TOWARDS UNBIASED DIFFUSION VARIATIONAL INVERSION VIA PRINCIPLED POSTERIOR MATCHING

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models provide expressive priors for imaging inverse problems, but scalable variational inference with an implicit, non-degenerate posterior remains challenging because the prior KL term depends on the score of the variational distribution. We propose *Principled Posterior Matching (PPM)*, a score-based variational framework that supports both particle-based optimization for individual observations and amortized inference through a reconstruction network for rapid, single-pass reconstruction. For two distributions evolved under the same forward diffusion process, we use the KL-dissipation relation to express their finite-time discrepancy in terms of an integrated Fisher divergence and a terminal KL term. We further derive a pathwise gradient representation using a score-projection identity under explicit differentiability, integrability, and reparameterization assumptions. When the prior and variational scores are exact, Monte Carlo sampling yields an unbiased estimator of the resulting gradient. In practice, these scores are calculated by pretrained and auxiliary models, and its performance is therefore affected only by score-estimation and optimization errors. We evaluate PPM on a range of computational imaging problems, including natural-image inpainting, super-resolution, and deblurring, fluorescence-microscopy reconstruction, and synthetic radio-interferometric black-hole imaging. Across these tasks, PPM produces more diverse reconstruction samples while maintaining competitive reconstruction fidelity relative to existing approaches. These results demonstrate that PPM provides a unified approach to sample-based variational inference and efficient amortized reconstruction across diverse imaging inverse problems.

1 INTRODUCTION

Diffusion models Ho et al. (2020); Song et al. (2020); Sohl-Dickstein et al. (2015) have been widely used as powerful priors for ill-posed inverse problems: recovering $\boldsymbol{x}$ from measurements $\boldsymbol{y} = \mathcal{A}(\boldsymbol{x}) + \boldsymbol{n}$, with forward operator $\mathcal{A}$ and measurement noise $\boldsymbol{n}$. A diffusion model learns to approximate the distribution of clean images by progressively adding noise and then training a neural network to reverse this corruption through denoising. When integrated into inverse problem solvers, diffusion priors not only yield high-quality point estimates but also facilitate full posterior exploration for uncertainty quantification, a capability crucial in real-world imaging applications.

Sampling-based solvers such as DPS Chung et al. (2022) guide reverse diffusion by approximating the time dependent likelihood $\nabla_{x_t} \log p(y|x_t) \simeq \nabla_{x_t} \log p(y|\hat{x}_0(x_t))$. Variational inference instead optimizes a posterior distribution $q_\varphi(\boldsymbol{x}_0|\boldsymbol{y})$ through a measurement-likelihood term and a prior KL: $D_{\mathrm{KL}}(q_\varphi \| p_0)$. By evaluating $p(\boldsymbol{y}|\boldsymbol{x}_0)$ directly, it avoids the approximation of $\nabla_{x_t} \log p(y|x_t)$. However, optimizing the prior KL requires the score of the variational distribution q_φ, which is generally intractable for an implicit reconstruction network.

Existing variational methods make this optimization tractable in different ways. The point-estimate formulation of RED-Diff Mardani et al. (2023) simplifies the variational score but does not explicitly model posterior spread, assuming that the variational posterior $q_\varphi(\mathbf{x}_0|\mathbf{y})$ is a degenerate Dirac delta distribution $\delta(\mathbf{x}_0 - \boldsymbol{\mu})$. DAVI Lee et al. (2024) captures implicit posterior distributions by learning an auxiliary score and adopting the integral KL objective of Diff-Instruct Luo et al. (2023). Although tractable, this time-integrated marginal KL is generally distinct from the original prior KL, yielding a different variational objective when combined with the measurement likelihood. This motivates

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.

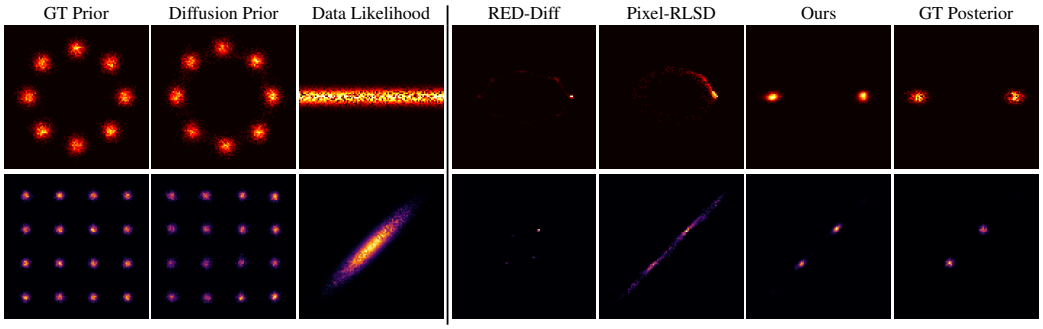

Figure 1: **Comparison of PPM and diffusion-based VI baselines (RED-Diff and RLSD) on 2D posterior estimation tasks.** The prior is a Gaussian mixture and the observation is a linear projection. In the displayed runs, PPM places samples near each visible mode of the reference posterior, whereas RED-Diff concentrates on fewer modes and the adapted RLSD baseline shows either measurement mismatch or weaker prior agreement. This toy comparison illustrates mode coverage for these two examples and is not a general claim of exact posterior recovery.

a distribution-level formulation whose tractable optimization is explicitly connected to the original posterior-KL structure.

In this paper, we introduce Principled Posterior Matching (PPM), a framework for more accurate diffusion-based variational inversion that derives its prior objective and score-based updates from the posterior-KL formulation. The KL-dissipation relation expresses the decrease in KL along a shared forward diffusion as an integrated Fisher divergence, making the diffusion endpoint terms explicit. For this Fisher objective, a score-projection identity yields a pathwise gradient without differentiating through the parameter dependence of the exact variational score, under the stated regularity and reparameterization assumptions. In practice, an auxiliary network estimates the evolving variational score. Combined with the original measurement likelihood, PPM provides a unified framework for optimizing either a particle set for a fixed observation or a stochastic reconstruction network across observations. The latter enables efficient inference on new measurements, generating each reconstruction sample in a single forward pass.

We evaluate PPM on a comprehensive set of inverse imaging benchmarks. We first validate our method on standard tasks—including image inpainting, super-resolution, and deblurring—using established datasets such as FFHQ Karras et al. (2019) and ImageNet Deng et al. (2009) at 256×256 resolution. Beyond natural images, we rigorously test PPM on two challenging scientific imaging applications: super-resolution fluorescence microscopy Qiao et al. (2021) and radio-interferometric imaging of black holes Chael et al. (2019); Sun & Bouman (2021); Sun et al. (2022). Across all experiments, PPM overall outperforms leading baselines, including DPS Chung et al. (2022), KL-based VI methods Mardani et al. (2023); Zilberstein et al. (2024), and the amortized approach DAVI Lee et al. (2024). Crucially, PPM not only achieves superior reconstruction fidelity but also delivers great sample diversity and accurate uncertainty quantification.

2 Background

2.1 Diffusion Models

Diffusion models Ho et al. (2020); Song et al. (2020); Sohl-Dickstein et al. (2015) gradually perturb data through a forward stochastic differential equation (SDE):

$$\mathrm{d}\boldsymbol{x}_t = f(\boldsymbol{x}_t, t)\,\mathrm{d}t + g(t)\,\mathrm{d}\boldsymbol{w}_t, \tag{1}$$

where $t \in [0, T]$, f and g are the drift and diffusion coefficients, and $\boldsymbol{w}_t$ is standard Wiener process. Let p_t denote the marginal distribution of $\boldsymbol{x}_t$, with $p_0 = p_{\mathrm{data}}$. The process is designed so that p_T is approximately Gaussian. Generation follows the corresponding reverse-time SDE:

$$\mathrm{d}\boldsymbol{x}_t = \left[f(\boldsymbol{x}_t, t) - g(t)^2 \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\right] \mathrm{d}t + g(t)\,\mathrm{d}\overline{\boldsymbol{w}}_t, \tag{2}$$

integrated from T to 0, where $\overline{\boldsymbol{w}}_t$ is reverse-time Wiener process. A neural network $s_\theta(\boldsymbol{x}_t, t)$ learns to approximate the score $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)$ through denoising score matching, enabling sampling via

Table 1: Comparison of posterior representations and approximation sources. "Evidence reported" describes the quantities evaluated in this paper and does not imply formal calibration.

Method	Posterior representation	Main approximation or modeling choice	Evidence reported
DPS	Reverse-diffusion samples	Approximate time-dependent likelihood score	Samples and reconstruction metrics
RED-Diff	Point estimate	Restricted, effectively zero-spread representation	Reconstruction metrics
RLSD	Finite particles	User-designed heuristic repulsive regularization	Diversity and reconstruction metrics
DAVI	Amortized network	Time-integrated marginal objective and learned score	Diversity and reconstruction metrics
PPM (Ours)	Particles or amortized network	Learned prior/scores, finite optimization trajectory	Diversity(w/ std.), reconstruction metrics

Eq. 2. For inverse problems, this learned prior is combined with the measurement likelihood to infer images consistent with the observations.

2.2 Diffusion Models for Inverse Problems

Owing to their strong ability to accurately approximate complex image distributions, diffusion priors have emerged as powerful data-driven priors for imaging inverse problems. There are two main categories of diffusion-model–based solvers for imaging inverse problems: (1) MCMC sampling Brooks et al. (2011), which uses score functions to guide gradient-based samplers and (2) Variational inference frameworks, including particle-based and amortized methods. The particle-based methods optimize the posterior distribution by score distillation Poole et al. (2022); Mardani et al. (2023); Zilberstein et al. (2024), which enforce similarity between reconstructed images and the diffusion prior by minimizing the KL divergence between them; the amortized variational inference framework Lee et al. (2024) aims to learn a parametric reconstruction network mapping observations directly to the posterior samples. Below, we provide an overview of each approach.

