as a denoising module inside an outer reconstruction procedure. Aligning the likelihood center with the clean-image coordinate and continuing its covariance for residual diffusion uncertainty define an explicit surrogate posterior path. The resulting dynamics combine base multiscale transport with a tunable Langevin component, providing a common mechanism for posterior tracking and task-dependent exploration. Endpoint consistency and a finite-horizon tracking estimate give analytical support for this construction.

PD-IMEX provides an efficient and stable realization of the dynamics. It uses one score evaluation per diffusion scale, treats the linear likelihood drift implicitly, and places the Langevin innovation after the data-consistency solve. Experiments show strong reconstruction quality at 100 score evaluations, demonstrate that implicit integration preserves coarse-grid efficiency, and show how the relative Langevin intensity adapts the fidelity–diversity behavior to different inverse problems. Together, these results provide a direct route from a pretrained multiscale diffusion prior to a theo-

retically supported and flexible posterior solver.

Appendix A. Proof of the exact posterior family.

Proof of Proposition 3.1. Let $q_\lambda = e^{-2\lambda}$. Since $\pi_\lambda^y = P_{q_\lambda} p_0(\cdot | y)$ and the Gaussian heat semigroup satisfies

$$\partial_q P_q \nu = \frac{1}{2} \Delta P_q \nu,$$

we have $q'_\lambda = -2q_\lambda$ and therefore

$$(A.1) \quad \partial_\lambda \pi_\lambda^y = -q_\lambda \Delta \pi_\lambda^y.$$

Let m_λ be the density of (3.6). Its Fokker–Planck equation is

$$(A.2) \quad \partial_\lambda m_\lambda = -\nabla \cdot ((q_\lambda + a_\lambda) m_\lambda \nabla \log \pi_\lambda^y) + a_\lambda \Delta m_\lambda.$$

Substituting $m_\lambda = \pi_\lambda^y$ and using $\pi_\lambda^y \nabla \log \pi_\lambda^y = \nabla \pi_\lambda^y$ gives

$$-\nabla \cdot ((q_\lambda + a_\lambda) \nabla \pi_\lambda^y) + a_\lambda \Delta \pi_\lambda^y = -q_\lambda \Delta \pi_\lambda^y = \partial_\lambda \pi_\lambda^y.$$

Thus π_λ^y solves (A.2). Since the initial law is $\pi_{\lambda_0}^y$, uniqueness of the Fokker–Planck solution implies $m_\lambda = \pi_\lambda^y$ on every compact interval. Finally, $q_\lambda \rightarrow 0$, and continuity of the Gaussian semigroup at zero gives $P_{q_\lambda} p_0(\cdot | y) \Rightarrow p_0(\cdot | y)$. $\square$

Appendix B. Endpoint and score consistency.

Proof of Proposition 3.3. Because $p_0 \in L^1(\mathbb{R}^d)$ and the Gaussian kernels form an approximate identity,

$$(B.1) \quad \|\pi_\lambda - p_0\|_{L^1} \rightarrow 0.$$

Let

$$L_\infty(z) = \exp \left[-\frac{\|Az - y\|^2}{2\sigma_d^2} \right].$$

Since $\Gamma_\lambda^{-1} \rightarrow \sigma_d^{-2} I$, we have $L_\lambda(z) \rightarrow L_\infty(z)$ pointwise, and $0 < L_\lambda, L_\infty \leq 1$. Write $f_\lambda = \pi_\lambda L_\lambda$ and $f_\infty = p_0 L_\infty$. Then

$$(B.2) \quad \begin{aligned} \|f_\lambda - f_\infty\|_{L^1} &\leq \|\pi_\lambda - p_0\|_{L^1} + \int \pi_\lambda |L_\lambda - L_\infty| dz \\ &\leq 2 \|\pi_\lambda - p_0\|_{L^1} + \int p_0 |L_\lambda - L_\infty| dz \rightarrow 0, \end{aligned}$$

where the final integral converges by dominated convergence. Consequently,

$$\mathcal{Z}_\lambda = \int f_\lambda dz \rightarrow \mathcal{Z}_\infty := \int f_\infty dz > 0.$$

