Norm-Controlled Likelihood Guidance for Diffusion-based Inverse Solver

Anonymous Authors[1]

Abstract

Diffusion-based inverse solvers approximate the posterior by combining a pretrained diffusion prior with an approximate likelihood guidance term. Across diverse host solvers, we identify a consistent diagnostic signal: for the same image and inverse task, different noise realizations lead to different reconstructions, and larger likelihood-guidance norms reliably predict worse perceptual quality. This observation motivates us to propose the *likelihood-score norm* as a general, quantitative error proxy. To leverage this proxy, we propose **Norm-Controlled Likelihood Guidance**, a set of three effective modules that steer the sampling trajectory toward smaller likelihood norms while remaining on the diffusion manifold. Theoretically, we justify our motivation by proving that under a tractable multimodal model, smaller likelihood norms can reliably imply smaller score approximation errors. Experimentally, we demonstrate that the proposed method can acheive consistent improvements across multiple diffusion-based solvers on FFHQ and ImageNet-256 datasets. Ablation studies and cost analyses further validate the design and show favorable quality–compute trade-offs.

1. Introduction

Many imaging and scientific pipelines are naturally posed as *inverse problems*: recovering an unknown signal x_0 from indirect and noisy measurements $y = \mathcal{A}(x_0)+\varepsilon$, where $\mathcal{A}$ is a forward operator modeling the measurement process and ε denotes additive noise. Common examples include inpainting, super-resolution, deblurring, MRI/CT reconstruction, and other linear or mildly non-linear sensing tasks. These problems are often *ill-posed* due to the non-invertibility of $\mathcal{A}$ and the presence of noise, meaning multiple or even infinitely many solutions may explain the observed data (Arridge et al., 2019; Stuart, 2010; Tikhonov, 1963). To this end, solving inverse problems has been typically formulated as a sampling task: drawing x_0 from the posterior distribution $p(x|y)$. A principled way to formalize this is through the Bayesian framework. In this view, prior knowledge about the unknown signal is captured by a distribution $p(x)$, while the measurement process is modeled by a likelihood $p(y|x)$. Combining these two components, the solution involves sampling from the posterior distribution: $\pi(x|y) \propto p(x)\, p(y|x)$. Although conceptually straightforward, this approach is challenging in practice: priors are often unknown and high-dimensional and is difficult to be explicitly modeled by data due to its high-dimensional and structurally complex nature.

To address this challenge, diffusion models provide a powerful framework for sampling from complex high-dimensional distributions via score-based dynamics (Ho et al., 2020; Song et al., 2021; Nichol & Dhariwal, 2021; Rombach et al., 2022). In diffusion-based inverse sampling, the goal is to sample from the posterior $\pi(x|y)$ by constructing a reverse-time process over a family of conditional distributions $\pi_t(x_t|y) \propto p_t(x_t)\, p(y|x_t)$, where t denotes the noise level. As $t \to 0$, the marginal $\pi_t(x_t|y)$ converges to the target posterior $\pi(x|y)$, enabling sampling through a continuous denoising trajectory. Then, the key of this approach is to learn the score of the marginal distribution: $\nabla_{x_t} \log \pi_t(x_t|y) = \nabla_{x_t} \log p_t(x_t) + \nabla_{x_t} \log p(y|x_t)$, where the prior score $\nabla_{x_t} \log p_t(x_t)$ is provided by a pretrained diffusion model through its learned score field $s_\theta(x_t, t) \approx \nabla_{x_t} \log p_t(x_t)$. In contrast, the difficulty lies in the computation of the likelihood score $\nabla_{x_t} \log p(y|x_t)$, which is largely intractable, especially for non-linear or ill-posed forward operators.

Existing diffusion inverse solvers often address this challenge by approximating the intractable likelihood term $p(y|x_t)$ with $p(y|\mathbb{E}[x_0|x_t])$, which leverages the likelihood at the data space ($t = 0$) as an approximation. This plug-in approach is exemplified by DPS (Chung et al., 2023), and many recent methods can be viewed as its variants (Song et al., 2023a; Tang et al., 2024; Huang et al., 2024b;a). The posterior mean $\mathbb{E}[x_0|x_t]$ is typically approximated via Tweedie's formula using the prior score $\nabla_{x_t} \log p_t(x_t)$. Concretely, DPS approximates the conditional posterior as:$\pi_t(x_t|y) \propto p_t(x_t) \cdot p\big(y|\hat{x}_0(x_t)\big)$, where

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the FoGen Workshop at ICML 2026. Do not distribute.

$\hat{x}_0(x_t) = x_t + \sigma_t^2 \nabla_{x_t} \log p_t(x_t)$ is the Tweedie-style estimate of $\mathbb{E}[x_0|x_t]$. Although this approach has proven effective, it introduces approximation errors that increase with the noise level t, often resulting in degraded reconstruction quality. Some recent works have attempted to mitigate this limitation, focusing on refining when and how the likelihood $p(y|x_t)$ is incorporated. For instance, DCPS introduces intermediate potentials to reduce long-step bias (Janati et al., 2024); MGPS employs midpoint guidance to improve per-step likelihood estimation (Moufad et al., 2025).

Despite their empirical success, these methods rely on heuristic algorithmic refinements that require task-specific tuning. A more fundamental, broadly applicable technique is needed, one that improves performance independently of the underlying algorithm. Ideally, a quantitative signal directly correlated with likelihood-approximation error would offer a principled, model-agnostic tool for guiding sampling and analysis. In this work, we demonstrate that the likelihood-score norm effectively tracks reconstruction quality, transforming likelihood guidance from a fixed heuristic into a measurable, general-purpose signal for runtime control and informed sampling. Our contributions are highlighted as follows:

- We introduce the *low-noise likelihood-score norm* as an in-loop diagnostic tool that consistently predicts reconstruction quality across datasets, tasks, and host solvers. Furthermore, we provide a rigorous theoretical analysis to validate the strong correlation between the likelihood-score norm and the true score-approximation error.
- Guided by this diagnostic/certificate logic, we propose Norm-Controlled Likelihood Guidance (NCLG) that exploits the intrinsic stochasticity of diffusion sampling to *refine toward* lower-norm trajectories without modifying the host prior transition: Restart Sampling (RS) provides a coarse restart into a favorable basin, Likelihood-Norm Stochastic Selection (NSS) performs repeated candidate selection among stochastic branches using predicted low-noise norm aggregates, and Liklihood-Norm Projection (LNP) applies a capped Langevin refinement that reduces instantaneous norms while retaining the diffusion prior term.
- We validate the framework on FFHQ and ImageNet inverse problems, showing consistent improvements over recent baselines with a favorable quality–compute trade-off; ablations further isolate which design choices are critical, including the importance of measuring/controlling norms in low-noise windows and the complementary roles of restart, selection, and refinement.

2. Related Work

Diffusion priors for inverse problems. Diffusion and score-based models learn a time-indexed score field that enables high-quality generation and serves as a strong implicit prior for inverse problems (Ho et al., 2020; Song et al., 2021; Nichol & Dhariwal, 2021; Rombach et al., 2022). Diffusion inverse solvers typically sample from an approximate posterior by combining the learned prior score with a task-dependent data-consistency (likelihood) term. Existing approaches can be broadly grouped into: *(i)* closed-form or projection-based solvers exploiting linear-Gaussian structure, e.g., DDRM/DDNM and ΠGDM (Kawar et al., 2022; Wang et al., 2023; Song et al., 2023b); *(ii)* general guidance methods that backpropagate an observation loss through a denoiser prediction $\hat{x}_0(x_t, t)$ (DPS and related variants) (Chung et al., 2023; Song et al., 2023a; Tang et al., 2024; Huang et al., 2024b; Bastek et al., 2025; Cheng et al., 2025; Jacobsen et al., 2024; Huang et al., 2024a). Some work apply more explicitly Bayesian samplers, including variational and Monte Carlo/SMC/MCMC constructions, which can improve posterior faithfulness at higher compute cost (Mardani et al., 2024; Dou & Song, 2024; Cardoso et al., 2024; Wu et al., 2024). Our method builds on guidance-based approaches while incorporating MCMC/MC modules that can be seamlessly integrated into both sequential samplers and broader Bayesian frameworks.

Likelihood approximations and where errors enter. For general operators and noise models, the likelihood score along the diffusion trajectory is intractable, so inverse solvers rely on approximations often via a Tweedie/denoiser plug-in that reduces likelihood evaluation to the image domain through $\hat{x}_0(x_t, t)$ (Chung et al., 2023; Boys et al., 2024). A key limitation of such plug-ins is that their accuracy is time-dependent, with larger diffusion noise generally inducing larger approximation error; this motivates methods that modify *how* and *when* the likelihood is injected. DCPS introduces intermediate potentials to reduce accumulated bias from long reverse-time jumps, and MGPS uses midpoint guidance to improve likelihood handling within each step (Janati et al., 2024; Moufad et al., 2025). Despite improvements, both methods retain residual approximation error. Our work extends this line by proposing an external wrapper that further reduces likelihood error through an orthogonal mechanism.

Diagnostics and alignment signals. Several works connect sample quality to score/noise alignment and provide post-hoc or trajectory-level diagnostics for diffusion-based conditional generation. Misalignment maps quantify conditional–unconditional score discrepancy and correlate it with perceptual recognizability (Boutin et al., 2023). CAS performs sample-level self-rejection using a condition-alignment score derived from probability-flow ODE quantities (Hong et al., 2024). Inversion-behavior studies relate initial noise, generated samples, and inversion biases that affect controllability (Staniszewski et al., 2025). While these are primarily post-hoc or sample-level, our approach introduces a novel

Table 1. **Spearman correlations between LPIPS and windowed likelihood-norm aggregates.** Each entry reports Spearman correlation ρ_S over two noise windows: Low-noise ($\mathcal{W}_{\text{low}}$) and High-noise ($\mathcal{W}_{\text{high}}$). Left: FFHQ. Right: ImageNet. Pearson analogs in Appendix Table 9.

Task	σ_y	FFHQ		ImageNet	
		Low-noise	High-noise	Low-noise	High-noise
SR 4×	0.05	**0.33**	0.19	**0.54**	0.09
	0.3	**0.28**	0.09	**0.31**	0.07
Center	0.05	**0.38**	0.14	**0.69**	0.14
	0.3	**0.33**	0.16	**0.78**	0.14
Half	0.05	**0.47**	0.19	**0.52**	0.12
	0.3	**0.36**	0.17	**0.48**	0.12

score-level alignment signal that is directly usable during sampling, as demonstrated next in Section 3.