Monte Carlo posterior sampling. Posterior sampling replaces the unconditional score in Eq. 2 with a conditional score, decomposed through Bayes' rule as

$$\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t \mid \boldsymbol{y}) \approx s_\theta(\boldsymbol{x}_t, t) + \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{y} \mid \boldsymbol{x}_t). \tag{3}$$

The time dependent likelihood is generally intractable. DPS Chung et al. (2022) approximates it by $p_t(\boldsymbol{y} \mid \boldsymbol{x}_t) \approx p(\boldsymbol{y} \mid \hat{\boldsymbol{x}}_0)$, where $\hat{\boldsymbol{x}}_0 = \mathbb{E}[\boldsymbol{x}_0 \mid \boldsymbol{x}_t]$ is a denoised estimate. Alternative approaches include stochastic PnP Sun et al. (2023), sequential MC Dou & Song (2024), and variable-splitting methods Xu & Chi (2024). These methods generally require iterative inference for each observation.

Variational inference (Particle-based). For a fixed observation, VI optimizes a posterior approximation $q_\varphi(\boldsymbol{x}_0 \mid \boldsymbol{y})$ by minimizing the posterior KL, equivalently minimizing

$$\mathcal{L}_{\mathrm{VI}}(\varphi) = \underbrace{-\mathbb{E}_{q_\varphi(\boldsymbol{x}_0 \mid \boldsymbol{y})}[\log p(\boldsymbol{y} \mid \boldsymbol{x}_0)]}_{\mathcal{L}_{\mathrm{data}}} + \underbrace{D_{\mathrm{KL}}(q_\varphi(\cdot \mid \boldsymbol{y}) \| p_0)}_{\mathcal{L}_{\mathrm{prior}}}, \tag{4}$$

ignoring the constant $\log p(\boldsymbol{y})$. This formulation evaluates the measurement likelihood directly on clean reconstructions, avoiding time dependent likelihood approximations. The remaining challenge is optimizing the diffusion-prior term. Point-estimate formulations such as SDS Poole et al. (2022) and RED-Diff Mardani et al. (2023) simplify prior optimization through restricted variational representations, but do not explicitly model posterior spread. RLSD Zilberstein et al. (2024) introduces repulsion between reconstructions to promote diversity.

Variational inference (Amortized). Rather than optimizing each observation separately, amortized VI trains a reconstruction network for reuse across observations. DAVI Lee et al. (2024) follows Diff-Instruct Luo et al. (2023) in replacing the prior KL with an integral KL objective:

$$\mathcal{L}_{\mathrm{IKL}}(\varphi) = \int_0^T \omega(t)\, D_{\mathrm{KL}}(q_{\varphi,t}(\cdot \mid \boldsymbol{y}) \| p_t)\, \mathrm{d}t, \tag{5}$$

where $q_{\varphi,t}$ is obtained by applying the forward diffusion to q_φ, and $\omega(t)$ weights diffusion times. Diff-Instruct derives a gradient expressed through the prior and variational scores. However, this time-integrated objective generally differs from the original prior KL at $t = 0$; their relationship depends on the diffusion dynamics, horizon, and weighting. This distinction motivates examining the KL–Fisher relation under forward diffusion, with the endpoint terms stated explicitly.

3 METHOD

This section develops Principled Posterior Matching (PPM) as a score-based variational framework for inverse problems. We first establish the finite-time relation between KL dissipation and Fisher divergence, including its terminal term. We then derive a pathwise gradient representation under explicit regularity and reparameterization assumptions. The result concerns the gradient of a specified diffused score-discrepancy objective; it does not imply that a finite particle set or an amortized network equals the Bayesian posterior. The practical algorithm further replaces exact distribution scores with pretrained and auxiliary neural estimates.

3.1 PROBLEM FORMULATION

Following Bayes' rule in Eq. 4, classical variational inference for an inverse problem can be written, up to the constant $\log p(\boldsymbol{y})$, as

$$\begin{aligned}\varphi^* &= \arg\min_{\varphi} D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}_0 \mid \boldsymbol{y}) \,\|\, p(\boldsymbol{x}_0 \mid \boldsymbol{y})) \\ &= \arg\min_{\varphi} \mathbb{E}_{q_\varphi(\boldsymbol{x}_0|\boldsymbol{y})}[-\log p(\boldsymbol{y} \mid \boldsymbol{x}_0)] + D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}_0 \mid \boldsymbol{y}) \,\|\, p_0(\boldsymbol{x}_0))\,.\end{aligned} \tag{6}$$

For additive Gaussian measurement noise, the first expectation is proportional to $\mathbb{E}_{q_\varphi}\|\boldsymbol{y}-\mathcal{A}(\boldsymbol{x}_0)\|_2^2$, up to constants independent of φ. It therefore enforces consistency with the observation through the forward model $\mathcal{A}$. The remaining difficulty is the prior KL term: when q_φ is implicit and non-degenerate, its gradient depends on the score of the variational distribution.

3.2 FINITE-TIME FISHER–KL RELATION

Let $p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$ denote the transition kernel of the forward diffusion. Let $q_{\varphi,t}(\boldsymbol{x}_t \mid \boldsymbol{y})$ and $p_t(\boldsymbol{x}_t)$ be obtained by evolving $q_{\varphi,0}(\boldsymbol{x}_0 \mid \boldsymbol{y})$ and $p_0(\boldsymbol{x}_0)$, respectively, under the same forward SDE. We use $t_{\min} \geq 0$ as the smallest diffusion time included in the score objective. Under the required regularity and integrability conditions, the relative entropy dissipates along the common forward diffusion according to the standard relative-entropy dissipation identity (Song et al., 2021).

$$\begin{aligned}&D_{\mathrm{KL}}(q_{\varphi,t_{\min}}(\cdot \mid \boldsymbol{y}) \,\|\, p_{t_{\min}}) \\ &\quad = D_{\mathrm{KL}}(q_{\varphi,T}(\cdot \mid \boldsymbol{y}) \,\|\, p_T) + \frac{1}{2}\int_{t_{\min}}^{T} g(t)^2 \mathbb{E}_{q_{\varphi,t}(\cdot|\boldsymbol{y})}\Big[\big\|s_{q_{\varphi,t}}(\boldsymbol{x}_t \mid \boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t)\big\|_2^2\Big]\,\mathrm{d}t,\end{aligned} \tag{7}$$

where

$$s_{p_t}(\boldsymbol{x}_t) := \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t), \quad s_{q_{\varphi,t}}(\boldsymbol{x}_t \mid \boldsymbol{y}) := \nabla_{\boldsymbol{x}_t} \log q_{\varphi,t}(\boldsymbol{x}_t \mid \boldsymbol{y}) \tag{8}$$

are exact distribution scores. The terminal term vanishes only when $q_{\varphi,T}(\cdot \mid \boldsymbol{y}) = p_T$, or in a limiting setting in which both processes converge to the same terminal law. With a finite diffusion horizon, any nonzero terminal term remains an approximation source if it is not explicitly optimized.

Optimizing the integrated term in Eq. 7 requires differentiating a discrepancy that contains the φ-dependent score $s_{q_{\varphi,t}}$. The following result provides a pathwise representation of this gradient.

Theorem 1 (Pathwise Gradient Equivalence)**.** *Fix $t > 0$. Assume that $\boldsymbol{x}_0 = G_\varphi(\boldsymbol{\xi}, \boldsymbol{y})$, where $\boldsymbol{\xi}$ has a φ-independent distribution indicating the number of particles, and that $\boldsymbol{x}_t = F_t(\boldsymbol{x}_0, \boldsymbol{\epsilon})$, where $\boldsymbol{\epsilon}$ is also sampled from a φ-independent distribution. Assume that F_t and G_φ are differentiable almost everywhere, $p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$ is positive and continuously differentiable in $\boldsymbol{x}_t$, differentiation may be exchanged with the integrals defining $q_{\varphi,t}$ and the expectations below, and all resulting terms are dominated by an integrable envelope. Let $d : \mathbb{R}^D \to \mathbb{R}$ be twice continuously differentiable, define $d'(\boldsymbol{r}) := \nabla_{\boldsymbol{r}} d(\boldsymbol{r})$, and assume that its required derivatives satisfy the same domination condition. Define*

$$\Delta_{\varphi,t}(\boldsymbol{x}_t, \boldsymbol{y}) := s_{q_{\varphi,t}}(\boldsymbol{x}_t \mid \boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t).$$

Then

$$\begin{aligned}\nabla_\varphi \mathbb{E}_{\boldsymbol{\xi},\boldsymbol{\epsilon}}[d(\Delta_{\varphi,t}(\boldsymbol{x}_t, \boldsymbol{y}))] = \nabla_\varphi \mathbb{E}_{\boldsymbol{\xi},\boldsymbol{\epsilon}}\Big[&- d'\big(\mathrm{sg}_\varphi[s_{q_{\varphi,t}}](\boldsymbol{x}_t \mid \boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t)\big)^\top \\ &\times \big(\mathrm{sg}_\varphi[s_{q_{\varphi,t}}](\boldsymbol{x}_t \mid \boldsymbol{y}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0)\big) + d\big(\mathrm{sg}_\varphi[s_{q_{\varphi,t}}](\boldsymbol{x}_t \mid \boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t)\big)\Big].\end{aligned} \tag{9}$$

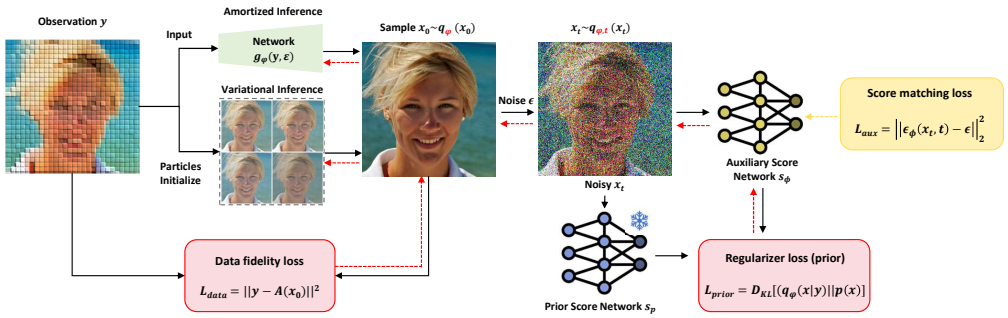

Figure 2: **Overview of PPM.** The variational approximation q_φ is parameterized by either a finite particle set or an amortized reconstruction network. A fixed pretrained network s_θ approximates the prior score, while an auxiliary network s_ϕ is trained to approximate the score of the current diffused variational distribution. PPM alternates between updating s_ϕ and updating φ using a data-fidelity term and a score-based gradient surrogate.