For all sufficiently large λ , $\mathcal{Z}_\lambda \geq \mathcal{Z}_\infty/2$, and

$$\left\| \frac{f_\lambda}{\mathcal{Z}_\lambda} - \frac{f_\infty}{\mathcal{Z}_\infty} \right\|_{L^1} \leq \frac{1}{\mathcal{Z}_\lambda} \|f_\lambda - f_\infty\|_{L^1} + \left| \frac{1}{\mathcal{Z}_\lambda} - \frac{1}{\mathcal{Z}_\infty} \right| \|f_\infty\|_{L^1} \rightarrow 0.$$

Since total variation is one half of the L^1 distance between densities, this proves (3.17).

For the score statement, Bayes' rule gives

$$\nabla \log p(z | y) = \nabla \log p_0(z) - \sigma_d^{-2} A^\top (Az - y).$$

Subtracting this identity from (3.14) yields

$$(B.3) \quad \bar{s}_\lambda(z; y) - \nabla \log p(z | y) = \nabla \log \pi_\lambda(z) - \nabla \log p_0(z) - A^\top (\Gamma_\lambda^{-1} - \sigma_d^{-2} I)(Az - y).$$

The first term converges locally uniformly by assumption. On every compact set, the second term converges uniformly because $\Gamma_\lambda^{-1} \rightarrow \sigma_d^{-2} I$ and $Az - y$ is bounded there. $\square$

Appendix C. Transport defect for the continued likelihood.

Proof of Proposition 3.5. Write $f_\lambda = \pi_\lambda$ and $u_\lambda = f_\lambda L_\lambda$, so that $\bar{\pi}_\lambda = u_\lambda / \int u_\lambda$. Since $\partial_\lambda f_\lambda = -q_\lambda \Delta f_\lambda$, centering removes the derivative of the normalizer and gives

$$(C.1) \quad r_\lambda = \mathcal{C}_{\bar{\pi}_\lambda} \left[\frac{\partial_\lambda u_\lambda + q_\lambda \Delta u_\lambda}{u_\lambda} \right].$$

The product rule contributes $2q_\lambda s_\lambda^{\text{pr}} \cdot s_\lambda^{\text{lik}}$. Moreover, direct differentiation of (3.12) yields

$$\partial_\lambda \log L_\lambda = \frac{1}{2} \dot{\eta}_\lambda \|s_\lambda^{\text{lik}}\|^2, \quad \Delta \log L_\lambda = -\text{tr}(H_\lambda).$$

Consequently,

$$\frac{\partial_\lambda L_\lambda + q_\lambda \Delta L_\lambda}{L_\lambda} = -q_\lambda \text{tr}(H_\lambda) + \left(q_\lambda + \frac{1}{2} \dot{\eta}_\lambda \right) \|s_\lambda^{\text{lik}}\|^2.$$

Substitution into (C.1) proves (3.21), because the spatially constant trace term vanishes under centering. $\square$

Appendix D. Entropy identity and finite-horizon tracking.

Proof of Theorem 3.6. For notational brevity, write $q = q_\lambda$, $a = a_\lambda$, $\pi = \bar{\pi}_\lambda$, $\mu = \mu_\lambda$, and $g = \mu/\pi$. The Fokker–Planck equation of (3.23) is

$$(D.1) \quad \begin{aligned} \partial_\lambda \mu &= -\nabla \cdot ((q + a)\mu \nabla \log \pi) + a \Delta \mu \\ &= -q \nabla \cdot (\mu \nabla \log \pi) + a \nabla \cdot (\mu \nabla \log g). \end{aligned}$$