3. Likelihood-Score Norm as Error Proxy

In this section, we present the core methodological insight of our work: leveraging the norm of the estimated likelihood score as a robust indicator of trajectory quality in diffusion-based posterior sampling. Since the exact likelihood score $\nabla_{x_t} \log p(y|x_t)$ is typically intractable, especially under non-linear forward operators, practical diffusion inverse solvers rely on differentiable proxies $\hat{g}_t(x_t)$ to approximate it. Formally, we define the instantaneous *likelihood-score norm* as $\ell_t := \|\hat{g}_t(x_t)\|_2$. Here, $\hat{g}_t(x_t)$ denotes the likelihood-score proxy provided by the host solver. For example, in DPS (Chung et al., 2023), $\hat{g}_t(x_t)$ is constructed by first computing a denoised estimate $\hat{x}_0(x_t) \approx \mathbb{E}[x_0|x_t]$ via the Tweedie formula and then evaluating the measurement residual:

$$\hat{g}_t(x_t) = \nabla_{x_t} \log p(y|\hat{x}_0(x_t))$$

This form, or its variants, is commonly used across solvers such as DCPS (Janati et al., 2024) and MGPS (Moufad et al., 2025) to guide posterior sampling when the true likelihood is unavailable.

For a sampling window defined by a set of noise levels $\mathcal{W} = \{t_1, t_2, \ldots, t_K\}$, we define the *aggregated norm* as:

$$\text{Norm}(\mathcal{W}) := \sum_{t \in \mathcal{W}} \|\hat{g}_t(x_t)\|_2^2. \tag{1}$$

In what follows, we will use this likelihood-score norm as an error proxy for score-based posterior diffusion solvers. we will validate this by showing its strong correlation with reconstruction quality, identifying where it is most informative, and supporting this with theoretical analysis.

3.1. Likelihood-Score Norm Reflects Generation Quality

We begin by investigating the correlation between reconstruction quality and the likelihood-score norm. Specifically, we first analyze the instantaneous norm ℓ_t at a fixed noise level $t = 0.3$ across multiple reconstruction runs initialized with different noise seeds, utilizing various host solvers (DPS, DCPS, MGPS). To identify the regime where this signal is most informative along the diffusion trajectory, we further define two distinct noise windows: a low-noise window $\mathcal{W}_{\text{low}} = (0, 0.3]$ and a high-noise window $\mathcal{W}_{\text{high}} = (0.7, 1]$, and compute the aggregated norm $\text{Norm}(\mathcal{W})$ across diverse tasks and datasets. We report the Spearman (ρ_{S}) and Pearson (ρ_{P}) correlations between these norms and the LPIPS perceptual metric. The results are summarized in Figure 1 and Table 1.

Result 1: Positive correlation. As shown in Figure 1, the likelihood-score norm is consistently and positively correlated with LPIPS all tested solvers. Since higher LPIPS values denote lower quality, higher norms reliably predict degraded perceptual fidelity, establishing the norm as an effective diagnostic signal.

Result 2: Low-Noise norms are more predictive. As shown in Table 1, while the aggregated likelihood-score norm still demonstrate positive correlation with LPIPS across various tasks and datasets, the correlations is significantly stronger in the low-noise window ($t \in \mathcal{W}_{\text{low}}$) compared to the high-noise window ($t \in \mathcal{W}_{\text{high}}$). This confirms that the likelihood-score norm are more informative when measured in the low-noise regime

These findings motivate a clear design principle: *leverage likelihood-score norm in the low-noise regime to guide sampling*. Building on this, the next section introduces a sampling framework that explicitly controls the likelihood norm during these critical steps to improve reconstruction fidelity.

3.2. Theoretical Demonstration

To demonstrate our empirical finding and support the utility of the proposed likelihood score in subsection 3.1, we further deliver the theoretical analysis on the mixture-of-Gaussians (MoG) data model, a synthetic setting that is widely deployed in many existing theoretical works (Karan et al., 2025; Liang et al., 2025; Chidambaram et al., 2024; Chen et al., 2024; Shah et al., 2023; Chen et al., 2023). We first characterize the conditional-score error $e(x)$ between the DPS-style plug-in approximation and the ground truth in Lemma 3.1. Building on this, we establish a rigorous bound on $e(x)$ controlled by the likelihood-score norm in the low-noise regime in Theorem 3.2.

Consider a Mixture-of-Gaussians (MoG) data model where the Gaussian corruption and Gaussian observation processes are defined as follows:

$$x = x_0 + \sigma_t z, \quad y = x_0 + \sigma_y \varepsilon, \quad z, \varepsilon \sim \mathcal{N}(0, I). \tag{2}$$

Here, the clean signal lies in a discrete support $x_0 \in$

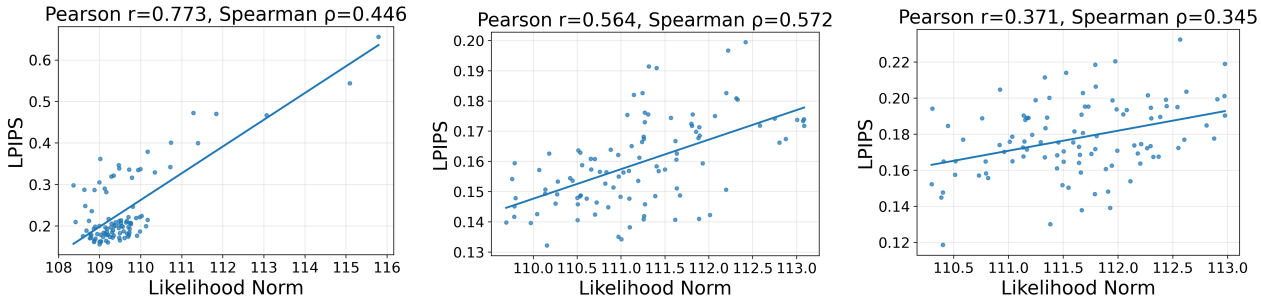

Figure 1. **Within-instance scatter: LPIPS vs. likelihood norm at $t = 0.3$.** Each point is one run with a different initial noise for the *same* image and task in FFHQ. Left: DPS (Chung et al., 2023); Middle: DCPS (Janati et al., 2024); Right: MGPS (Moufad et al., 2025). Larger likelihood norms are associated with larger LPIPS (worse). It is clear that the likelihood norm is positively correlated with the LPIPS score, providing evidence of using the likelihood norm to selectively perform data generation with better quality.

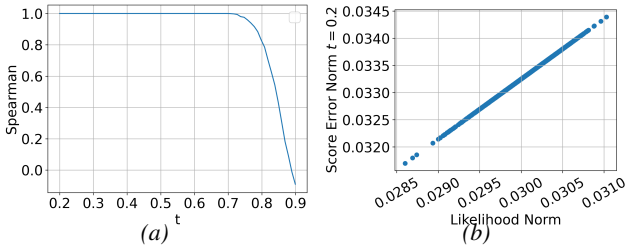

Figure 2. **MOG toy study: likelihood-norm vs. conditional-score error.** (a): Spearman correlation between $\|\hat{g}_t\|_2$ and $\|\hat{s}_t - s_t^\star\|_2$ as a function of time t. (b): Scatter plots of $\|\hat{g}_t\|_2$ and $\|\hat{s}_t - s_t^\star\|_2$ at $t = 0.2$.

$\{\mu^{(1)}, \ldots, \mu^{(K)}\} \subset \mathbb{R}^d$ with prior probabilities $\mathbb{P}(x_0 = \mu^{(k)}) = \pi_k$. We define the posterior means as $\bar{\mu}(x) := \mathbb{E}[x_0|x]$ and $\bar{\mu}(x, y) := \mathbb{E}[x_0|x, y]$. We now derive the true conditional score and the DPS-estimated score, explicitly computing their approximation error:

Lemma 3.1. *The* true *conditional score at noise level t satisfies $s^\star(x) := \nabla_x \log p_t(x|y) = \frac{1}{\sigma_t^2}\big(\bar{\mu}(x, y) - x\big)$. A stylized DPS-style plug-in conditional score ("predict $\hat{x}_0$ via Tweedie, then apply likelihood") is*

$$\tilde{s}(x) := \frac{1}{\sigma_t^2}\big(\bar{\mu}(x) - x\big) + s^{\text{lik}}(x), \tag{3}$$

$$s^{\text{lik}}(x) := \nabla_x \log p(y|\bar{\mu}(x)) .$$

Then the conditional-score error $e(x) := \tilde{s}(x) - s^\star(x)$ reads

$$e(x) = \frac{1}{\sigma_t^2}\big(\bar{\mu}(x) - \bar{\mu}(x, y)\big) + s^{\text{lik}}(x). \tag{4}$$

Lemma 3.1 decomposes the conditional-score error into two distinct components (Eq. 4): the *mean gap* term $\bar{\mu}(x) - \bar{\mu}(x, y)$ and the likelihood-score proxy $s^{\text{lik}}(x)$. While the mean gap is generally intractable, in the low-noise regime, this term becomes non-dominant and effectively negligible. As illustrated in Figure 2, in the low-noise regime ($t < 0.3$), the likelihood-score norm $\|s^{\text{lik}}(x)\|_2$ and the total score error $\|e(x)\|_2$ exhibit a strong affine relationship, with a Spearman coefficient near 1. Theorem 3.2 below formalizes this observation, establishing a rigorous bound on $e(x)$ where the mean gap manifests as an ambiguity floor controlled by the geometry of the mixture:

Theorem 3.2 (Likelihood-norm certificate (low-noise MoG); proof in Appendix E)**.** *Fix a reference mode $\mu^\star = \mu^{(k^\star)}$ and define the separation radii $\Delta_{\min} := \min_{j \neq k^\star} \|\mu^{(j)} - \mu^\star\|_2$ and $\Delta_{\max} := \max_{j \neq k^\star} \|\mu^{(j)} - \mu^\star\|_2$. Work on the measurement-consistency event $\|y - \mu^\star\|_2 \leq r_y$ with $r_y \leq \Delta_{\min}/4$. Assume we are in the* low-noise *regime in the sense that the diffusion scale at this time satisfies $\sigma_t \leq \frac{1}{8}\Delta_{\min}$. Then there exist explicit constants and thresholds τ, such that for any x satisfying $\|s^{\text{lik}}(x)\|_2 \leq \tau$, the conditional-score error $e(x) := \tilde{s}(x) - s^\star(x)$ obeys*

$$\|e(x)\|_2 \leq \delta_{\text{floor}}^{\text{lik}} + \|s^{\text{lik}}(x)\|_2, \tag{5}$$

where in the low-noise regime, the ambiguity floor $\delta_{\text{floor}}^{\text{lik}}$ can be written as $\mathcal{O}\left(\exp(-\Delta_{\min}^2/(4\sigma_y^2))\right)$, and $\|s^{\text{lik}}(x)\|_2$ can be written as $\mathcal{O}(1/\sigma_y^2)$.