The operator sg_φ *blocks only the explicit dependence of the exact variational score on* φ*; gradients through the sampled path* $\boldsymbol{x}_t$ *are retained.*

Theorem 1 is an equality for exact distribution scores. Sampling the fixed base variables $(\boldsymbol{\xi}, \boldsymbol{\epsilon})$ gives an unbiased Monte Carlo estimator of the displayed pathwise gradient under the stated assumptions. For K particles, let $k \sim \mathrm{Uniform}\{1, \dots, K\}$, set $\boldsymbol{\xi} = k$, and define $G_\varphi(k, \boldsymbol{y}) = \boldsymbol{\mu}_k$. The base index is independent of $\varphi = \{\boldsymbol{\mu}_k\}_{k=1}^K$, and Gaussian perturbation gives $q_{\varphi,t} = K^{-1} \sum_k p_t(\cdot \mid \boldsymbol{\mu}_k)$, which has a smooth density for $t > 0$.

Proof sketch. For fixed φ, differentiating the diffused marginal

$$q_{\varphi,t}(\boldsymbol{x}_t \mid \boldsymbol{y}) = \int p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0) q_\varphi(\mathrm{d}\boldsymbol{x}_0 \mid \boldsymbol{y})$$

with respect to $\boldsymbol{x}_t$ gives the conditional-score identity

$$s_{q_{\varphi,t}}(\boldsymbol{x}_t \mid \boldsymbol{y}) = \mathbb{E}[\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0) \mid \boldsymbol{x}_t, \boldsymbol{y}] .$$

Consequently, for every integrable vector function $\boldsymbol{u}$,

$$\mathbb{E}\big[\boldsymbol{u}^\top \big\{ s_{q_{\varphi,t}} - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0) \big\}\big] = 0.$$

Writing this expectation over the fixed base variables and differentiating it pathwise accounts for both the explicit φ-dependence of $s_{q_{\varphi,t}}$ and the dependence through the sampled path. Choosing $\boldsymbol{u} = d'(\mathrm{sg}_\varphi[s_{q_{\varphi,t}}] - s_{p_t})$ and rearranging gives Eq. 9. Appendix A provides the full directional-derivative argument. □

In practice, the fixed pretrained diffusion network $s_\theta(\boldsymbol{x}_t, t)$ approximates $s_{p_t}(\boldsymbol{x}_t)$, and an auxiliary network $s_\phi(\boldsymbol{x}_t, t)$ approximates $s_{q_{\varphi,t}}(\boldsymbol{x}_t \mid \boldsymbol{y})$. We train $s_\phi(\boldsymbol{x}_t, t)$ by denoising score matching:

$$\mathcal{L}_{\mathrm{aux}}(\phi) = \mathbb{E}_{\substack{t \sim \rho(t),\ \boldsymbol{x}_0 \sim q_\varphi(\cdot \mid \boldsymbol{y}) \\ \boldsymbol{x}_t \sim p_t(\cdot \mid \boldsymbol{x}_0)}} \Big[\lambda_{\mathrm{aux}}(t) \, \| s_\phi(\boldsymbol{x}_t, t) - s_{\mathrm{cond},t}(\boldsymbol{x}_t, \boldsymbol{x}_0) \|_2^2 \Big] , \tag{10}$$

where ρ is the time-sampling density on $[t_{\min}, T]$ and $s_{\mathrm{cond},t}(\boldsymbol{x}_t, \boldsymbol{x}_0) := \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$. For the Gaussian perturbation $\boldsymbol{x}_t = \alpha_t \boldsymbol{x}_0 + \sigma_t \boldsymbol{\epsilon}$, with $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$, this target is

$$s_{\mathrm{cond},t}(\boldsymbol{x}_t, \boldsymbol{x}_0) = -\frac{\boldsymbol{x}_t - \alpha_t \boldsymbol{x}_0}{\sigma_t^2} = -\frac{\boldsymbol{\epsilon}}{\sigma_t} . \tag{11}$$

All network outputs below are expressed in score units. If a diffusion implementation predicts noise, its output is converted to score units before score differences are formed. We initialize s_ϕ from s_θ and update it on samples from the current variational approximation. Replacing the exact scores by s_θ and s_ϕ makes the practical update a plug-in approximation to Theorem 1; score-estimation and incomplete-optimization errors therefore remain.

Inference after optimization. The two modes have different test-time semantics. In particle-based VI, Algorithm 1 is itself a per-observation inference procedure. For a fixed measurement $\boldsymbol{y}$, the optimized set $\{\boldsymbol{\mu}_k\}_{k=1}^K$ is returned as an empirical approximation, and a reported sample is obtained by selecting a particle uniformly. Its empirical mean and pixel-wise standard deviation are computed directly from this set; the optimized particles are not a reusable sampler for a new observation. In amortized inference, the trained network is reusable. For a new observation $\boldsymbol{y}$, draw $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, h\boldsymbol{I})$ and compute $\boldsymbol{x} = g_\varphi(\boldsymbol{y} + \boldsymbol{\eta})$. Repeating this forward pass with independent perturbations produces an empirical sample set. When $h = 0$, the network instead returns a single deterministic reconstruction. These outputs are samples from the learned variational approximation, not exact posterior samples.

3.3 Unified Optimization Framework

Let $\rho(t)$ be the time-sampling density and define $a(t) = g(t)^2/[2\rho(t)]$. Sampling $t \sim \rho$ then converts the Fisher integral in Eq. 7 into an expectation. With $\boldsymbol{r}_t = s_\phi(\boldsymbol{x}_t, t) - s_\theta(\boldsymbol{x}_t, t)$, the implemented scalar surrogate is

$$\mathcal{L}_{\mathrm{PPM}}(\varphi) = \mathcal{L}_{\mathrm{prior}}(\varphi) + \lambda_{\mathrm{data}}\mathcal{L}_{\mathrm{data}}(\varphi),$$

$$\mathcal{L}_{\mathrm{data}}(\varphi) = \mathbb{E}_{\boldsymbol{x}_0 \sim q_\varphi(\cdot|\boldsymbol{y})}\left[\|\boldsymbol{y} - \mathcal{A}(\boldsymbol{x}_0)\|_2^2\right], \quad \mathcal{L}_{\mathrm{prior}}(\varphi) = \mathbb{E}_{\substack{t\sim\rho(t),\ \boldsymbol{x}_0\sim q_\varphi(\cdot|\boldsymbol{y}) \\ \boldsymbol{x}_t\sim p_t(\cdot|\boldsymbol{x}_0)}}\left[a(t)D(\varphi)\right], \quad (12)$$

$$D(\varphi) = -d'(\boldsymbol{r}_t)^\top \left(s_\phi(\boldsymbol{x}_t, t) - s_{\mathrm{cond},t}(\boldsymbol{x}_t, \boldsymbol{x}_0)\right) + d(\boldsymbol{r}_t)$$

During the φ update, θ and ϕ are fixed, while differentiation retains the dependence of both score-network outputs on $\boldsymbol{x}_t$. Thus $\mathcal{L}_{\mathrm{prior}}$ is an auto-differentiation surrogate for the integrated score-based gradient, rather than a probability divergence whose scalar value is itself minimized exactly. For amortized training, an additional outer expectation over $\boldsymbol{y} \sim \mathcal{D}_{\boldsymbol{y}}$ is understood in Eq. 12.

If the exact scores are used, the Monte Carlo gradient of Eq. 12 is unbiased for the weighted integrated gradient specified above. Connecting this gradient to the original prior KL additionally requires the terminal term in Eq. 7 to vanish or to be included.

The squared discrepancy $d(\boldsymbol{r}) = \|\boldsymbol{r}\|_2^2$ is the choice connected to the Fisher–KL relation in Eq. 7. The pathwise manipulation in Theorem 1 can also be written for other differentiable scalar functions d, but replacing the squared discrepancy does not automatically define or minimize a known probability divergence. Establishing such correspondences is left to future work.

Particle-based variational inference. For a fixed observation $\boldsymbol{y}$, the parameters are the image particles $\varphi = \{\boldsymbol{\mu}_k\}_{k=1}^K$. The particles define an empirical distribution at $t = 0$ and a smooth Gaussian mixture after positive-time perturbation. Alternating score learning and particle updates produces multiple candidate reconstructions for the same observation, subject to finite-particle, score-estimation, and optimization error.

Amortized inference. The amortized parameter φ denotes the weights of a reconstruction network g_φ. During training and sampling, the perturbation $\boldsymbol{\eta}$ supplies the random input that induces $q_\varphi(\boldsymbol{x}_0|\boldsymbol{y})$ for a fixed observation. Initializing g_φ from a compatible pretrained diffusion U-Net accelerates training, and the resulting model provides single-pass reconstructions for new observations.