The second equality follows from $\nabla \cdot (\mu \nabla \log g) = \Delta \mu - \nabla \cdot (\mu \nabla \log \pi)$. By the definition of the defect,

$$(D.2) \quad \partial_\lambda \pi = -q \Delta \pi + r_\lambda \pi.$$

Differentiate $H(\lambda) = \int \mu \log g \, dz$. Since $\int \partial_\lambda \mu \, dz = 0$,

$$(D.3) \quad H'(\lambda) = \int \partial_\lambda \mu \log g \, dz - \int \mu \partial_\lambda \log \pi \, dz.$$

The Langevin contribution in (D.1) is

$$(D.4) \quad a \int \nabla \cdot (\mu \nabla \log g) \log g \, dz = -a \int \|\nabla \log g\|^2 \, d\mu.$$

The base transport contribution is

$$(D.5) \quad \begin{aligned} q \int \mu \nabla \log \pi \cdot \nabla \log g \, dz &= q \int \nabla \pi \cdot \nabla g \, dz \\ &= -q \int g \Delta \pi \, dz = -q \int \frac{\Delta \pi}{\pi} \, d\mu. \end{aligned}$$

Meanwhile, (D.2) gives

$$(D.6) \quad - \int \mu \partial_\lambda \log \pi \, dz = q \int \frac{\Delta \pi}{\pi} \, d\mu - \int r_\lambda \, d\mu.$$

The first terms in (D.5) and (D.6) cancel, yielding the exact identity (3.26).

Apply the log-Sobolev inequality (3.27) to $f = \sqrt{g}$. Since $\int g \, d\pi = 1$,

$$H(\lambda) = \text{Ent}_\pi(g), \quad \int \|\nabla \sqrt{g}\|^2 \, d\pi = \frac{1}{4} I(\mu|\pi),$$

and hence

$$(D.7) \quad I(\mu|\pi) \geq 2\rho_\lambda H(\lambda).$$

Combining (3.26), (3.28), and (D.7),

$$H'(\lambda) \leq -(2a_\lambda \rho_\lambda - c_\lambda) H(\lambda) + \varepsilon_\lambda = -k_\lambda H(\lambda) + \varepsilon_\lambda.$$

The integrating-factor form of Gronwall's inequality gives (3.30). If the cumulative contraction diverges and its weighted defect forcing vanishes, then $H(\lambda) \rightarrow 0$. Pinsker's inequality, the triangle inequality, and Proposition 3.3 then yield convergence to the posterior in total variation.

For completeness, the entropy variational inequality

$$\int F \, d\nu \leq \text{KL}(\nu|\pi) + \log \int e^F \, d\pi$$

with $F = -\theta_\lambda r_\lambda$ yields (3.31) pointwise in λ . □

Appendix E. Post-solve Langevin innovation and weak error.

Proof of Proposition 3.7. Write $R_k = (I + \Omega_k H_{k+1})^{-1}$ and $\tilde{z}_{k+1} = \hat{x}_{0,k} + \varrho_k(z_k - \hat{x}_{0,k}) + \Omega_k b_{k+1}$. Then $z_{k+1} = R_k \tilde{z}_{k+1} + \sqrt{V_k} \xi_k$, so its conditional covariance is exactly $V_k I$. If the innovation is instead included in the implicit solve, the update becomes $R_k(\tilde{z}_{k+1} + \sqrt{V_k} \xi_k)$ and has conditional covariance $V_k R_k^2$, since R_k is symmetric. This proves the covariance comparison.

On a fixed finite interval, Taylor expansion of the scalar integrals in (3.37)–(3.38) gives, uniformly in k ,

$$(E.1) \quad \varrho_k = 1 - \kappa_k h_k + O(h_k^2), \quad \Omega_k = \chi_k h_k + O(h_k^2), \quad V_k = 2a_k h_k + O(h_k^2),$$

because the exponential factor defining V_k is $1 + O(h_k)$ uniformly over the step. This proves (3.40).