Theorem 3.2 establishes that the likelihood-score norm $\|s^{\text{lik}}(x)\|_2$ effectively upper-bounds the true conditional-score error $\|e(x)\|_2$, subject to a baseline floor $\delta_{\text{floor}}^{\text{lik}}$. The decomposition in the theorem reveals the distinct behavior of these terms in the low-noise regime:

- **Vanishing Ambiguity Floor ($\delta_{\text{floor}}^{\text{lik}}$):** The floor term scales exponentially as $\mathcal{O}\big(\exp(-\Delta_{\min}^2/(4\sigma_y^2))\big)$. This exponential decay captures the rapid vanishing of mode ambiguity once the diffusion noise σ_t drops below the separation threshold.
- **Dominant Controllable Proxy ($\|s^{\text{lik}}(x)\|_2$):** In contrast, the likelihood norm is dominated by the term scaling as $\mathcal{O}(1/\sigma_y^2)$, which remains significant even in the low-noise limit. The inequality effectively compares the vanishing exponential part against the significant inverse-variance part. Consequently, the bound tightens to $\|e(x)\|_2 \approx \|s^{\text{lik}}(x)\|_2$, theoretically justifying the likelihood norm as a direct proxy for the intractable score error, matching the near-linear alignment observed at low t in Figure 2.

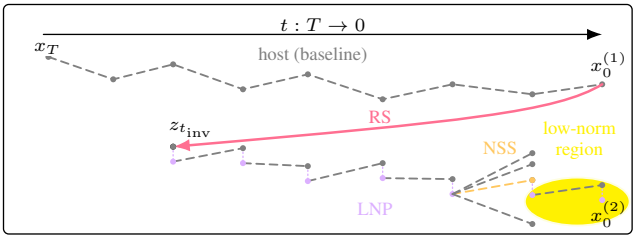

Figure 3. **Conceptual illustration of likelihood-norm-guided inverse sampling.** A baseline host trajectory (gray dashed) may not reach a low likelihood-norm region. Our method restarts via RS by partially inverting a first-pass sample to an earlier latent, then applies per-step LNP and NSS to steer the trajectory toward lower likelihood norms.

4. Norm-Controlled Likelihood Guidance

In this section, we introduce a plug-in framework that leverages the certificate signal in Theorem 3.2 within existing solvers. In specific, : *we exploit a shared degree of freedom across diffusion inverse solvers to search for or refine toward lower-norm sampling trajectories while keeping the host sampler unchanged.* This degree of freedom arises from intrinsic stochasticity, such as noise injection, SDE sampling, or randomized correctors.

Our algorithm framework is summarized in Alg. 1, which exhibits three key modules: uses a two-pass restart to partially invert a first-pass sample x_0 back to an earlier latent x_t (with $t > 0$), initializing a coarse, lower-likelihood basin where subsequent control becomes effective; within this basin, Likelihood-Norm Stochastic Selection (NSS) *reactively selects* lower-norm trajectory among the host's stochastic paths; and Low-Norm Projection (LNP) *actively steers* the state via a short capped Langevin refinement that reduces the likelihood norm. Full module pseudocode is deferred to Appendix G.9.

4.1. Local norm control via refinement and selection

We first introduce two modules locally control the likelihood norm: Likelihood-Norm Projection (LNP) and Likelihood-Norm Stochastic Selection (NSS). In Fig. 3, NSS corresponds to choosing the dashed continuation that is predicted to enter a lower-norm region, and LNP corresponds to the short purple corrective moves applied along the chosen path.

Likelihood-Norm Projection via Normalized Langevin (LNP). LNP is motivated by the naive idea of reducing the likelihood norm via gradient descent directly on x_t. However, such an update can drive x_t away from the diffusion marginal $p(x_t)$, leading to inaccurate sampling. To avoid this issue, we adopt corresponding Gibbs-style formulation (Geman & Hwang, 1986; Hajek, 1988; Erdogdu et al., 2018) that jointly enforces diffusion-marginal consistency and low likelihood norm, and sample from it using Langevin dynamics. Specifically, at timestep t,

Algorithm 1 Norm-Controlled Likelihood Guidance

Require: Timesteps $t = T, T-1, \ldots, 1$.
Require: BASESTEP: host reverse step $(x_t,\ t \to t-1; \eta_t)$.
Require: $x_0^{(1)}$: First-round final sample using BASESTEP.
Require: PRIOREST(x_t, t): Estimator of $\nabla \log p(x_t)$.
Require: $\hat{g}_t(x_t)$: Estimator of $\nabla \log p(y|x_t)$.
Require: RS inverse encoder $\mathcal{S}^{-1}$, noise η_{inv}: Inverse sampling from x_0 to $x_{\eta_{\text{inv}}}$ using prior score PRIOREST.
Require: NSS norm probe runner $\mathcal{S}_{\text{probe}}$, N, $\mathcal{T}_{\text{sel}}$, window $\mathcal{W}$, probe noise η_{sol}. The Detailed algorithm is shown in Alg. 4.
Require: LNP $\gamma, \varepsilon, R_{\text{proj}}, \eta_{\text{mcmc}}, \tau_{\text{start}}$ ▷ one ordinary host run

$t_{\text{inv}} \leftarrow \lceil \tau_{\text{inv}} T \rceil$
$x_{t_{\text{inv}}} \leftarrow \mathcal{S}^{-1}(x_0^{(1)}; \eta_{\text{inv}})$ ▷ truncate at t_{inv}
$t_0 \leftarrow t_{\text{inv}}$
for $t = t_0, t_0-1, \ldots, 1$ **do**
  if LNP enabled **and** $t \le \lfloor \tau_{\text{start}} T \rfloor$ **then**
    for $r = 1$ to R_{proj} **do**
      $\hat{s}_t \leftarrow$ PRIORSCOREEST(x_t, t)
      $\alpha_t \leftarrow \gamma / (\|\hat{g}_t(x_t)\|_2 + \varepsilon)$
      Sample $\xi \sim \mathcal{N}(0, I)$
      $x_t \leftarrow x_t + \eta_{\text{mcmc}}(\hat{s}_t + \alpha_t \hat{g}_t) + \sqrt{2\eta_{\text{mcmc}}}\,\xi$
    end for
  end if
  if then$t \in \mathcal{T}_{\text{sel}}$
    Sample candidate noises $\{\eta^{(n)}\}_{n=1}^{N} \sim \mathcal{N}(0, I)$
    $\{S^{(n)}\}_{n=1}^{N} \leftarrow \mathcal{S}_{\text{probe}}(x_t, t, \mathcal{W}, \eta_{\text{sol}}, \{\eta^{(n)}\})$
    $\eta_t \leftarrow \eta^{(\arg\min_n S^{(n)})}$
  else
    Sample $\eta_t \sim \mathcal{N}(0, I)$
  end if
  $x_{t-1} \leftarrow$ BASESTEP$(x_t,\ t \to t-1; \eta_t)$
  $x_t \leftarrow x_{t-1}$
end for
return x_0

we define a tempered approximate posterior $q_t(x_t) \propto p_t(x_t) \cdot \text{LIKEEST}(x_t, t;\ y, \mathcal{A})^{\alpha_t}$, where $p_t(x_t)$ is the diffusion marginal, $\text{LIKEEST}(x_t, t;\ y, \mathcal{A})$ is a host proxy for $p_t(y|x_t)$ where $\nabla \log \text{LIKEEST}(x_t, t;\ y, \mathcal{A}) = \hat{g}_t(x_t)$, and α_t is a dynamic scaling factor that caps the likelihood term's magnitude by a user-defined projection gain γ, ensuring $|\alpha_t \hat{g}_t(x_t)| \le \gamma$; see Appendix G.7 for details. To sampling toward q_t, we apply R_{proj} unadjusted Langevin dynamics:

$$x_t \leftarrow x_t + \eta_{\text{mcmc}}(\hat{s}_t + \alpha_t \hat{g}_t) + \sqrt{2\eta_{\text{mcmc}}}\,\xi,$$

with $\xi \sim \mathcal{N}(0, I)$ and $\eta_{\text{mcmc}} := 1/\gamma$. The parameter γ controls the trade-off between likelihood-norm reduction and preservation of the diffusion marginal. In practice, we set γ to a moderate value (default $\gamma = 50$) to retain the influence of the prior score $\hat{s}_t$, ensuring that x_t remains close to the diffusion marginal $p(x_t)$ and supports accurate downstream sampling. By default we apply LNP only in the low-noise portion (small t), where the likelihood-norm signal is most predictive (Section 3); see Algorithm 1 and 3.

Likelihood-Norm Stochastic Selection (NSS). NSS exploits the intrinsic stochasticity of diffusion samplers (where each denoising step typically injects Gaussian noise) by selecting noise that yield lower predicted low-noise norm

aggregates. Specifically, NSS proposes N noise candidates $\{\eta^{(n)}\}_{n=1}^{N}$, rolls out each corresponding trajectory, and assigns a score based on the *low-noise* norm aggregate:

$$S^{(n)} := \sum_{\tau \in \mathcal{W}_t} \|\hat{g}_\tau^{(n)}\|_2^2, \qquad \mathcal{W}_t := \{\, \tau \in \mathcal{W}_{\text{low}} : \tau \le t \,\},$$

where $\mathcal{W}_{\text{low}}$ denotes a noise regime where the likelihood-score norm is most predictive, as discussed in Section 3. The candidate with the smallest score is selected, and the corresponding noise $\eta^{(\arg\min_n S^{(n)})}$ is used for the actual host solver update. Intuitively, NSS transforms local stochasticity into a selective mechanism for controlling the likelihood norm during sampling.

LNP/NSS are local mechanisms that benefit from better global initialization. We highlight that LNP/NSS is a *local* correction on x_t. Starting from an arbitrary high-noise state, the posterior over x_0 can be highly multi-modal, and local MCMC/Selection can be inefficient when the initialization is poor. This motivates a complementary mechanism that constructs a more favorable restart state for LNP/NSS and we show it in the next subsection.

4.2. Global low-norm basin construction via restart

We now introduce Restart Sampling (RS), a host-agnostic *global* module that coarsely reduces the likelihood norm by inverting a sampled x_0 back to a latent state x_t by running the denoising process backward, which then serves as a favorable initialization for subsequent local control. Prior work such as ReGuidance (Karan et al., 2025) demonstrates that reversing to the highest noise state x_T can yield more observation-consistent trajectories. Building on this idea, our method is the first to design RS explicitly from the perspective of the likelihood-score norm. We provide a principled construction of the restart state and explain how norm-guided restarts improve downstream control.

Restart Sampling via Partial Inverse Encoding (RS). The RS procedure is simple: run the host solver once to obtain a reconstruction $x_0^{(1)}$, then invert $x_0^{(1)}$ *without conditioning on the observation* by running the deterministic reverse process (e.g., DDIM or ODE) backward from $t = 0$ to an intermediate noise level $t_{\text{inv}} := \lceil \tau_{\text{inv}} T \rceil$ (default $\tau_{\text{inv}} = 0.7$), resulting in a restart latent $z_{t_{\text{inv}}}$ finally, restart sampling from $z_{t_{\text{inv}}}$ again via the host solver and apply local control thereafter (Algorithm 1 and 2).