4 Experiment

In this section, we compare our proposed method, PPM, with state-of-the-art (SoTA) diffusion model-based methods for solving inverse problems, particularly those employing variational inference. Our experiments aim to demonstrate PPM's capability to generate diverse reconstructions while maintaining fidelity, thereby recovering the full posterior and surpassing baselines that typically yield homogeneous results.

4.1 Toy Examples of 2D Posterior Estimation

We first validated PPM on two simple 2D posterior-estimation tasks. In each task, the hidden variable $x \in \mathbb{R}^{2\times 1}$ follows a Gaussian-mixture prior, and measurements obey the linear model

$$y = Fx + n, \quad (13)$$

where F is a linear projection matrix and $n \sim \mathcal{N}(0, \sigma^2 I)$ is Gaussian noise. We compare PPM against two related baselines, RED-Diff Mardani et al. (2023) and RLSD Zilberstein et al. (2024), using the same pre-trained pixel-space diffusion prior. To ensure fair compaison, we adapt RLSD by replacing its latent diffusion backbone with our pixel-space prior and incorporating its repulsive term into a pixel-space SDS loss, so that all methods operate under identical conditions.

Figure 1 highlights the significant differences in each method's ability to capture multimodal posteriors. RED-Diff exhibits severe mode collapse: its samples converge to accurate point estimates but lack diversity. RLSD produces more varied samples, yet its empirical repulsive term leads to either measurement inconsistency (top example) or poor prior adherence (bottom example). In contrast, PPM faithfully recovers every posterior mode, yielding both accurate and diverse samples that closely match the ground-truth bimodal distribution.

4.2 Computational Photography

We then evaluated PPM on various computational photography tasks, including box inpainting, motion deblurring, and super-resolution, using natural images.

Datasets and Pretrained Models We primarily evaluate on two natural-image datasets with distinct characteristics—FFHQ (256 × 256) Karras et al. (2019) and ImageNet (256 × 256) Deng et al. (2009)—using 32 randomly selected images from each validation set. For FFHQ, we use the diffusion prior released by DPS Chung et al. (2022), and for ImageNet we employ the model from Dhariwal & Nichol (2021). Both are used off-the-shelf, without task-specific fine-tuning. RLSD Zilberstein et al. (2024) requires a latent diffusion backbone, so we use Stable Diffusion v2.1 Rombach et al. (2022) following its original setup.

Baselines We compare PPM against state-of-the-art computational imaging methods that leverage diffusion priors, spanning both variational-inference and MCMC-sampling paradigms. Our VI baselines include RED-Diff Mardani et al. (2023) and RLSD Zilberstein et al. (2024), while our MCMC baselines are ΠGDM Song et al. (2022) and DPS Chung et al. (2022). All methods use the same pixel-space diffusion prior—except RLSD, which retains its Stable Diffusion prior—and are evaluated with hyperparameters set to their original defaults to ensure a fair comparison.

Evaluation Metrics We assess reconstruction fidelity using PSNR and SSIM, and quantify sampling diversity using the pairwise cosine similarity among N reconstruction samples from the same observation. The final diversity is the average of all O observations:

$$1 - \frac{1}{|O|} \sum_{o} \frac{2}{N(N-1)} \sum_{i<j} \frac{\langle x_{oi}, x_{oj} \rangle}{\|x_{oi}\|_2 \, \|x_{oj}\|_2}.$$

For quantitative and efficient comparison, we estimate 1,000 observations in ImageNet and FFHQ validation dataset and reconstruct 8 samples from each observation.

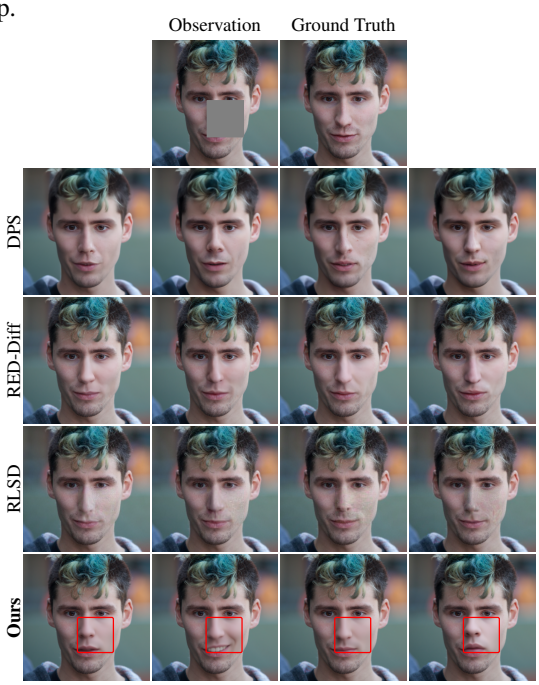

Figure 3: **Box inpainting on FFHQ.** Top: a masked observation and its ground truth. Subsequent rows show four samples each from DPS, RED-Diff, RLSD, and PPM (top to bottom). PPM produces more diverse, high-fidelity completions in the masked region (red boxes), while the baselines show limited variation.

Inpainting Box-inpainting results are assessed on FFHQ validation images. Qualitative comparisons with 80×80 masked boxes are visualized in Fig. 3, while quantitative metrics evaluated on larger 128×128 center masks are reported in Table 2. We compare PPM against two VI baselines, RED-Diff Mardani et al. (2023) and RLSD Zilberstein et al. (2024). PPM significantly surpasses

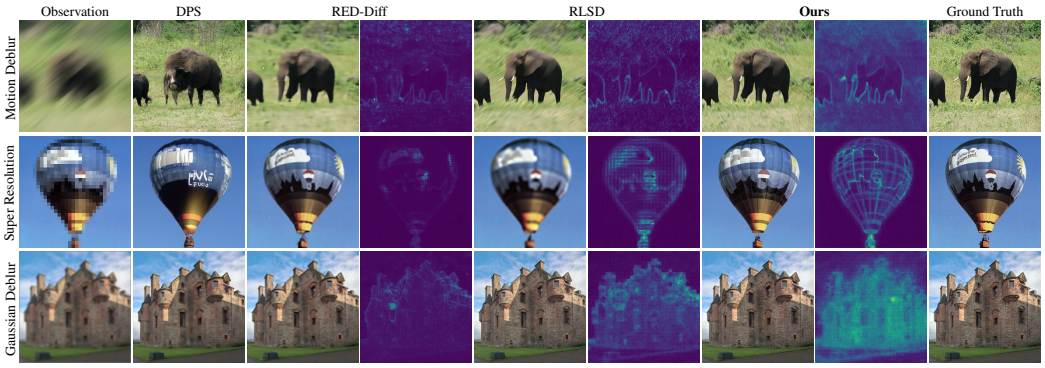

Figure 4: **Comparison of PPM, DPS, and diffusion-based VI baselines (RED-Diff, RLSD) on ImageNet motion deblurring, super-resolution and Gaussian deblurring.** For each VI method, we show a reconstruction sample (left) and standard-deviation uncertainty map (right). PPM achieves higher fidelity than DPS, accurately recovering fine details such as elephant tusks, balloon logos and human faces, and provides more spatially aligned uncertainty estimates than competing VI methods: RED-Diff shows mode collapse, and RLSD has strong artifacts in uncertainty maps.

Table 2: **Quantitative comparison of PPM and baselines across super-resolution, motion deblurring, and box inpainting on FFHQ.** All methods use the same pretrained diffusion model, except RLSD, which employs Stable Diffusion. The best and second best results are highlighted.

Method	Super Resolution			Motion Deblurring			Box Inpainting		
	PSNR ↑	SSIM ↑	Diversity ↑	PSNR ↑	SSIM ↑	Diversity ↑	PSNR ↑	SSIM ↑	Diversity ↑
Sampling									
DDRM	22.62	0.62	0.001	-	-	-	22.11	0.78	0.004
DPS	21.02	0.57	0.010	20.34	0.55	0.006	23.43	0.80	0.009
ΠGDM	23.92	0.67	0.011	25.82	0.75	0.006	23.25	0.86	0.010
Particle									
RED-Diff	26.21	0.70	0.001	29.02	0.84	0.001	24.69	0.87	0.002
Ours(VI)	26.71	0.72	0.013	28.18	0.83	0.009	28.73	0.97	0.016
Amortized inference									
DAVI	24.49	0.70	0.006	27.69	0.82	0.003	24.58	0.83	0.005
Ours(AI)	24.85	0.73	0.008	29.17	0.85	0.004	25.26	0.88	0.008

both baselines in diversity while simultaneously delivering superior reconstruction quality (PSNR and SSIM). In challenging cases where critical features like facial contours (Fig. 3) are masked, RED-Diff collapses to nearly identical outputs and RLSD struggles to trade off data fidelity against prior adherence. By contrast, PPM produces varied yet plausible reconstructions (e.g., different lip shapes) thanks to its principled score-based divergence and tailored optimization strategy.

Motion Deblurring, Super Resolution, and Gaussian Denoising For motion deblurring, we follow Chung et al. (2022) by convolving each image with a randomly sampled 61×61 motion kernel (variance = 0.3^2). For super-resolution, we downsample images by a factor of 8. Additionally, for Gaussian denoising, we corrupt the images with additive white Gaussian noise with a standard deviation of $\sigma = 0.2$. We evaluate PPM on these tasks using FFHQ and ImageNet validation sets, comparing against DPS Chung et al. (2022), ΠGDM Song et al. (2022), RED-Diff Mardani et al. (2023), RLSD Zilberstein et al. (2024), and the amortized baseline DAVI Lee et al. (2024).

Figure 4 shows that PPM delivers sharper, more observation-consistent reconstructions than DPS, evident in fine details such as the elephant's tusks and the balloon's logo—thanks to PPM's exact likelihood term versus DPS's approximation. Compared to VI methods (RED-Diff and RLSD), PPM also excels at uncertainty quantification: its standard-deviation maps accurately reflect positional uncertainty (e.g., the elephant's tusks, the boy's facial edge), whereas RED-Diff collapses modes and RLSD introduces p persistent speckle artifacts due to latent-space optimization.