Tweedie's identity gives $\hat{x}_{0,k} - z = q_k s_k^{\text{pr}}(z)$. Smoothness in λ and (E.1) yield

$$\tilde{z}_{k+1} = z + \chi_k h_k (s_k^{\text{pr}}(z) + b_k) + O((1 + \|z\|) h_k^2), \quad R_k = I - \chi_k h_k H_k + O(h_k^2).$$

Consequently,

$$(E.2) \quad \mathbb{E}[z_{k+1} - z \mid z_k = z] = (q_k + a_k)\bar{s}_{\lambda_k}(z; y)h_k + O((1 + \|z\|)h_k^2),$$

$$(E.3) \quad \text{Cov}(z_{k+1} \mid z_k = z) = 2a_k h_k I + O(h_k^2),$$

where the first line uses $\bar{s}_{\lambda_k}(z; y) = s_k^{\text{PT}}(z) - H_k z + b_k$. These are the first conditional moments of (3.23) to first order. $\square$

Proof of Theorem 3.8. Set

$$B_\lambda(z) = (q_\lambda + a_\lambda)\bar{s}_\lambda(z; y), \quad \mathcal{L}_\lambda f = B_\lambda \cdot \nabla f + a_\lambda \Delta f.$$

Fix k and condition on $z_k^{\text{alg}} = z$. The affine Gaussian update can be written as

$$\Delta_k := z_{k+1}^{\text{alg}} - z = m_k(z) + \sqrt{V_k}\xi_k, \quad \xi_k \sim \mathcal{N}(0, I),$$

where (E.2)–(E.3) give

$$(E.4) \quad m_k(z) = B_{\lambda_k}(z)h_k + O((1 + \|z\|)h_k^2), \quad V_k = 2a_{\lambda_k}h_k + O(h_k^2).$$

The raw second moment, rather than only the covariance, is needed below. Since $\mathbb{E}[\Delta_k \Delta_k^\top] = V_k I + m_k m_k^\top$,

$$(E.5) \quad \mathbb{E}[\Delta_k \Delta_k^\top] = 2a_{\lambda_k}h_k I + O((1 + \|z\|^2)h_k^2).$$

The centered Gaussian third moment vanishes. Expanding around m_k and using (E.4) therefore gives, for some integer m independent of the grid,

$$(E.6) \quad \|\mathbb{E}[\Delta_k^{\otimes 3}]\| \leq C(1 + \|z\|^m)h_k^2, \quad \mathbb{E}\|\Delta_k\|^4 \leq C(1 + \|z\|^m)h_k^2.$$

The first estimate is a signed tensor-moment bound, whereas the second is the absolute fourth-moment bound required for the Taylor remainder.

For a function f whose spatial derivatives through order four have the stated polynomial-growth bounds, a third-degree Taylor expansion with a fourth-order remainder gives

$$\begin{aligned} f(z + \Delta_k) &= f(z) + Df(z)[\Delta_k] + \frac{1}{2}D^2f(z)[\Delta_k, \Delta_k] \\ &\quad + \frac{1}{6}D^3f(z)[\Delta_k, \Delta_k, \Delta_k] + R_4(z, \Delta_k). \end{aligned}$$

The remainder satisfies

$$|R_4(z, \Delta_k)| \leq C(1 + \|z\|^m + \|\Delta_k\|^m) \|\Delta_k\|^4.$$

Gaussian moment estimates and (E.4) therefore bound its conditional expectation by $C(1 + \|z\|^{m'})h_k^2$ for some m' . Taking conditional expectations in the Taylor expansion and applying (E.4)–(E.6) yields

$$(E.7) \quad \widehat{Q}_k f(z) = f(z) + h_k \mathcal{L}_{\lambda_k} f(z) + O((1 + \|z\|^m)h_k^2).$$