How RS enables effective local norm control. We formalize the theoretical justification of RS through a norm transfer bound, which connects the restart state's likelihood score norm to the first-pass sample's observation residual $\|y - \mathcal{A}x_0^{(1)}\|_2$ under DPS host. For deterministic reverse process from $t = 0$ to $t = \tau$, a key property is an *inverse-consistency* event: $\hat{x}_0(z_\tau, \tau) \approx x_0^{(1)}$, where $\hat{x}_0(z_\tau, \tau)$ is a the denoiser prediction for $\mathbb{E}[x_0|z_\tau]$ which maps noisy latent z_τ to final sampled x_0 often via Tweedie's formula (Tweedie, 1957). Combining these yields a norm transfer bound of the form

$$\|\hat{g}_t\|_2 \;\le\; C \cdot \Big(\|y - \mathcal{A}x_0^{(1)}\|_2 + \|\hat{x}_0(z_\tau, \tau) - x_0^{(1)}\|_2 \Big),$$

where C is a Jacobian-scaled constant from the Tweedie plug-in. See Appendix G.9 for the full derivation.

Therefore, a restart state that preserves a small observation residual also yields a smaller likelihood-score norm, providing a favorable MCMC initialization that makes even a single round of LNP sufficient and effective.

Truncated restart to preserve inverse consistency. In practice, RS truncates the reverse diffusion process at $t_{\text{inv}} < T$ rather than fully inverting to the maximum noise level T. This reflects a limitation of denoiser-based posterior-mean estimators in diffusion inverse solvers: as diffusion noise increases, the inverse-consistency event in Eq. (4.2) becomes increasingly fragile (Janati et al., 2024; Moufad et al., 2025; Bastek et al., 2025; Huang et al., 2024a).

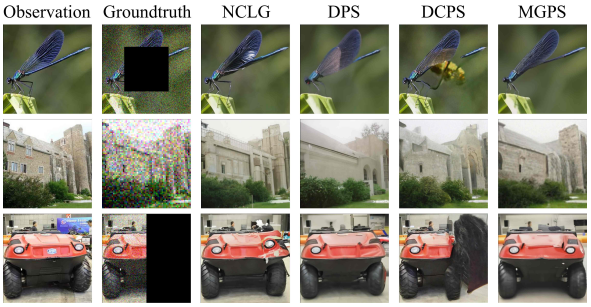

Figure 4. **Qualitative comparison of norm-controlled likelihood guidance on ImageNet.** Each row shows a different inverse problem instance with additive Gaussian noise ($\sigma_y = 0.3$). From left to right: the ground-truth image x_0, the corresponding observation y, and reconstructions from our method (NCLG) and three host solvers (DPS / DCPS / MGPS). **Top to bottom:** image inpainting, $4\times$ super-resolution, and image outpainting.

5. Experiments

We evaluate the proposed norm-control wrappers (Module A: RS, Module B: LNP, Module C: NSS) as plug-in modules on top of several host inverse samplers. Our experiments address two questions: (i) whether norm control improves reconstruction quality across hosts, datasets, and inverse tasks, and (ii) whether the ablation trends agree with the certificate viewpoint in Section 3.2.

5.1. Setup

We report results on **FFHQ** and **ImageNet256** for three standard inverse tasks: *Center* and *Half* inpainting, and *SR* $4\times$ super-resolution. Observations are corrupted by additive Gaussian noise with $\sigma_y \in \{0.05, 0.3\}$. We use **DPS** (Chung

Table 2. **LPIPS results** (lower is better) for different hosts and our Norm-Controlled Likelihood Guidance (NCLG) on FFHQ and ImageNet256. Columns correspond to task (*Center*, *Half*, *SR* 4×) and observation noise levels $\sigma_y \in \{0.05, 0.3\}$. For each +NCLG entry, parentheses report the relative change in % versus its host, $100 \cdot (\text{NCLG} - \text{host})/\text{host}$ (negative is better).

Method	**FFHQ** (LPIPS ↓)						**ImageNet256** (LPIPS ↓)					
	Center		Half		SR 4×		Center		Half		SR 4×	
	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$
DPS	0.153	0.316	0.246	0.325	0.243	0.282	0.378	0.483	0.443	0.495	0.453	0.537
+NCLG	**0.106** (−31%)	0.131 (−59%)	0.159 (−35%)	0.174 (−46%)	0.156 (−36%)	0.208 (−26%)	0.300 (−21%)	0.324 (−33%)	0.361 (−19%)	0.397 (−20%)	0.342 (−25%)	0.371 (−31%)
DCPS	0.122	0.287	0.207	0.268	0.227	0.319	0.326	0.401	0.359	0.391	0.441	0.472
+NCLG	0.108 (−11%)	0.153 (−47%)	**0.143** (−31%)	0.197 (−26%)	**0.135** (−41%)	0.257 (−19%)	**0.290** (−11%)	0.328 (−18%)	**0.310** (−14%)	0.328 (−16%)	0.341 (−23%)	0.358 (−24%)
MGPS	0.138	0.259	0.232	0.279	0.235	0.261	0.341	0.431	0.366	0.387	0.436	0.485
+NCLG	0.109 (−21%)	0.128 (−51%)	0.145 (−38%)	**0.168** (−40%)	0.147 (−37%)	0.168 (−36%)	0.295 (−13%)	0.331 (−23%)	0.315 (−14%)	**0.320** (−17%)	**0.331** (−24%)	**0.352** (−27%)
DAPS	0.142	0.274	0.218	0.264	0.229	0.257	0.339	0.425	0.357	0.420	0.396	0.455
+NCLG	0.115 (−19%)	**0.126** (−54%)	0.151 (−31%)	0.181 (−31%)	0.150 (−34%)	**0.165** (−36%)	0.306 (−10%)	**0.301** (−29%)	0.316 (−11%)	0.331 (−21%)	0.347 (−12%)	0.360 (−21%)

et al., 2023), **DCPS** (Janati et al., 2024), **MGPS** (Moufad et al., 2025), and **DAPS** (Zhang et al., 2025) in our codebase as host samplers (Section 4). For each host, we report results for the base solver and its NCLG-augmented variant, which incorporates three modules: Restart Sampling (RS), Noise Selection (NSS), and Langevin Projection (LNP). Table 2 reports LPIPS scores (lower is better), a perceptual metric that can serve as an empirical proxy for score error, motivated by our theoretical result that the likelihood-score norm predicts this error. The ablations additionally report **FID** and a **low-t aggregated likelihood norm** computed from the host likelihood-score estimator $\hat{g}_t \approx \nabla_{x_t} \log p_t(y|x_t)$, aggregated over a low-noise window consistent with Section 3. Detailed setup is shown in Appendix. G.

5.2. Norm control improves multiple hosts consistently

Table 2 summarizes LPIPS (lower is better) for four host solvers and their norm-controlled variants (+NCLG) across two datasets, three tasks, and two observation-noise levels. Appendix. B shows more metrics (PSNR, SSIM, FID) among more inverse tasks (Gaussian Deblur, HDR, motion deblur, JPEG2). NCLG improves the host in *every* reported setting, with relative reductions ranging from 10% to 59%. The gains are systematically larger under noisier measurements: averaged over all tasks, hosts, and datasets, NCLG reduces LPIPS by about 0.065 at $\sigma_y = 0.05$ (approximately 23% relative reduction) and by about 0.109 at $\sigma_y = 0.3$ (approximately 31% relative reduction). This trend is consistent with our diagnostic premise: as the measurement becomes less reliable, likelihood guidance is more prone to producing large likelihood-score norms, so explicitly steering trajectories toward lower-norm regions yields larger perceptual gains. Finally, these results support that likelihood-norm control is *complementary* to schedule-level refinements such as DCPS/MGPS: even when the host changes *when/how* guidance is injected, explicitly regulating the likelihood-score norm further improves reconstruction quality. This verifies the generalizability of our method.

Table 4 summarizes wall-clock time and peak GPU memory of a 5090 laptop on a representative hard setting (FFHQ *Center*, $\sigma_y = 0.3$) for NCLG modules. RS increases runtime mainly because it executes a second sampling pass, while keeping the peak memory essentially unchanged since it reuses the host sampler without introducing additional large state. LNP adds only a small per-step overhead (extra gradients / a few inner updates), leading to a modest time increase and a limited memory bump. NSS is the most expensive component because it evaluates multiple noise candidates by rolling out future trajectories to compute low-noise norm aggregates.

5.3. RS and LNP are complementary

Table 3 ablates RS and LNP on FFHQ with a DPS host, and reports LPIPS/FID together with the low-t likelihood norm. We observe three informative trends which validates our theoritical claims:

(i) RS reduces low-t likelihood norms and improves quality. RS (row "+A") substantially lowers the low-t likelihood norm relative to the DPS baseline, and improves LPIPS accordingly. This is consistent with the mechanism emphasized in Section 4.2 and Theorem 3.2.

(ii) RS relies on *deterministic* and *partial* inverse encoding. Our design of RS in Section 4.2 relies on an *inverse-consistency* condition (Eq. (4.2)) to ensure the effectiveness of reverse diffusion (running the sampling process backward from x_0). To validate this requirement, we show that injecting large noise during inversion ($\eta_{\text{inv}} = 1.0$) or fully reversing to the initial timestep ($\tau_{\text{inv}} = 1.0$), which breaks this consistency, results in negligible LPIPS improvement and significantly higher likelihood-score norms.

(iii) LNP alone can be harmful, but becomes effective after RS. Applying LNP without RS (row "+B") increases the likelihood norm and degrades LPIPS/FID, indicating that naive MCMC-style correction may suffer from a poor initialization. In contrast, RS+LNP (row "+A+B") achieves the best overall results and the smallest low-t norms, supporting the interpretation that RS provides a favorable initialization for the LNP Markov corrections.

Table 3. **Ablation of RS depth and LNP on FFHQ with DPS host**. We report three metrics per setting: LPIPS (lower is better), FID (lower is better), and low-t aggregated likelihood norm (lower is better). **Best** in each column is bold; second best is underlined. Module A = RS, B = LNP.