Figure 6 presents a qualitative comparison of the amortized inference results on the FFHQ dataset. We contrast our PPM framework against the KL-divergence-based baseline, DAVI. The pixel-wise uncertainty maps demonstrate that our method captures meaningful posterior diversity, particularly

in ambiguous regions (e.g., edges and textures), whereas the KL-based approach exhibits signs of posterior collapse with largely suppressed uncertainty.

Table 2 confirms that PPM achieves the best overall balance of PSNR, SSIM, and diversity. While RLSD slightly outperforms in PSNR/SSIM for super-resolution—due to its higher-resolution Stable Diffusion prior (512 × 512)—its artifacts undermine visual quality. MCMC sampling methods like DPS match PPM's diversity but fall short in fidelity. In summary, PPM outperforms both variational and sampling-based baselines across computational photography tasks, delivering superior reconstruction quality and reliable uncertainty estimates.

4.3 Radio Interferometric Black Hole Imaging

We applied PPM to reconstruct and quantify uncertainty in black hole images from very long baseline interferometry (VLBI) measurements. Using a general-relativistic magnetohydrodynamics (GRMHD) simulated Sagittarius A* black hole image, we emulate a synthetic observation of the Event Horizon Telescope (EHT) array, which comprises nine telescopes worldwide to form an Earth-sized interferometer. Ignoring atmospheric turbulence, the measurement model that maps the true image x to the observed visibilities y is

$$y = M\mathcal{F}\{x\} + n, \tag{14}$$

where $\mathcal{F}$ is the Fourier transform, M selects the measured frequency components, and n is additive Gaussian noise. Because the EHT sampling is extremely sparse in the Fourier domain (Fig. 5(a)), this defines a highly ill-posed inverse problem: enforcing only data consistency yields the classical dirty image, riddled with sidelobe artifacts (top right of Fig. 5).

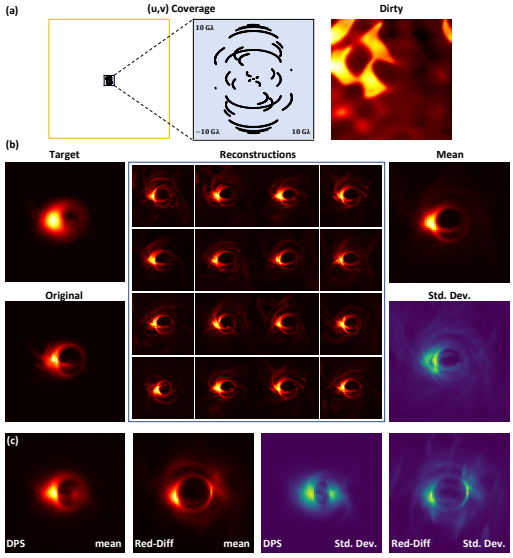

Figure 5: **Black hole interferometric imaging on synthetic EHT data.** This ill-posed VLBI inverse problem reconstructs images from sparse Fourier samples. (a) EHT (u, v) coverage and prior-free dirty image. (b) GRMHD ground truth (blurred to EHT resolution), 16 PPM posterior samples, mean reconstruction and std map. (c) Baseline comparison.

Robust uncertainty quantification is therefore critical before making scientific inferences. Our PPM reconstruction follows the InverseBench Zheng et al. (2025) protocol, using a diffusion prior trained on approximately 50,000 synthetic black hole images. Figure 5 (b) shows the ground-truth GRMHD image, the target blurred to the EHT's resolution, 16 independent PPM posterior samples, and the resulting mean reconstruction alongside its standard-deviation map. We further compare PPM with baselines in Fig. 5(c). While RED-Diff suffers from severe mode collapse (indicated by the suppressed uncertainty map) and DPS yields blurry reconstructions, only PPM faithfully captures the key morphology—ring diameter, azimuthal position of the bright crescent, and the black hole's swirling signature, demonstrating PPM's ability to deliver accurate reconstructions with reliable uncertainty estimates in challenging VLBI black hole imaging scenarios.

5 Conclusion

In this paper, we presented Principled Posterior Matching (PPM), a principled framework that addresses the fundamental limitations of existing variational diffusion-based inverse problem solvers. By identifying that the mode collapse in prior works stems from biased approximations of the KL divergence, we proposed a rigorous alternative based on the integration of Fisher divergence. This enables an unbiased gradient estimator, allowing for the exact minimization of the variational objective without structural collapse. PPM unifies variational and amortized inference, enabling both faithful posterior recovery and efficient, unsupervised generation. Validated across computational and scientific imaging tasks—including microscopy and black-hole imaging—PPM consistently outperforms baselines. Its superior fidelity and reliable uncertainty quantification establish it as a robust foundation for trustworthy imaging.

AI USE STATEMENT

In this work, we did not use generative AI tools for research ideation, experimental design, methodology development, or writing the core content of this paper. We used generative AI tools only for minor grammar and language polishing of the manuscript text. We have reviewed all AI-assisted text to ensure it accurately reflects our original intent and findings. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REFERENCES

Asad Aali, Marius Arvinte, Sidharth Kumar, and Jonathan I Tamir. Solving inverse problems with score-based generative priors learned from noisy data. *arXiv preprint arXiv:2305.01166*, 2023.

Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of markov chain monte carlo*. CRC press, 2011.

Andrew Chael, Katie Bouman, Michael Johnson, Maciek Wielgus, Lindy Blackburn, Chi-Kwan Chan, Joseph Rachid Farah, Daniel Palumbo, and Dominic Pesce. eht-imaging: v1. 1.0: Imaging interferometric data with regularized maximum likelihood. *Zenodo*, 2019.

Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

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

Zehao Dou and Yang Song. Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *The Twelfth International Conference on Learning Representations*, 2024.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 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*, pp. 4401–4410, 2019.

Sojin Lee, Dogyun Park, Inho Kong, and Hyunwoo J Kim. Diffusion prior-based amortized variational inference for noisy inverse problems. In *European Conference on Computer Vision*, pp. 288–304. Springer, 2024.

Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. Diff-instruct: A universal approach for transferring knowledge from pre-trained diffusion models. *Advances in Neural Information Processing Systems*, 36:76525–76546, 2023.

Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391*, 2023.

Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022.

Chang Qiao, Di Li, Yuting Guo, Chong Liu, Tao Jiang, Qionghai Dai, and Dong Li. Evaluation and development of deep neural networks for image super-resolution in optical microscopy. *Nature methods*, 18(2):194–202, 2021.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.

Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*, 2022.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34:1415–1428, 2021.

He Sun and Katherine L Bouman. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 2628–2637, 2021.

He Sun, Katherine L Bouman, Paul Tiede, Jason J Wang, Sarah Blunt, and Dimitri Mawet. α-deep probabilistic inference (α-dpi): efficient uncertainty quantification from exoplanet astrometry to black hole feature extraction. *The Astrophysical Journal*, 932(2):99, 2022.

Yu Sun, Zihui Wu, Yifan Chen, Berthy T Feng, and Katherine L Bouman. Provable probabilistic imaging using score-based generative priors. *arXiv preprint arXiv:2310.10835*, 2023.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.

Xingyu Xu and Yuejie Chi. Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction. *arXiv preprint arXiv:2403.17042*, 2024.

Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy T Feng, Caifeng Zou, Yu Sun, Nikola Kovachki, Zachary E Ross, et al. Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. *arXiv preprint arXiv:2503.11043*, 2025.

Mingyuan Zhou, Huangjie Zheng, Zhendong Wang, Mingzhang Yin, and Hai Huang. Score identity distillation: Exponentially fast distillation of pretrained diffusion models for one-step generation. In *Forty-first International Conference on Machine Learning*, 2024.

Nicolas Zilberstein, Morteza Mardani, and Santiago Segarra. Repulsive latent score distillation for solving inverse problems. *arXiv preprint arXiv:2406.16683*, 2024.

A PROOFS

In this section we present a detailed proof of the Gradient Equivalence Theorem. The proof is based on the so-called **Score-projection identity** which was first found in Vincent (2011) to bridge denoising score matching and denoising auto-encoders. Later the identity is applied for deriving distillation methods based on Fisher divergences Zhou et al. (2024). We follow the derivation logic of Zhou et al. Zhou et al. (2024) and adapt it to our PPM framework.

Score-projection identity. Let $\boldsymbol{u}(\cdot, \varphi)$ be a vector-valued function. Using the notations of the Gradient Equivalence Theorem, under mild conditions, the identity holds:

$$\mathbb{E}_{\substack{\boldsymbol{x}_0 \sim q_\varphi(\cdot|\boldsymbol{y}) \\ \boldsymbol{x}_t|\boldsymbol{x}_0 \sim p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)}} \Big[f(\boldsymbol{x}_t, \boldsymbol{x}_0, \varphi) \Big] = 0, \quad \forall \varphi \tag{A.1}$$
$$f(\boldsymbol{x}_t, \boldsymbol{x}_0, \varphi) = \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T \big\{ s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{x}_0) \big\}.$$

Next, we turn to prove the Gradient Equivalence Theorem.