We next derive the corresponding expansion for the exact transition. Set $s = \lambda_k$ and $t = \lambda_{k+1}$, and write $\mathbb{E}_{s,z}$ for expectation conditional on $Z_s = z$. Dynkin's formula gives

$$Q_k f(z) = f(z) + \int_s^t \mathbb{E}_{s,z}[\mathcal{L}_r f(Z_r)] dr.$$

Applying the time-dependent Dynkin formula to $g(r, x) = \mathcal{L}_r f(x)$ gives

$$\mathbb{E}_{s,z}[\mathcal{L}_r f(Z_r)] = \mathcal{L}_s f(z) + \int_s^r \mathbb{E}_{s,z} [((\partial_u + \mathcal{L}_u)\mathcal{L}_u f)(Z_u)] du.$$

Substitution into the preceding identity yields

$$\begin{aligned} Q_k f(z) &= f(z) + h_k \mathcal{L}_s f(z) \\ &\quad + \int_s^t \int_s^r \mathbb{E}_{s,z} [((\partial_u + \mathcal{L}_u)\mathcal{L}_u f)(Z_u)] du dr. \end{aligned}$$

The coefficient regularity implies that $((\partial_u + \mathcal{L}_u)\mathcal{L}_u f)(x)$ has polynomial growth. The exact moment bounds and the area $h_k^2/2$ of the integration region then give

$$(E.8) \quad Q_k f(z) = f(z) + h_k \mathcal{L}_{\lambda_k} f(z) + O((1 + \|z\|^m) h_k^2).$$

Comparing (E.7) and (E.8) gives

$$(E.9) \quad |(\widehat{Q}_k - Q_k)f(z)| \leq C_f(1 + \|z\|^m) h_k^2.$$

Finally, define the backward functions

$$u_k(z) := \mathbb{E}[\phi(Z_\Lambda) \mid Z_{\lambda_k} = z].$$

They satisfy $u_k = Q_k u_{k+1}$ and, by assumption, have spatial derivatives through order four with polynomial-growth bounds uniform in k . Since the exact and numerical processes have the same initial law, telescoping gives

$$\mathbb{E}\phi(z_K^{\text{alg}}) - \mathbb{E}\phi(Z_\Lambda) = \sum_{k=0}^{K-1} \mathbb{E} \left[(\widehat{Q}_k - Q_k) u_{k+1}(z_k^{\text{alg}}) \right].$$

Applying (E.9), the numerical moment bounds, and

$$\sum_{k=0}^{K-1} h_k^2 \leq h \sum_{k=0}^{K-1} h_k = h(\Lambda - \lambda_0)$$

proves (3.41). $\square$

Acknowledgments and declarations. The authors used OpenAI Codex to assist with language editing, L^AT_EX preparation, code development, and experimental workflow organization. The authors assume responsibility for all content.

REFERENCES

- [1] B. D. Anderson, Reverse-time diffusion equation models, *Stochastic Process. Appl.*, 12 (1982), pp. 313–326.
- [2] D. Bakry, I. Gentil, and M. Ledoux, *Analysis and Geometry of Markov Diffusion Operators*, Grundlehren der Mathematischen Wissenschaften 348, Springer, Cham, 2014.
- [3] J. Batson and L. Royer, Noise2Self: Blind denoising by self-supervision, in *Proceedings of ICML*, 2019, pp. 524–533.
- [4] Black Forest Labs, S. Batifol, A. Blattmann, et al., FLUX.1 Kontext: Flow matching for in-context image generation and editing in latent space, arXiv:2506.15742, 2025.
- [5] B. Boys, J. Pidstrigach, M. Girolami, S. Reich, A. Mosca, and O. D. Akyildiz, Tweedie moment projected diffusions for inverse problems, *Trans. Mach. Learn. Res.*, 2024.