Method	Center						SR 4×					
	$\sigma_y = 0.05$			$\sigma_y = 0.3$			$\sigma_y = 0.05$			$\sigma_y = 0.3$		
	LPIPS ↓	FID ↓	Norm↓	LPIPS ↓	FID ↓	Norm↓	LPIPS ↓	FID ↓	Norm↓	LPIPS ↓	FID ↓	Norm↓
DPS (baseline)	0.153	30.83	963.84	0.316	31.55	938.87	0.243	29.51	79.10	0.282	32.45	85.22
+A	0.149	29.17	107.20	0.274	30.11	112.51	0.218	27.43	24.85	0.259	30.92	45.64
+A (η_{inv}=1.0)	0.154	31.04	970.43	0.312	31.43	947.52	0.239	29.44	78.73	0.285	32.35	86.51
+A (τ_{inv}=1.0)	0.153	30.91	958.31	0.318	31.49	940.17	0.246	30.03	79.44	0.283	32.88	83.95
+A (τ_{inv}=1.0 ODE)	0.147	37.51	754.29	0.280	42.07	826.93	0.225	36.24	65.33	0.268	38.14	71.13
+B	0.165	43.51	1343.77	0.337	48.20	1585.94	0.271	42.64	93.21	0.305	45.15	104.29
+A + B	**0.114**	**27.65**	**102.93**	**0.138**	**28.80**	**106.74**	**0.156**	**26.35**	**20.34**	**0.208**	**27.94**	**39.21**
+A (τ_{inv}=1.0) + B	0.170	31.47	810.77	0.342	35.64	871.74	0.265	30.67	72.43	0.311	31.20	74.39

Table 4. **Computation cost and LPIPS** on FFHQ *Center* inpainting with $\sigma_y = 0.3$. Time is wall-clock per sample (seconds), Peak GPU is peak memory usage, LPIPS is perceptual error (lower is better). Module A = RS, B = LNP, C = NSS.

Host	Variant	Time [s] ↓	Peak GPU [GB] ↓	LPIPS ↓
DPS	baseline	**5.25**	**2.81**	0.316
	+A (RS)	11.24	**2.81**	0.274
	+B (LNP)	13.48	3.25	0.138
	+C (NSS)	29.12	10.24	0.131
DCPS	baseline	14.86	**2.81**	0.287
	+A (RS)	28.51	**2.81**	0.225
	+B (LNP)	30.27	3.25	0.166
	+C (NSS)	48.28	10.24	0.153
MGPS	baseline	10.49	3.14	0.259
	+A (RS)	23.10	3.14	0.234
	+B (LNP)	24.93	3.57	0.138
	+C (NSS)	39.15	10.55	**0.128**

5.4. Comparison to REGUIDANCE (Karan et al., 2025)

We also compare the REGUIDANCE Wrapper, referred as the variant "+ A (τ_{inv}=1.0 ODE)". ReGuidance first samples an initial solution x_0 using a host solver, then runs the deterministic reverse diffusion process without the observation y to reach the *highest-noise* latent x_T, and finally reruns sampling from x_T using a *deterministic variant* of the host solver. Empirically, this baseline improves LPIPS over the host, confirming that two-pass restart can help. However, it is consistently worse than our default RS+LNP configuration, especially on FID. We attribute this gap to two theoretical constraints emphasized by ReGuidance: **(i) Deterministic requirement.** ReGuidance requires deterministic second-pass sampling, and reports that its benefit can diminish when combined with stochastic one (Karan et al., 2025). This restriction is detrimental for distribution-level metrics (FID), whereas our RS is compatible with stochastic second-pass sampling. **(ii) Full reversion to high noise.** Full reversion traverses the highest-noise regime may break the inverse-consistency (Eq. (4.2)) that RS exploits.

6. Conclusion

We introduced *Norm-Controlled Likelihood Guidance* (NCLG), a training-free wrapper for diffusion inverse-problem solvers that treats the host's likelihood-score proxy as a solver-exposed diagnostic and control signal. Empirically, we find that windowed *low-noise* aggregates of the likelihood-score norm are consistently predictive of perceptual reconstruction quality. To explain why this diagnostic is meaningful, we analyze a tractable multi-mode Gaussian-mixture toy in which the true conditional score is available, and establish a matching certificate: when the likelihood-score norm is small (in the low-noise regime), the Tweedie plug-in conditional score must be close to the true conditional score up to an explicit floor. Guided by this viewpoint, we propose three lightweight and composable modules including restart sampling (RS), likelihood-norm projection (LNP), and likelihood-norm stochastic selection (NSS), that target trajectories with smaller likelihood norms.

Across FFHQ and ImageNet-256 inverse problems (inpainting and 4× super-resolution with Gaussian deblurring) and multiple host samplers, these modules deliver consistent quality gains and favorable quality–compute trade-offs, with the strongest improvements typically arising from combining RS and LNP and optional additional gains from NSS when extra compute is available. Limitations and future work include: (i) the most informative norm signals arise late in the schedule and NSS-style probing introduces rollout overhead, motivating more adaptive budget allocation (e.g., early stopping and dynamic windowing); (ii) our certificate analysis is mechanistic and derived for Tweedie/denoiser plug-in likelihood guidance, so extending comparable guarantees to alternative likelihood-score estimators remains an important direction; and (iii) while we observe useful signals beyond the low-noise window and robust gains on representative hosts and tasks, broader coverage and a sharper characterization of the high-noise regime remain open.

Impact Statements

This work makes diffusion-based inverse solvers more reliable and controllable by exposing a universal, in-loop diagnostic (likelihood-score norm) and using it to steer sampling with effective modules. The result is higher-quality reconstructions at lower tuning burden, with potential impact on medical imaging, microscopy, and scientific sensing pipelines.

References

Arridge, S., Maass, P., Öktem, O., and Schönlieb, C.-B. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.

Bastek, J.-H., Sun, W., and Kochmann, D. Physics-informed diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=tpYeermigp.

Boutin, V., Fel, T., Singhal, L., Mukherji, R., Nagaraj, A., Colin, J., and Serre, T. Diffusion models as artists: Are we closing the gap between humans and machines? In *Proceedings of the 40th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2023.

Boys, B., Girolami, M., Pidstrigach, J., Reich, S., Mosca, A., and Akyildiz, O. D. Tweedie moment projected diffusions for inverse problems. *Transactions on Machine Learning Research*, 2024. arXiv:2310.06721.

Cardoso, G., Janati, Y., Le Corff, S., and Moulines, E. Monte carlo guided diffusion for bayesian linear inverse problems. In *International Conference on Learning Representations*, 2024.

Chen, S., Chewi, S., Lee, H., Li, Y., Lu, J., and Salim, A. The probability flow ode is provably fast, 2023. URL https://arxiv.org/abs/2305.11798.

Chen, S., Nagaraj, D., Wu, H., Yang, J., Zhang, A. R., Piech, C., and Ravikumar, P. Learning general Gaussian mixtures with efficient score matching. *arXiv preprint arXiv:2404.18893*, 2024.

Cheng, C., Han, B., Maddix, D. C., Ansari, A. F., Stuart, A., Mahoney, M. W., and Wang, B. Gradient-free generation for hard-constrained systems. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=teE4pl9ftK.

Chidambaram, M., Gatmiry, K., Chen, S., Lee, H., and Lu, J. What does guidance do? a fine-grained analysis in a simple setting, 2024. URL https://arxiv.org/abs/2409.13074.

Chung, H., Kim, J., McCann, M. T., Klasky, M. L., and Ye, J. C. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023.

Dhariwal, P. and Nichol, A. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, 2021. arXiv:2105.05233.

Dou, Z. and Song, J. Filtering posterior sampling for linear inverse problems with diffusion models. In *International Conference on Learning Representations*, 2024.

Erdogdu, M. A., Mackey, L., and Shamir, O. Global non-convex optimization with discretized diffusions. In *Advances in Neural Information Processing Systems 31*, pp. 9671–9680, 2018. URL https://papers.nips.cc/paper/8175-global-non-convex-optimization-with-discreti

Geman, S. and Hwang, C.-R. Diffusions for global optimization. *SIAM Journal on Control and Optimization*, 24 (5):1031–1043, 1986. doi: 10.1137/0324060.

Hajek, B. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 13(2):311–329, 1988. doi: 10.1287/moor.13.2.311.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Klambauer, G., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017.

Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.

Hong, C., Cha, B., and Oh, T.-H. CAS: A probability-based approach for universal condition alignment score. In *International Conference on Learning Representations*, 2024.

Huang, J., Yang, G., Wang, Z., and Park, J. J. DiffusionPDE: Generative PDE-solving under partial observation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a. URL https://openreview.net/forum?id=z0I2SbjN0R.

Huang, Y., Liang, Z., Li, H., Chen, K., and Chen, Y. Stochastic control guidance: Plug-and-play guidance for non-differentiable rules in diffusion models. In *Proceedings of the 41st International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2024b.

Jacobsen, C., Zhuang, Y., and Duraisamy, K. Cocogen: Physically-consistent and conditioned score-based generative models for forward and inverse problems, 2024. URL https://arxiv.org/abs/2312.10527.

Janati, Y., Moufad, B., Durmus, A., Moulines, E., and Olsson, J. Divide-and-conquer posterior sampling for denoising diffusion priors. In *Advances in Neural Information Processing Systems*, 2024.

Karan, A., Shah, K., and Chen, S. Reguidance: A simple diffusion wrapper for boosting sample quality on hard inverse problems, 2025. URL https://arxiv.org/abs/2506.10955.

Kawar, B., Elad, M., Ermon, S., and Song, J. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022.

Li, Z., Dou, H., Fang, S., Han, W., Deng, Y., and Yang, L. Physics-aligned field reconstruction with diffusion bridge. In *The Thirteenth International Conference on Learning Representations*.

Liang, Y., Sha, Z., Shi, Z., Song, Z., Wan, M., and Zhou, Y. Unraveling the smoothness properties of diffusion models: A gaussian mixture perspective. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025.

Mardani, M., Song, J., Kautz, J., and Vahdat, A. A variational perspective on solving inverse problems with diffusion models. In *International Conference on Learning Representations*, 2024.

Moufad, B., Janati, Y., Bedin, L., Durmus, A. O., Douc, R., Moulines, E., and Olsson, J. Variational diffusion posterior sampling with midpoint guidance. In *International Conference on Learning Representations*, 2025.

Nichol, A. and Dhariwal, P. Improved denoising diffusion probabilistic models. In *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2021.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

Shah, K., Chen, S., and Klivans, A. Learning mixtures of gaussians using the DDPM objective. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Song, J., Meng, C., and Ermon, S. Loss-guided diffusion models for plug-and-play controllable generation. In *Proceedings of the 40th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2023a.

Song, J., Vahdat, A., Mardani, M., and Kautz, J. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*, 2023b.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

Staniszewski, L., Kuciński, L., and Deja, K. There and back again: On the relation between noises, images, and their inversions in diffusion models. In *International Conference on Learning Representations*, 2025.

Stuart, A. M. Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559, 2010.

Tang, H., Fang, J., Wang, Y., and Yang, X. Solving general noisy inverse problem via posterior sampling: A policy gradient viewpoint. *arXiv preprint arXiv:2403.10585*, 2024.

Tikhonov, A. N. Solution of incorrectly formulated problems and the regularization method. *Sov Dok*, 4:1035–1038, 1963.