Proof of Gradient Equivalence Theorem. We start by applying the chain rule for the total derivative with respect to φ. The function $d(\cdot)$ depends on φ both directly through the score function $s_{q_{\varphi,t}}$ and indirectly through the distribution $\boldsymbol{x}_t \sim q_{\varphi,t}$ (as $\boldsymbol{x}_t$ depends on $\boldsymbol{x}_0 \sim q_\varphi(\cdot|\boldsymbol{y})$). This gives two terms:

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{x}_t \sim q_{\varphi,t}} \frac{\partial}{\partial \varphi} d \left(s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t) \right) \\
&= \mathbb{E}_{\substack{\boldsymbol{x}_0 \sim q_\varphi(\cdot|\boldsymbol{y}) \\ \boldsymbol{x}_t|\boldsymbol{x}_0 \sim p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)}} \Bigg[d' \left(s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t) \right)^T \frac{\partial}{\partial \varphi} s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) \\
&+ \frac{\partial}{\partial \boldsymbol{x}_t} d(s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) - s_{p_t}(\boldsymbol{x}_t)) \frac{\partial \boldsymbol{x}_t}{\partial \varphi} \Bigg]
\end{aligned} \tag{A.2}$$

We differentiate (A.1) on both sides w.r.t φ to resolve the first term in (A.2). Since the expectation is zero for all φ, its derivative is also zero. We apply the total derivative (multivariate chain rule), noting that f depends on φ directly, as well as indirectly through $\boldsymbol{x}_0$ and $\boldsymbol{x}_t$ (both depend on φ):

$$0 = \frac{\partial}{\partial \varphi} \mathbb{E}_{\substack{\boldsymbol{x}_0 \sim q_\varphi(\cdot|\boldsymbol{y}) \\ \boldsymbol{x}_t|\boldsymbol{x}_0 \sim p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)}} \Big[f(\boldsymbol{x}_t, \boldsymbol{x}_0, \varphi) \Big] \tag{A.3}$$
$$= \mathbb{E} \Big[\frac{\partial f}{\partial \boldsymbol{x}_0} \frac{\partial \boldsymbol{x}_0}{\partial \varphi} + \frac{\partial f}{\partial \boldsymbol{x}_t} \frac{\partial \boldsymbol{x}_t}{\partial \varphi} + \frac{\partial f}{\partial \varphi} \Big] \tag{A.4}$$
$$\begin{aligned}
&= \mathbb{E} \Big[\frac{\partial}{\partial \boldsymbol{x}_0} \Big\{ \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T \{ -\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{x}_0) \} \Big\} \frac{\partial \boldsymbol{x}_0}{\partial \varphi} \\
&+ \frac{\partial}{\partial \boldsymbol{x}_t} \Big\{ \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T \{ s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{x}_0) \} \Big\} \frac{\partial \boldsymbol{x}_t}{\partial \varphi} \\
&+ \frac{\partial}{\partial \varphi} \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) + \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T \frac{\partial}{\partial \varphi} \{ s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) \} \Big]
\end{aligned} \tag{A.5}$$
$$\begin{aligned}
&= \mathbb{E} \Big[\frac{\partial}{\partial \varphi} \Big\{ \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T \{ s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{x}_0) \} \Big\} \\
&+ \boldsymbol{u}(\boldsymbol{x}_t, \varphi)^T \frac{\partial}{\partial \varphi} \{ s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y}) \} \Big].
\end{aligned} \tag{A.6}$$

By rearranging the terms in (A.4) and (A.5), and by applying the stop-gradient operator sg to $s_{q_{\varphi,t}}$ within the derivative terms (treating it as fixed relative to the path differentiation), the expression can be consolidated into (A.6). We then rewrite (A.6) as follows:

$$\mathbb{E}_{\boldsymbol{x}_t\sim q_{\varphi,t}}\boldsymbol{u}(\boldsymbol{x}_t,\varphi)^T\frac{\partial}{\partial\varphi}\{s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y})\}$$
$$=-\frac{\partial}{\partial\varphi}\mathbb{E}_{\substack{\boldsymbol{x}_0\sim q_\varphi(\cdot|\boldsymbol{y})\\ \boldsymbol{x}_t|\boldsymbol{x}_0\sim p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)}}\Big\{\boldsymbol{u}(\boldsymbol{x}_t,\varphi)^T$$
$$\{s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y})-\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)\}\Big\} \tag{A.7}$$

Let $\boldsymbol{u}(\boldsymbol{x}_t,\varphi)=d'\left(s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y})-s_{p_t}(\boldsymbol{x}_t)\right)$ and substitute this specific function $\boldsymbol{u}$ into the identity (A.7):

$$\mathbb{E}_{\boldsymbol{x}_t\sim q_{\varphi,t}}\Big[d'\left(s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y})-s_{p_t}(\boldsymbol{x}_t)\right)^T\frac{\partial}{\partial\varphi}\{s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y})\}\Big] \tag{A.8}$$
$$=\mathbb{E}_{\boldsymbol{x}_t\sim q_{\varphi,t}}\Big[d'\left(s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y})-s_{p_t}(\boldsymbol{x}_t)\right)^T\frac{\partial}{\partial\varphi}\{s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y})\}\Big] \tag{A.9}$$
$$=-\frac{\partial}{\partial\varphi}\mathbb{E}_{\substack{\boldsymbol{x}_0\sim q_\varphi(\cdot|\boldsymbol{y})\\ \boldsymbol{x}_t|\boldsymbol{x}_0\sim p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)}}\Big\{d'\left(s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y})-s_{p_t}(\boldsymbol{x}_t)\right)^T$$
$$\{s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y})-\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)\}\Big\} \tag{A.10}$$

Since φ does not participate in the derivative of the outer $d(\cdot)$, equation (A.8) is equivalent to (A.9). Moreover, φ does not appear in the differentiation with respect to $\boldsymbol{x}_t$ for the second term of (A.2):

$$\mathbb{E}_{\boldsymbol{x}_t\sim q_{\varphi,t}}\Big[\frac{\partial}{\partial\boldsymbol{x}_t}d(s_{q_{\varphi,t}}(\boldsymbol{x}_t|\boldsymbol{y})-s_{p_t}(\boldsymbol{x}_t))\frac{\partial\boldsymbol{x}_t}{\partial\varphi}\Big]$$
$$=\frac{\partial}{\partial\varphi}\mathbb{E}_{\substack{\boldsymbol{x}_0\sim q_\varphi(\cdot|\boldsymbol{y})\\ \boldsymbol{x}_t|\boldsymbol{x}_0\sim p_t(\boldsymbol{x}_t|\boldsymbol{x}_0)}}\Big[d\left(s_{q_{\mathrm{sg}[\varphi],t}}(\boldsymbol{x}_t|\boldsymbol{y})-s_{p_t}(\boldsymbol{x}_t)\right)\Big] \tag{A.11}$$

By combining (A.10) and (A.11) into (A.2), we get exactly the **Gradient Equivalence Theorem** stated in the text. This completes the proof.

B Additional Results on Fluorescent Microscopy

Beyond standard computational photography tasks, we also applied our method to real-world scientific imaging challenges in biomedicine and astronomy. We evaluated PPM on super-resolution fluorescent microscopy, a critical tool for visualizing subcellular structures. Here, the observation y is a wide-field microscope image (approximately 200 nm resolution), whose measurement model is

$$y=\mathrm{PSF}\circledast x+n, \tag{1}$$

where x is the underlying high-resolution fluorescence signal, PSF is the microscope point-spread function, and n is additive Gaussian noise. Accurately recovering x, along with precise uncertainty quantification, is essential for resolving the fine-grained dynamics of subcellular structures, such as organelles, and their interactions. In our experiment, we primarily benchmark PPM with RED-Diff. Both methods use the same diffusion prior pretrained on the BioSR dataset Aali et al. (2023), which comprises over 10,000 256×256 super-resolution images of diverse subcellular structures—microtubules, endoplasmic reticulum (ER), clathrin-coated pits (CCPs), and F-actin—captured with a structured illumination microscope (approximately 100 nm resolution).

Figure 7 demonstrates that PPM delivers higher-fidelity reconstructions than RED-Diff, faithfully rendering thin filaments in microtubules and the mesh-like ER. Quantitatively, PPM also achieves superior PSNR and SSIM scores. Crucially, PPM's uncertainty maps align with imaging physics: confidence peaks at structure centers and decreases toward blurred edges, revealing hollow structures. This behavior reflects the influence of microscope's PSF, which preserves feature presence while smearing precise boundaries. PPM accurately captures this boundary uncertainty, whereas RED-Diff's estimates fail to indicate edge ambiguity. These results underscore PPM's reliability for nanometer-scale biomedical imaging, where uncertainty quantification is indispensable.

Algorithm 1 Principled Posterior Matching (PPM) for Inverse Problems

Require: Fixed prior-score network s_θ; auxiliary score network s_ϕ, initialized from s_θ; schedule $\{\alpha_t, \sigma_t, g(t)\}$; time density ρ supported on $[t_{\min}, T]$; $d(\boldsymbol{r}) = \|\boldsymbol{r}\|_2^2$; data weight λ_{data}; auxiliary weight $\lambda_{\text{aux}}(t)$; observation-perturbation variance $h \geq 0$; total iterations.