- [6] A. Buades, B. Coll, and J.-M. Morel, A non-local algorithm for image denoising, in *Proceedings of CVPR*, 2005, pp. 60–65.
- [7] G. Cardoso, Y. Janati El Idrissi, S. Le Corff, and E. Moulines, Monte Carlo guided denoising diffusion models for Bayesian linear inverse problems, in *International Conference on Learning Representations*, 2024.
- [8] J. Chang, C. Duan, Y. Jiao, R. Li, J. Z. Yang, and C. Yuan, Provable diffusion posterior sampling for Bayesian inversion, arXiv:2512.08022, 2025.
- [9] T. Chan, A. Marquina, and P. Mulet, High-order total variation-based image restoration, *SIAM J. Sci. Comput.*, 22 (2000), pp. 503–516.
- [10] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, and J. C. Ye, Diffusion posterior sampling for general noisy inverse problems, in *International Conference on Learning Representations*, 2023.
- [11] H. Chung, B. Sim, D. Ryu, and J. C. Ye, Improving diffusion models for inverse problems using manifold constraints, in *Advances in Neural Information Processing Systems*, 35 (2022), pp. 25683–25696.
- [12] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, Image denoising by sparse 3-D transform-domain collaborative filtering, *IEEE Trans. Image Process.*, 16 (2007), pp. 2080–2095.
- [13] P. Dhariwal and A. Nichol, Diffusion models beat GANs on image synthesis, in *Advances in Neural Information Processing Systems*, 34 (2021), pp. 8780–8794.
- [14] J. L. Doob, *Classical Potential Theory and Its Probabilistic Counterpart*, Grundlehren der Mathematischen Wissenschaften 262, Springer, New York, 1984, <https://doi.org/10.1007/978-1-4612-5208-5>.
- [15] Z. Dou and Y. Song, Diffusion posterior sampling for linear inverse problem solving: A filtering perspective, in *International Conference on Learning Representations*, 2024.
- [16] P. Dupuis and R. S. Ellis, *A Weak Convergence Approach to the Theory of Large Deviations*, Wiley Series in Probability and Statistics, John Wiley & Sons, New York, 1997.
- [17] S. Gu, L. Zhang, W. Zuo, and X. Feng, Weighted nuclear norm minimization with application to image denoising, in *Proceedings of CVPR*, 2014, pp. 2862–2869.
- [18] J. Ho, A. Jain, and P. Abbeel, Denoising diffusion probabilistic models, in *Advances in Neural Information Processing Systems*, 33 (2020), pp. 6840–6851.
- [19] H. Ji, C. Liu, Z. Shen, and Y. Xu, Robust video denoising using low-rank matrix completion, in *Proceedings of CVPR*, 2010, pp. 1791–1798.
- [20] Y. Jiao, N. Li, C. Cai, Y. Chen, and G. Li, Provable diffusion-based posterior sampling for linear inverse problems via DDIM, arXiv:2607.19333, 2026.
- [21] T. Karras, S. Laine, and T. Aila, A style-based generator architecture for generative adversarial networks, in *Proceedings of CVPR*, 2019, pp. 4401–4410.
- [22] B. Kavar, G. Vaksman, and M. Elad, SNIPS: Solving noisy inverse problems stochastically, in *Advances in Neural Information Processing Systems*, 34 (2021), pp. 21757–21769.
- [23] B. Kavar, M. Elad, S. Ermon, and J. Song, Denoising diffusion restoration models, in *Advances in Neural Information Processing Systems*, 35 (2022).
- [24] S. Laine, T. Karras, J. Lehtinen, and T. Aila, High-quality self-supervised deep image denoising, in *Advances in Neural Information Processing Systems*, 32 (2019).
- [25] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, Noise2Noise: Learning image restoration without clean data, in *Proceedings of ICML*, 2018, pp. 2971–2980.
- [26] T. Li, A. Abdulle, and W. E, Effectiveness of implicit methods for stiff stochastic differential equations, *Commun. Comput. Phys.*, 3 (2008), pp. 295–307.
- [27] T. Li and K. He, Back to basics: Let denoising generative models denoise, in *Proceedings of CVPR*, 2026, pp. 36115–36125.
- [28] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, SwinIR: Image restoration using Swin Transformer, in *Proceedings of ICCV*, 2021, pp. 1833–1844.
- [29] M. Mardani, J. Song, J. Kautz, and A. Vahdat, A variational perspective on solving inverse problems with diffusion models, in *International Conference on Learning Representations*, 2024.
- [30] B. Moufadi, Y. Janati El Idrissi, L. Bedin, A. O. Durmus, R. Douc, E. Moulines, and J. Olson, Variational diffusion posterior sampling with midpoint guidance, in *International Conference on Learning Representations*, 2025.
- [31] T. Pang, Z. Shen, and R. Zhang, Forward-evolution error analysis and adaptive design for matrix-valued diffusion models, arXiv:2608.15103, 2026.
- [32] N. Parikh and S. Boyd, Proximal algorithms, *Found. Trends Optim.*, 1 (2014), pp. 127–239, <https://doi.org/10.1561/2400000003>.
- [33] M. Pereyra, Proximal Markov chain Monte Carlo algorithms, *Stat. Comput.*, 26 (2016), pp. 745–