Tweedie, M. C. Statistical properties of inverse gaussian distributions. i. *The Annals of Mathematical Statistics*, 28(2):362–377, 1957.

Wang, Y., Yu, J., Wang, Y., Liu, C., and Dong, C. Zero-shot image restoration using denoising diffusion null-space model. In *International Conference on Learning Representations*, 2023.

Wu, Z., Sun, Y., Chen, Y., Zhang, B., Yue, Y., and Bouman, K. L. Principled probabilistic imaging using diffusion models as plug-and-play priors. In *Advances in Neural Information Processing Systems*, 2024.

Zhang, B., Chu, W., Berner, J., Meng, C., Anandkumar, A., and Song, Y. Improving diffusion inverse problem solving with decoupled noise annealing. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20895–20905. IEEE, June 2025. doi: 10.1109/cvpr52734.2025.01946. URL http://dx.doi.org/10.1109/CVPR52734.2025.01946.

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

A. Qualitative Results

We show two concrete examples comparing two hosts (DPS, DCPS) with our norm-controlled modules in Fig. 5. We also provide more visualization comparing more host solvers and their NCLG variants in Fig. 6 and 7.

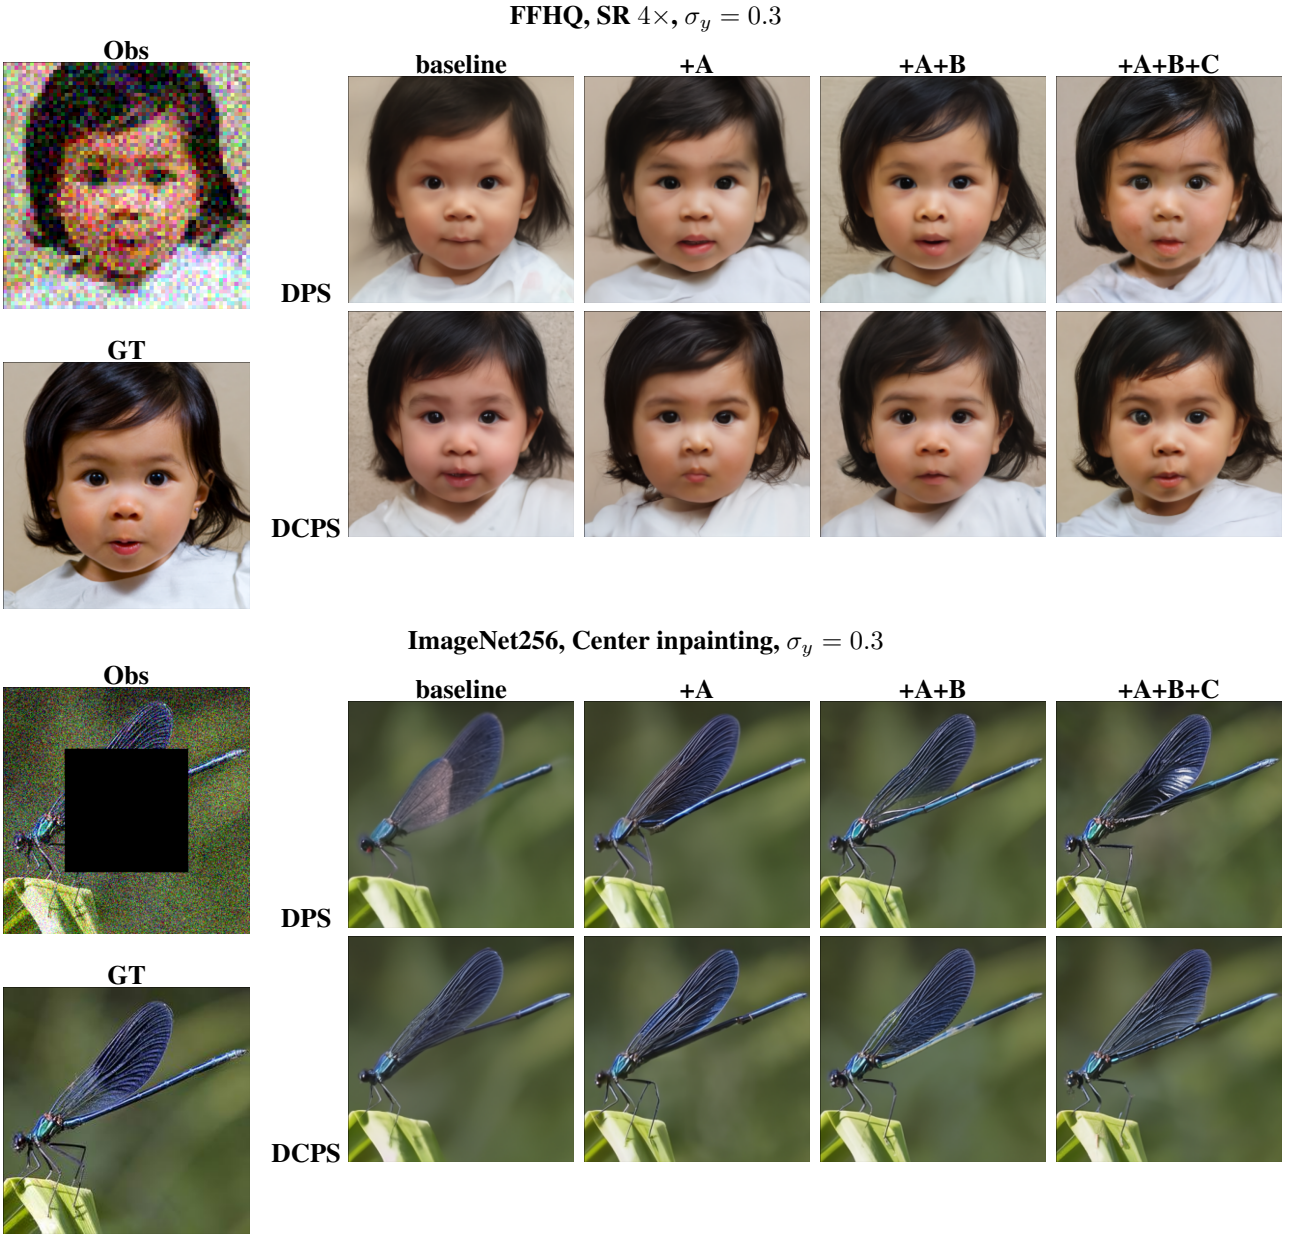

Figure 5. **Additional qualitative results (single example).** For each task, we show the noisy observation (Obs) and ground truth (GT) once (left). Reconstructions are shown for two hosts (rows: DPS, DCPS) under baseline and progressively adding our norm-control modules (columns: baseline, +A: RS, +A+B: RS+LNP, +A+B+C: RS+LNP+NSS).

B. More Qualitative Results

In this section, we provide more qualitative results to demonstrate the performance of NCLG. Here, Table. 5 and 6 show more metrics results beyond LPIPS using PSNR, SSIM and FID. These results show that NCLG improve upon base solvers among general metrics rather than optimizing a dedicated metric. We also conduct experiment on more inverse problem tasks including Gaussian deblur, Jpeg and HDR in Table 7 and 8.

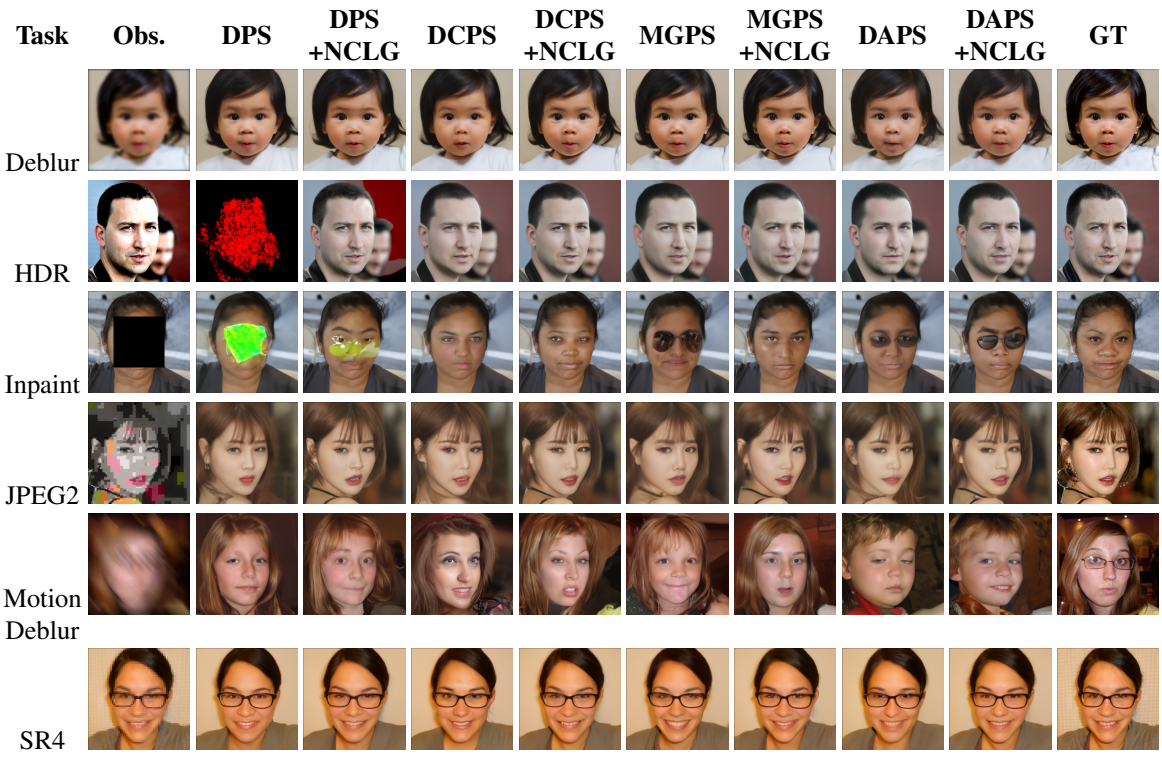

Figure 6. Additional qualitative comparisons on FFHQ. Columns show the observation, each host baseline, the corresponding host with NCLG, and the ground truth.

C. Background and Related Work

C.1. Diffusion Models and Diffusion Inverse Solvers

Score-based diffusion models learn a time-indexed score field $\nabla_{x_t} \log p_t(x_t)$ for the noisy latent x_t, either in discrete time via denoising diffusion probabilistic models (DDPMs) or in continuous time via stochastic differential equations (SDEs) (Ho et al., 2020; Song et al., 2021; Nichol & Dhariwal, 2021; Rombach et al., 2022). Given an observation $y = A(x_0) + \text{noise}$, an unconditional diffusion prior can be used as a powerful image prior for solving inverse problems by sampling from the posterior over clean images, $p(x_0|y)$.