Require: Mode: *Particle* for a fixed observation $\boldsymbol{y}$, or *Amortized* for observations drawn from the empirical distribution $\mathcal{D}_{\boldsymbol{y}}$.

if Mode is *Particle* **then**
  Initialize $\varphi \leftarrow \{\boldsymbol{\mu}_k\}_{k=1}^K$ from the measurement or a rough inverse.
else
  Initialize the reconstruction-network parameters φ, using the pretrained U-Net weights when architectures are compatible.
end if
while not converged **do**
  Sample from the current variational approximation:
  if Mode is *Particle* **then**
    Sample $k \sim \text{Uniform}\{1, \ldots, K\}$ and set $\boldsymbol{x}_0 \leftarrow \boldsymbol{\mu}_k$ for the fixed $\boldsymbol{y}$.
  else
    Sample $\boldsymbol{y} \sim \mathcal{D}_{\boldsymbol{y}}$ and $\boldsymbol{\eta} \sim \mathcal{N}(\boldsymbol{0}, h\boldsymbol{I})$.
    Set $\widetilde{\boldsymbol{y}} \leftarrow \boldsymbol{y} + \boldsymbol{\eta}$ and $\boldsymbol{x}_0 \leftarrow g_\varphi(\widetilde{\boldsymbol{y}})$.
  end if
  Sample $t \sim \rho(t)$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$.
  Set $\boldsymbol{x}_t \leftarrow \alpha_t \boldsymbol{x}_0 + \sigma_t \boldsymbol{\epsilon}$ and $\widehat{\boldsymbol{s}}_{\text{cond}} \leftarrow -\boldsymbol{\epsilon}/\sigma_t$.
  Update the auxiliary score:
  $\mathcal{L}_{\text{aux}} \leftarrow \lambda_{\text{aux}}(t) \|s_\phi(\boldsymbol{x}_t, t) - \widehat{\boldsymbol{s}}_{\text{cond}}\|_2^2$.
  $\phi \leftarrow \text{OptimizerStep}(\phi, \mathcal{L}_{\text{aux}})$.
  Update the variational parameters:
  Hold θ and ϕ fixed while retaining gradients through $\boldsymbol{x}_t$ and $\boldsymbol{x}_0$.
  $\boldsymbol{r}_t \leftarrow s_\phi(\boldsymbol{x}_t, t) - s_\theta(\boldsymbol{x}_t, t), \qquad a(t) \leftarrow g(t)^2/[2\rho(t)]$.
  $\mathcal{L}_{\text{prior}} \leftarrow a(t) \left[-d'(\boldsymbol{r}_t)^\top \{s_\phi(\boldsymbol{x}_t, t) - \widehat{\boldsymbol{s}}_{\text{cond}}\} + d(\boldsymbol{r}_t)\right]$.
  $\mathcal{L}_{\text{data}} \leftarrow \|\boldsymbol{y} - \mathcal{A}(\boldsymbol{x}_0)\|_2^2$.
  $\varphi \leftarrow \text{OptimizerStep}(\varphi, \mathcal{L}_{\text{prior}} + \lambda_{\text{data}} \mathcal{L}_{\text{data}})$.
end while
return In particle mode, $\{\boldsymbol{\mu}_k\}_{k=1}^K$ for the current observation; in amortized mode, the trained reconstruction network g_φ.

C Theoretical Analysis of Bias in Baselines

In this section, we provide a rigorous theoretical comparison between our proposed Principled Posterior Matching (PPM) and three leading categories of baselines. We demonstrate that while these methods offer practical utility, they all rely on biased approximations of the true posterior objective. In contrast, PPM minimizes the exact Kullback-Leibler (KL) divergence via an unbiased Fisher gradient estimator.

Bias in Optimization-based VI: RED-Diff and RLSD Optimization-based methods like RED-Diff Mardani et al. (2023) and RLSD Zilberstein et al. (2024) formulate the inverse problem as a variational optimization. However, they share a fundamental theoretical flaw in their handling of entropy.

- **Missing Exact Entropy (RED-Diff):** RED-Diff implicitly models the posterior as a Dirac delta distribution. This effectively removes the entropy term $H(q)$ from the KL divergence ($D_{\text{KL}} = \text{Energy} - \text{Entropy}$), collapsing the problem to **Maximum A Posteriori (MAP)** estimation. This theoretically explains its severe mode-seeking behavior. Consider the standard decomposition of the KL divergence between the variational distribution $q(\boldsymbol{x})$ and the true posterior $p(\boldsymbol{x}|\boldsymbol{y})$:

$$D_{\text{KL}}(q(\boldsymbol{x}) \| p(\boldsymbol{x}|\boldsymbol{y})) = \underbrace{-\mathbb{E}_q[\log p(\boldsymbol{x}|\boldsymbol{y})]}_{\text{Energy Term}} - \underbrace{H(q)}_{\text{Entropy Term}}, \tag{2}$$

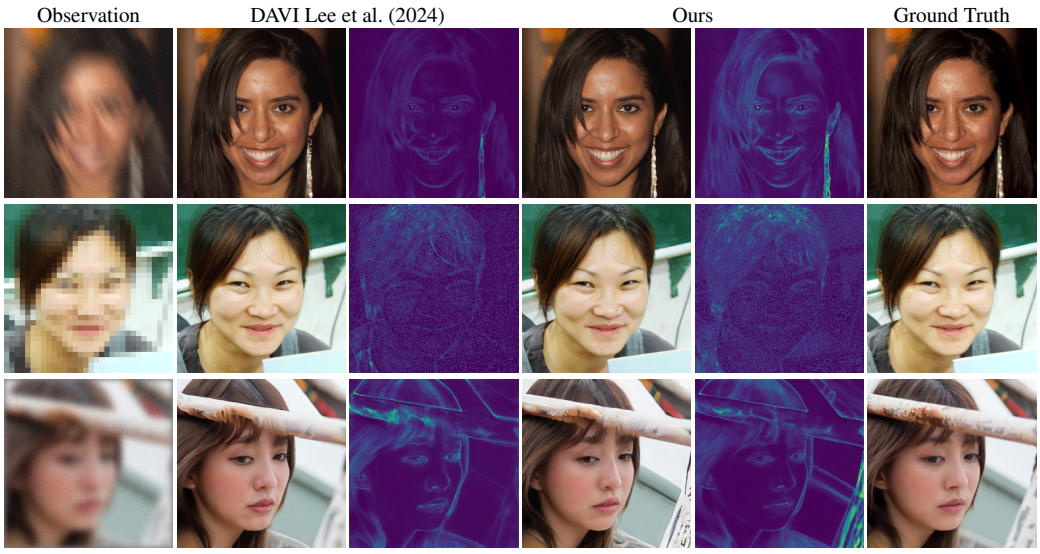

Figure 6: **Qualitative comparison of amortized inference on the FFHQ validation set.** We compare our PPM framework against the KL-divergence-based baseline, DAVI Lee et al. (2024). The rows correspond to three inverse problems: **Motion Deblurring** (top), $8\times$ **Super-Resolution** (middle), and **Gaussian Denoising** (bottom). For each task, we visualize the degraded observation, the reconstruction and pixel-wise uncertainty map from DAVI, the reconstruction and uncertainty map from our method, and the ground truth. Our method yields significantly sharper structural details (e.g., facial features and hair textures) and provides more informative uncertainty estimates that capture the rich diversity of local details, whereas the IKL-based approach tends to over-smooth the results.

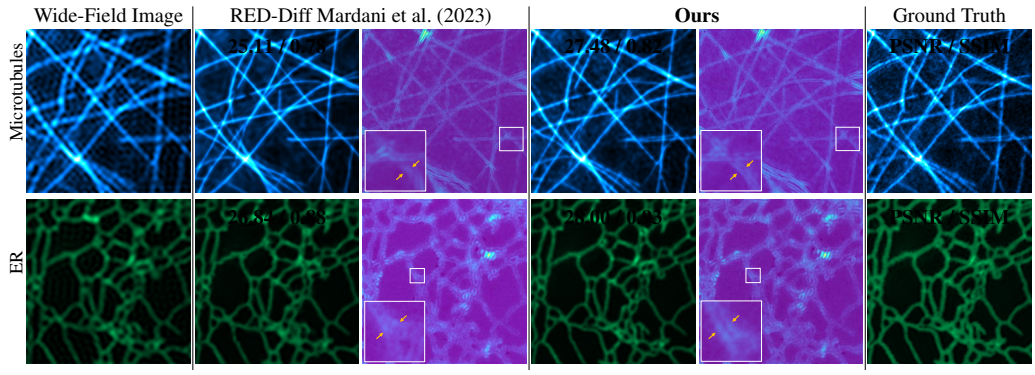

Figure 7: **Fluorescent super-resolution microscopic imaging results.** We compare our method with RED-Diff on microscopic images of Microtubules and ER samples. For each method, the reconstruction (with PSNR/SSIM scores) and its corresponding uncertainty map are reported. Our uncertainty maps accurately characterize the physical blur caused by the Point Spread Function in biological imaging: the uncertainty is lower at the center of the reconstructed structures and higher at the edges, effectively capturing the transition between the confirmed structures and the background.

where $H(q) = -\mathbb{E}_q[\log q(\boldsymbol{x})]$ denotes the differential entropy. RED-Diff and similar methods implicitly assume $q(\boldsymbol{x})$ is a point estimate, formalized as a Dirac delta distribution $q(\boldsymbol{x}) = \delta(\boldsymbol{x} - \boldsymbol{\mu})$. Under this assumption:

1. **Entropy Vanishes from Optimization:** The entropy of a Dirac distribution (or a Gaussian with fixed $\sigma \to 0$) is either undefined ($-\infty$) or a constant independent of the location parameter $\boldsymbol{\mu}$. Consequently, its gradient with respect to the optimization variable $\boldsymbol{\mu}$ is zero: $\nabla_{\boldsymbol{\mu}} H(\delta(\boldsymbol{x} - \boldsymbol{\mu})) = 0$.

2. **Energy Collapses to Log-Posterior:** The expectation in the energy term simplifies via the sifting property of the Dirac delta:

$$\mathbb{E}_{\boldsymbol{x}\sim\delta(\boldsymbol{x}-\boldsymbol{\mu})}[\log p(\boldsymbol{x}|\boldsymbol{y})] = \log p(\boldsymbol{\mu}|\boldsymbol{y}). \tag{3}$$

Substituting these back into the objective, the minimization problem becomes:

$$\min_{\boldsymbol{\mu}} D_{\mathrm{KL}}(q\|p) \iff \max_{\boldsymbol{\mu}} \log p(\boldsymbol{\mu}|\boldsymbol{y}). \tag{4}$$

This confirms that without the entropy term, minimizing the KL divergence is mathematically strictly equivalent to MAP estimation, which inherently seeks the single most probable mode rather than the full distribution.