- 760, <https://doi.org/10.1007/s11222-015-9567-4>.
- [34] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, High-resolution image synthesis with latent diffusion models, in *Proceedings of CVPR*, 2022, pp. 10684–10695.
 - [35] L. I. Rudin, S. Osher, and E. Fatemi, Nonlinear total variation based noise removal algorithms, *Phys. D*, 60 (1992), pp. 259–268.
 - [36] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al., ImageNet large scale visual recognition challenge, *Int. J. Comput. Vis.*, 115 (2015), pp. 211–252.
 - [37] Y. Song and S. Ermon, Generative modeling by estimating gradients of the data distribution, in *Advances in Neural Information Processing Systems*, 32 (2019).
 - [38] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, Score-based generative modeling through stochastic differential equations, in *International Conference on Learning Representations*, 2021.
 - [39] J. Song, A. Vahdat, M. Mardani, and J. Kautz, Pseudoinverse-guided diffusion models for inverse problems, in *International Conference on Learning Representations*, 2023.
 - [40] D. Strong and T. Chan, Edge-preserving and scale-dependent properties of total variation regularization, *Inverse Problems*, 19 (2003), pp. S165–S187, <https://doi.org/10.1088/0266-5611/19/6/059>.
 - [41] Y. Sun, Z. Wu, Y. Chen, B. T. Feng, and K. L. Bouman, Provable probabilistic imaging using score-based generative priors, *IEEE Trans. Comput. Imaging*, 10 (2024), pp. 1290–1305, <https://doi.org/10.1109/TCI.2024.3449114>.
 - [42] Y. Wang, J. Yu, and J. Zhang, Zero-shot image restoration using denoising diffusion null-space model, in *International Conference on Learning Representations*, 2023.
 - [43] L. Wu, B. Trippe, C. Naesseth, D. Blei, and J. P. Cunningham, Practical and asymptotically exact conditional sampling in diffusion models, in *Advances in Neural Information Processing Systems*, 36 (2023), pp. 31372–31403.
 - [44] Z. Wu, Y. Sun, Y. Chen, B. Zhang, Y. Yue, and K. Bouman, Principled probabilistic imaging using diffusion models as plug-and-play priors, in *Advances in Neural Information Processing Systems*, 37 (2024), pp. 118389–118427.
 - [45] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, Restormer: Efficient transformer for high-resolution image restoration, in *Proceedings of CVPR*, 2022, pp. 5728–5739.
 - [46] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising, *IEEE Trans. Image Process.*, 26 (2017), pp. 3142–3155.
 - [47] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, The unreasonable effectiveness of deep features as a perceptual metric, in *Proceedings of CVPR*, 2018, pp. 586–595.
 - [48] B. Zhang, W. Chu, J. Berner, C. Meng, A. Anandkumar, and Y. Song, Improving diffusion inverse problem solving with decoupled noise annealing, in *Proceedings of CVPR*, 2025, pp. 20895–20905.
 - [49] Y. Zhu, K. Zhang, J. Liang, J. Cao, B. Wen, R. Timofte, and L. Van Gool, Denoising diffusion models for plug-and-play image restoration, in *Proceedings of the CVPR Workshops*, 2023, pp. 1219–1229.