Under the usual forward noising and reverse-time formulation, conditioning on y leads to the decomposition

$$\nabla_{x_t} \log p_t(x_t|y) = \nabla_{x_t} \log p_t(x_t) + \nabla_{x_t} \log p_t(y|x_t), \tag{6}$$

where the prior score $\nabla_{x_t} \log p_t(x_t)$ is provided by the pretrained diffusion model, while the likelihood term $\nabla_{x_t} \log p_t(y|x_t)$ is typically intractable for general forward operators and noise models. Equation (6) is the common starting point for almost all diffusion-based inverse solvers, and different methods can be viewed as different approximations or parametrizations of the conditional likelihood term.

Existing solvers can be grouped by how they handle $\nabla_{x_t} \log p_t(y|x_t)$. *Closed-form/projection methods* exploit linear-Gaussian structure and analytic conditional distributions. DDRM (Kawar et al., 2022), DDNM (Wang et al., 2023), and ΠGDM (Song et al., 2023b) diagonalize the forward operator via SVD or pseudoinverses and update latent variables in measurement-aligned coordinates, yielding efficient and often exact conditional gradients in that regime. More recent Tweedie-type constructions approximate the measurement term by projecting the intractable posterior $p(x_0|x_t, y)$ onto tractable families, yielding closed-form or nearly closed-form likelihood scores in linear inverse problems (Boys et al., 2024).

A second family, which we refer to as *general guidance*, treats the likelihood term through a surrogate gradient evaluated at a predicted clean image. Diffusion Posterior Sampling (DPS) differentiates an observation loss through the model's predictor $\hat{x}_0(x_t, t)$ and scales it by the diffusion coefficients (Chung et al., 2023). Loss-Guided Diffusion (LGD) averages gradients over a local Gaussian around $\hat{x}_0$ to stabilize the guidance (Song et al., 2023a). Policy-gradient based guidance (DPG)

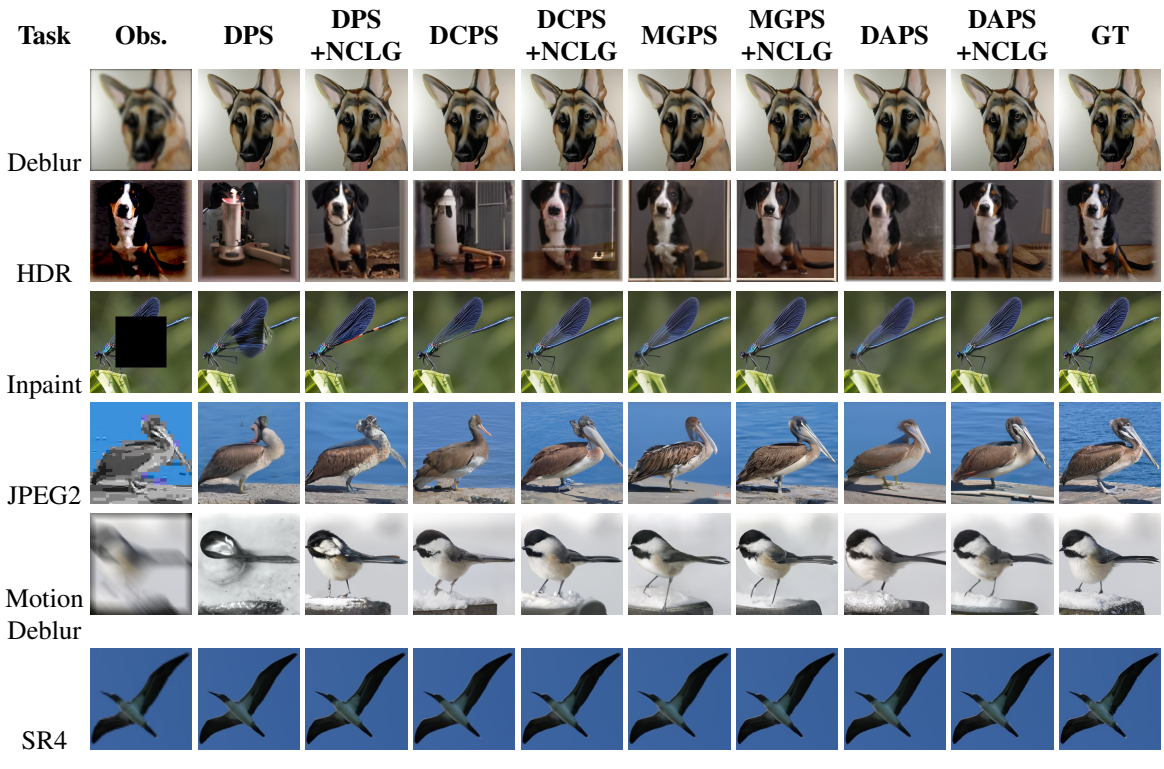

Figure 7. Additional qualitative comparisons on ImageNet. Columns show the observation, each host baseline, the corresponding host with NCLG, and the ground truth.

re-casts posterior sampling as a reinforcement-learning problem on noisy states (Tang et al., 2024), and Stochastic Control Guidance (SCG) formulates guidance as an optimal control problem, enabling plug-and-play guidance for non-differentiable rules (Huang et al., 2024b). All of these methods insert an approximate likelihood gradient into the reverse dynamics without modifying the underlying diffusion prior.

A third line of work adopts a more explicitly Bayesian viewpoint. RED-diff (Mardani et al., 2024) optimizes a Dirac variational approximation to the posterior with a diffusion-based regularizer, which tends to be mode-seeking by design. Sequential Monte Carlo samplers such as Filtering Posterior Sampling (FPS-SMC) and Monte Carlo Guided Diffusion (MCGDiff) construct asymptotically exact samplers by propagating and reweighting particles along the schedule (Dou & Song, 2024; Cardoso et al., 2024). PnP-DM builds an MCMC sampler over latent diffusion states and interprets diffusion priors as plug-and-play priors for probabilistic imaging (Wu et al., 2024). These probabilistic methods can provide stronger guarantees but are typically more computationally demanding.

The norm-controlled likelihood view developed in this work is compatible with all three families. It assumes a host solver that exposes (i) a base reverse step and (ii) an estimator of $\nabla_{x_t} \log p_t(y|x_t)$ (or a proxy thereto), and then applies training-free modules that adjust the *norm* of this term without changing how it is estimated.

C.2. Guidance, Likelihood Approximations, and Scheduling

For a fixed inverse task and observation model, the main difficulty lies in approximating the likelihood term in Equation (6). Even in simple linear-Gaussian settings, the exact conditional likelihood $p_t(y|x_t)$ depends on the intractable denoising distribution $p_t(x_0|x_t)$, so its gradient with respect to x_t is hard to evaluate. A common pattern in general guidance methods is to approximate $\nabla_{x_t} \log p_t(y|x_t)$ via a chain rule through a surrogate for the posterior mean:

$$\widehat{g}_t(x_t) \;\approx\; \nabla_{x_t} \log p_t(y|x_t) = J_t(x_t)^\top \nabla_{x_0} \log p(y|x_0)\big|_{x_0=\hat{x}_0(x_t,t)}, \tag{7}$$

where $\hat{x}_0(x_t, t) \approx \mathbb{E}_\theta[x_0|x_t]$ is the model's denoiser and $J_t(x_t)$ is its Jacobian with respect to x_t. In DPS, for example, $\nabla_{x_0} \log p(y|x_0)$ is computed as the gradient of an observation loss (e.g., squared residual in the linear-Gaussian case) between $A(\hat{x}_0)$ and y, and the resulting gradient is backpropagated through $\hat{x}_0$ (Chung et al., 2023). LGD, DPG, and SCG

Table 5. **PSNR/SSIM/FID** for different hosts and our Norm-Controlled Likelihood Guidance (NCLG) on FFHQ. Columns correspond to task (*Center*, *Half*, *SR* 4×) and observation noise levels $\sigma_y \in \{0.05, 0.3\}$.

Method	**FFHQ** (PSNR/SSIM/FID)					
	Center		Half		SR 4×	
	$\sigma_y = 0.05$	$\sigma_y = 0.3$	$\sigma_y = 0.05$	$\sigma_y = 0.3$	$\sigma_y = 0.05$	$\sigma_y = 0.3$
DPS	22.00/0.74/30.83	16.20/0.55/60.84	15.72/0.66/44.63	13.84/0.52/55.72	22.46/0.65/42.18	18.72/0.58/49.86
+NCLG	23.17/0.79/26.92	20.55/0.74/32.74	17.05/0.75/31.10	16.62/0.70/36.84	24.98/0.75/31.74	21.08/0.69/39.52
DCPS	22.86/0.77/28.11	16.95/0.58/55.62	16.43/0.71/35.87	14.55/0.58/49.18	23.38/0.69/37.26	17.42/0.52/56.44
+NCLG	23.10/0.79/27.10	19.62/0.70/36.10	17.62/0.78/27.80	15.90/0.65/41.02	26.18/0.79/27.48	19.54/0.63/46.31
MGPS	22.48/0.76/29.23	17.38/0.61/50.94	15.98/0.68/39.68	14.28/0.56/46.35	23.09/0.68/38.54	19.22/0.61/45.90
+NCLG	23.05/0.79/27.24	20.72/0.75/31.98	17.58/0.79/28.10	16.95/0.72/34.96	25.70/0.77/29.46	22.02/0.73/34.20
DAPS	22.35/0.76/29.54	17.10/0.59/53.74	16.27/0.70/37.12	14.64/0.59/47.28	23.27/0.69/37.88	19.40/0.62/45.10
+NCLG	22.90/0.78/27.88	20.81/0.75/31.35	17.36/0.77/29.20	16.74/0.71/36.08	25.56/0.77/29.82	22.18/0.74/33.40

Table 6. **PSNR/SSIM/FID** for different hosts and our Norm-Controlled Likelihood Guidance (NCLG) on ImageNet256. Columns correspond to task (*Center*, *Half*, *SR* 4×) and observation noise levels $\sigma_y \in \{0.05, 0.3\}$.