- **Surrogate Entropy via Repulsion (RLSD):** While RLSD mitigates mode collapse by using multiple particles, it fails to optimize the true entropy of the variational distribution. Instead, it introduces an ad-hoc repulsive regularization (e.g., kernel-based repulsion) to force diversity. This repulsion acts as a heuristic proxy for entropy ($H(q)$) but does not correspond to the true score of the variational distribution $\nabla \log q_t$. Consequently, the resulting uncertainty is artificial and dependent on the repulsion hyperparameters, rather than being derived from the true posterior landscape like in PPM.

Bias in DPS: Laplacian Approximation of Likelihood Score Diffusion Posterior Sampling (DPS) Chung et al. (2022) approximates samples from $p(\boldsymbol{x}|\boldsymbol{y})$ by modifying the reverse diffusion score with a likelihood guidance term.

- **The Laplacian Approximation:** Since the likelihood $p(\boldsymbol{y}|\boldsymbol{x}_t)$ is intractable, DPS approximates it using a clean data estimate $\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t) = \mathbb{E}[\boldsymbol{x}_0|\boldsymbol{x}_t]$, assuming $\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{y}|\boldsymbol{x}_t) \approx \nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{y}|\hat{\boldsymbol{x}}_0(\boldsymbol{x}_t))$.
- **Approximation Error:** This violates Jensen's inequality by treating the expectation of the likelihood as the likelihood of the expectation. This introduces systematic score estimation error, particularly in early diffusion stages. Unlike PPM, DPS accumulates this error over the trajectory, leading to potential artifacts.

Bias in DAVI: Integral KL (IKL) vs. Exact KL DAVI Lee et al. (2024) employs the IKL divergence to train amortized generators. The detail analysis is demonstrated in Appendix D.

- **Independence Assumption:** IKL integrates marginal KL divergences: $\int \omega_t D_{\mathrm{KL}}(q_t\|p_t)\mathrm{d}t$. This treats each timestep independently and ignores the temporal dependencies (chain rule) of the diffusion process.
- **The Gap:** The true path KL divergence is **not equal** to the weighted sum of marginal KLs. By neglecting temporal correlations, IKL serves only as a heuristic proxy, resulting in biased posterior estimation compared to PPM's exact Fisher integration.

D Derivation of the bias in IKL

In this section, we formally prove that replacing the standard prior KL divergence term with the Integral KL (IKL) objective fundamentally alters the optimization target. We show that this substitution is mathematically equivalent to optimizing against a **distorted prior** $p'(\boldsymbol{x}) \propto p(\boldsymbol{x})^\beta$, where $\beta < 1$ is a scaling factor resulting from signal attenuation.

D.1 Problem Formulation

Standard Variational Inference (VI) minimizes the KL divergence between the variational posterior $q_\varphi(\boldsymbol{x}|\boldsymbol{y})$ and the true posterior $p(\boldsymbol{x}|\boldsymbol{y})$:

$$\mathcal{J}_{\mathrm{VI}} = \underbrace{\mathbb{E}_{q_\varphi}[-\log p(\boldsymbol{y}|\boldsymbol{x})]}_{\mathcal{L}_{\text{Likelihood}}} + \underbrace{D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}|\boldsymbol{y}) \parallel p(\boldsymbol{x}))}_{\mathcal{L}_{\text{Prior}}}. \tag{5}$$

IKL-based methods replace $\mathcal{L}_{\text{Prior}}$ with an integrated objective:

$$\mathcal{J}_{\text{Mod}} = \mathbb{E}_{q_\varphi}[-\log p(\boldsymbol{y}|\boldsymbol{x})] + \int_0^T \omega(t) D_{\mathrm{KL}}(q_t(\boldsymbol{x}_t|\boldsymbol{y}) \parallel p_t(\boldsymbol{x}_t))\mathrm{d}t. \tag{6}$$

D.2 Lemma: KL Contraction under Diffusion

To relate the divergence at time t to the divergence at $t = 0$, we analyze the behavior of the KL divergence under the diffusion process. While the exact relation depends on the complex geometry of the distributions, we derive the scaling law under a **Gaussian Mean Shift assumption**. This serves as a standard first-order approximation for local structural errors in inverse problems.

Assumption (VP-SDE & Mean Shift). We assume a Variance Preserving (VP) diffusion schedule where $\alpha_t^2+\sigma_t^2 = 1$. We model the prior $p(\boldsymbol{x})$ as a standard normal, and the variational approximation $q(\boldsymbol{x}|\boldsymbol{y})$ as a distribution with a structural error (mean shift) $\boldsymbol{\Delta}$:

$$p(\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}; \mathbf{0}, \mathbf{I}), \quad q(\boldsymbol{x}|\boldsymbol{y}) = \mathcal{N}(\boldsymbol{x}; \boldsymbol{\Delta}, \mathbf{I}). \tag{7}$$

Derivation. 1. **At time $t = 0$:** The KL divergence is:

$$D_{\mathrm{KL}}(q_0 \parallel p_0) = \frac{1}{2}\|\boldsymbol{\Delta}\|^2. \tag{8}$$

2. **At time t:** Applying the transition kernel $\boldsymbol{x}_t = \alpha_t \boldsymbol{x}_0 + \sigma_t \boldsymbol{\epsilon}$, the distributions become:

$$p_t(\boldsymbol{x}_t) = \mathcal{N}(\boldsymbol{x}_t; \mathbf{0}, (\alpha_t^2 + \sigma_t^2)\mathbf{I}) = \mathcal{N}(\boldsymbol{x}_t; \mathbf{0}, \mathbf{I}), \tag{9}$$

$$q_t(\boldsymbol{x}_t|\boldsymbol{y}) = \mathcal{N}(\boldsymbol{x}_t; \alpha_t\boldsymbol{\Delta}, (\alpha_t^2 + \sigma_t^2)\mathbf{I}) = \mathcal{N}(\boldsymbol{x}_t; \alpha_t\boldsymbol{\Delta}, \mathbf{I}). \tag{10}$$

The KL divergence between these diffused distributions is:

$$D_{\mathrm{KL}}(q_t \parallel p_t) = \frac{1}{2}\|\alpha_t\boldsymbol{\Delta} - \mathbf{0}\|^2 = \alpha_t^2 \left(\frac{1}{2}\|\boldsymbol{\Delta}\|^2\right). \tag{11}$$

Result. We obtain the quadratic scaling relationship:

$$D_{\mathrm{KL}}(q_t(\boldsymbol{x}_t|\boldsymbol{y}) \parallel p_t(\boldsymbol{x}_t)) \approx \alpha_t^2 \cdot D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}|\boldsymbol{y}) \parallel p(\boldsymbol{x})). \tag{12}$$

D.3 Derivation of the Implicit Prior

Substituting the scaling law (Eq. 12) into the IKL integral, we rewrite the modified objective:

$$\begin{aligned} \mathcal{J}_{\mathrm{IKL}} &\approx \int_0^T \omega(t) \cdot \alpha_t^2 \cdot D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}|\boldsymbol{y}) \parallel p(\boldsymbol{x}))\mathrm{d}t \\ &= \underbrace{\left(\int_0^T \omega(t)\alpha_t^2 \mathrm{d}t\right)}_{\beta} \cdot D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}|\boldsymbol{y}) \parallel p(\boldsymbol{x})). \end{aligned} \tag{13}$$

We define β as the **Effective Prior Weight**. The total loss becomes:

$$\mathcal{J}_{\mathrm{Mod}} \approx \mathbb{E}_{q_\varphi}[-\log p(\boldsymbol{y}|\boldsymbol{x})] + \beta D_{\mathrm{KL}}(q_\varphi(\boldsymbol{x}|\boldsymbol{y}) \parallel p(\boldsymbol{x})). \tag{14}$$

We seek the form of the **Implicit Posterior** $p'(\boldsymbol{x}|\boldsymbol{y})$ such that the objective corresponds to a standard VI form $\beta D_{\mathrm{KL}}(q_\varphi \| p')$. Expanding the terms:

$$\begin{aligned} \mathcal{J}_{\mathrm{Mod}} &= \beta \left(\frac{1}{\beta}\mathbb{E}_{q_\varphi}[-\log p(\boldsymbol{y}|\boldsymbol{x})] + \mathbb{E}_{q_\varphi}[\log q_\varphi - \log p(\boldsymbol{x})]\right) \\ &= \beta \mathbb{E}_{q_\varphi}\left[\log q_\varphi(\boldsymbol{x}|\boldsymbol{y}) - \left(\frac{1}{\beta}\log p(\boldsymbol{y}|\boldsymbol{x}) + \log p(\boldsymbol{x})\right)\right]. \end{aligned} \tag{15}$$

The term inside the parenthesis corresponds to the log-density of the implicit target p':

$$\log p'(\boldsymbol{x}|\boldsymbol{y}) = \frac{1}{\beta}\log p(\boldsymbol{y}|\boldsymbol{x}) + \log p(\boldsymbol{x}) + C. \tag{16}$$

D.4 CONCLUSION

Exponentiating the result, we find that minimizing the IKL objective targets a distorted posterior:

$$p'(\boldsymbol{x}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{x})^{\frac{1}{\beta}} \cdot p(\boldsymbol{x}). \tag{17}$$

Equivalently, by factoring out the exponent $1/\beta$, this implies an optimization against a **Distorted Prior** $p'(\boldsymbol{x})$:

$$p'(\boldsymbol{x}|\boldsymbol{y}) \propto \left(p(\boldsymbol{y}|\boldsymbol{x}) \cdot p(\boldsymbol{x})^{\beta}\right)^{\frac{1}{\beta}} \implies p'(\boldsymbol{x}) \propto p(\boldsymbol{x})^{\beta}. \tag{18}$$

Since α_t^2 decays to 0, we typically have $\beta < 1$. This means the effective prior $p(\boldsymbol{x})^{\beta}$ is a **flattened (high-temperature) version** of the true prior.