Method	**ImageNet256** (PSNR/SSIM/FID)					
	Center		Half		SR 4×	
	$\sigma_y = 0.05$	$\sigma_y = 0.3$	$\sigma_y = 0.05$	$\sigma_y = 0.3$	$\sigma_y = 0.05$	$\sigma_y = 0.3$
DPS	17.62/0.61/54.88	14.86/0.46/69.42	14.42/0.55/56.28	12.84/0.43/66.40	20.96/0.52/71.64	18.34/0.44/78.96
+NCLG	19.14/0.71/43.06	18.02/0.66/46.52	15.92/0.66/44.38	15.08/0.59/49.88	23.18/0.65/55.12	21.30/0.61/58.74
DCPS	18.54/0.68/46.91	16.02/0.55/59.18	15.46/0.63/47.52	14.21/0.54/53.12	21.28/0.55/67.28	19.24/0.49/70.22
+NCLG	19.62/0.74/40.28	17.92/0.65/48.10	16.38/0.70/39.72	16.08/0.67/41.22	23.24/0.66/53.88	21.66/0.64/55.44
MGPS	18.33/0.66/48.72	15.61/0.52/64.27	15.28/0.62/48.36	14.36/0.55/52.06	21.36/0.56/66.02	18.96/0.47/72.48
+NCLG	19.46/0.73/41.36	17.84/0.64/48.76	16.18/0.68/41.10	16.32/0.69/39.58	23.51/0.68/51.96	21.92/0.66/53.06
DAPS	18.42/0.67/47.95	15.78/0.53/62.80	15.54/0.64/46.92	13.76/0.50/57.84	22.08/0.60/60.74	19.72/0.52/66.18
+NCLG	19.31/0.72/42.18	18.34/0.67/44.95	16.12/0.68/41.84	15.96/0.66/41.94	23.02/0.64/54.62	21.54/0.63/56.32

can all be viewed as different design choices for the loss, its smoothing, and the estimator of the value function in (7) (Song et al., 2023a; Tang et al., 2024; Huang et al., 2024b).

While conceptually simple, the point-mass substitution $p_t(y|x_t) \approx p(y|\hat{x}_0(x_t, t))$ and the chain-rule form in (7) introduce an approximation error whose magnitude depends on the noise level and the time interval on which the approximation is applied. Tweedie-type methods such as TMPD refine this approximation by matching higher-order moments of $p(x_0|x_t)$ and deriving closed-form expressions for the likelihood gradient in linear inverse problems, thereby improving the fidelity of the conditional score (Boys et al., 2024). However, these refinements still propagate through noisy latents x_t, and the quality of the approximation remains time-dependent.

A complementary line of work modifies *where* and *when* the likelihood is injected. Divide-and-Conquer Posterior Sampling (DCPS) shows that the point-mass substitution used in DPS induces a bias whose magnitude grows with the reverse-time jump length. To control this bias, DCPS subdivides the time horizon into short sub-intervals and introduces intermediate potentials $g_\ell(x_{t_\ell})$ on the latent states themselves. Each segment is then sampled from a local bridged posterior using a short sequence of Langevin steps combined with a variational update, which significantly reduces the accumulated approximation error (Janati et al., 2024). Midpoint Guidance Posterior Sampling (MGPS) is a single-step variant that inserts a midpoint latent x_m between $x_{t_{k+1}}$ and x_{t_k} and uses x_m to estimate the likelihood contribution at t_k (Moufad et al., 2025). Its analysis reveals a tension: injecting the observation very late (near $t = 0$) yields more accurate likelihood evaluations but provides fewer steps for the guidance signal to act, whereas injecting it early (large t) supplies more iterations but relies on a noisier, more biased approximation.

Probabilistic methods also implicitly encode guidance schedules. FPS-SMC constructs a sequence of intermediate posteriors over noisy latents and propagates particles through the diffusion schedule with importance weighting and resampling (Dou

Table 7. **Predicted LPIPS/PSNR/SSIM/FID** for different hosts and our Norm-Controlled Likelihood Guidance (NCLG) on FFHQ. Each cell is reported as *LPIPS / PSNR / SSIM / FID*. All tasks use observation noise $\sigma_y = 0.05$; JPEG uses QF= 2.

	FFHQ (LPIPS / PSNR / SSIM / FID)			
Method	**Gauss. Deblur**	**Motion Deblur**	**JPEG (QF=2)**	**HDR**
DPS	0.192 / 24.61 / 0.772 / 24.38	0.181 / 24.94 / 0.781 / 23.52	0.332 / 22.27 / 0.691 / 35.84	0.358 / 21.84 / 0.673 / 39.47
+NCLG	0.146 / 26.08 / 0.816 / 18.92	0.136 / 26.44 / 0.828 / 17.63	0.231 / 24.18 / 0.752 / 25.71	0.189 / 24.76 / 0.796 / 21.48
DCPS	0.158 / 25.68 / 0.803 / 21.66	0.149 / 26.03 / 0.817 / 20.88	0.181 / 24.74 / 0.773 / 27.42	0.153 / 25.22 / 0.789 / 24.96
+NCLG	0.133 / 26.59 / 0.831 / 18.47	0.124 / 26.97 / 0.842 / 17.58	0.146 / 25.78 / 0.809 / 22.63	0.126 / 26.14 / 0.824 / 20.91
MGPS	0.171 / 25.41 / 0.792 / 22.18	0.154 / 26.19 / 0.814 / 21.07	0.179 / 24.93 / 0.781 / 28.16	0.118 / 26.48 / 0.833 / 20.36
+NCLG	0.128 / 26.78 / 0.834 / 18.31	0.115 / 27.29 / 0.851 / 17.19	0.132 / 26.02 / 0.817 / 22.14	0.092 / 27.33 / 0.868 / 16.82
DAPS	0.169 / 25.52 / 0.798 / 22.44	0.160 / 25.87 / 0.809 / 21.52	0.188 / 24.58 / 0.772 / 28.91	0.166 / 25.07 / 0.784 / 25.17
+NCLG	0.129 / 26.71 / 0.833 / 18.72	0.121 / 27.06 / 0.846 / 17.84	0.141 / 25.89 / 0.812 / 22.47	0.127 / 26.39 / 0.826 / 20.58

Table 8. **Predicted LPIPS/PSNR/SSIM/FID** for different hosts and our Norm-Controlled Likelihood Guidance (NCLG) on ImageNet256. Each cell is reported as *LPIPS / PSNR / SSIM / FID*. All tasks use observation noise $\sigma_y = 0.05$; JPEG uses QF= 2.

	ImageNet256 (LPIPS / PSNR / SSIM / FID)			
Method	**Gauss. Deblur**	**Motion Deblur**	**JPEG (QF=2)**	**HDR**
DPS	0.338 / 21.72 / 0.536 / 52.41	0.307 / 22.84 / 0.584 / 48.26	0.463 / 19.74 / 0.442 / 67.83	0.492 / 18.91 / 0.413 / 72.94
+NCLG	0.276 / 23.31 / 0.603 / 42.18	0.243 / 24.21 / 0.649 / 37.56	0.334 / 21.68 / 0.532 / 50.27	0.292 / 21.94 / 0.576 / 46.11
DCPS	0.279 / 22.94 / 0.592 / 44.37	0.246 / 24.01 / 0.638 / 40.62	0.326 / 22.07 / 0.566 / 49.58	0.212 / 23.28 / 0.671 / 39.14
+NCLG	0.249 / 23.78 / 0.631 / 38.96	0.217 / 24.72 / 0.672 / 34.81	0.286 / 22.96 / 0.608 / 42.73	0.183 / 24.31 / 0.714 / 33.92
MGPS	0.287 / 22.66 / 0.577 / 45.02	0.241 / 24.18 / 0.646 / 39.87	0.338 / 21.96 / 0.558 / 51.11	0.183 / 23.96 / 0.697 / 36.28
+NCLG	0.247 / 23.92 / 0.635 / 39.24	0.205 / 24.88 / 0.681 / 34.18	0.287 / 22.88 / 0.604 / 43.06	0.151 / 24.86 / 0.737 / 31.17
DAPS	0.276 / 22.81 / 0.586 / 43.68	0.232 / 24.34 / 0.652 / 38.94	0.259 / 23.14 / 0.631 / 41.37	0.219 / 23.41 / 0.684 / 38.02
+NCLG	0.239 / 24.04 / 0.641 / 37.92	0.198 / 24.97 / 0.687 / 33.54	0.224 / 24.08 / 0.672 / 35.46	0.186 / 24.29 / 0.721 / 32.84

& Song, 2024). MCGDiff designs Monte Carlo moves that are tailored to diffusion-based priors in linear inverse problems (Cardoso et al., 2024). In all cases, the contribution of the likelihood term to the sampler is strongest in regions where its gradient norm is large, which empirically concentrates at relatively high noise levels (large t) in many imaging tasks. Our experiments and analysis corroborate this view: for a fixed image and task, the aggregated likelihood-norm along a trajectory is dominated by the early high-noise steps and strongly correlates with downstream metrics such as LPIPS, SSIM, and PSNR. In this sense, large-t segments are both the most influential for guidance and the most prone to approximation error, motivating explicit control of the likelihood-norm rather than implicitly trusting large gradients wherever they arise.

C.3. Noise/Score Alignment in Conditional Diffusion

A growing body of work links downstream quality to how well scores and noises are aligned with conditioning signals. Boutin et al. define a *misalignment map* that aggregates, over time, the ℓ_2 distance between normalized conditional scores $\delta_t^y(x, y) = \nabla_{x_t} \log p_\theta^*(x_t|y)$ and unconditional scores $\delta_t(x) = \nabla_{x_t} \log p_\theta(x_t)$ (Boutin et al., 2023). They show that this normalized score discrepancy strongly correlates with recognizability and that regions highlighted by the misalignment map track human-diagnostic features, suggesting that better alignment between conditional and unconditional scores leads to better perceptual quality.

Hong et al. propose CAS, a probability-based condition-alignment score defined from the conditional probability along the diffusion process (Hong et al., 2024). CAS can be evaluated for diverse conditioning modalities (text, images, instructions, audio) and enables sample-level self-rejection of misaligned generations without extra models, again demonstrating that better-aligned conditional and unconditional distributions tend to produce higher-quality samples.

From a different angle, Staniszewski et al. analyze the relation between the initial noise x_T, the generated image, and the latent obtained by DDIM inversion (Staniszewski et al., 2025). They identify a systematic bias in early inversion steps: the inverted noises become less diverse, particularly over smooth regions, making the inversion space less manipulable than the original noise space. Replacing a few early inversion steps by a short forward diffusion decorrelates latents and improves editing and interpolation, again tying the structure of x_T to downstream controllability and quality.

Table 9. **Pearson correlations between LPIPS and windowed likelihood-norm aggregates (appendix).** Each entry reports Pearson correlation ρ_P followed by a representative aggregate magnitude in parentheses. Low-noise and High-noise use timestep windows $\mathcal{W}_{\text{low}}$ and $\mathcal{W}_{\text{high}}$ (defined in text); All (Full) aggregates over the entire reverse trajectory.

Task	σ_y	Low-noise	High-noise	All (Full)
		FFHQ		
SR 4×	0.05	0.46(14.71)	0.29(209.96)	0.31(81.60)
	0.30	0.29(34.64)	0.15(220.21)	0.14(89.17)
Center	0.05	0.64(40.62)	0.16(2947.66)	0.17(959.32)
	0.30	0.78(112.03)	0.21(2942.57)	0.22(980.57)
Half	0.05	0.67(31.00)	0.30(2506.68)	0.31(827.17)
	0.30	0.52(94.41)	0.24(2567.35)	0.27(851.79)
		ImageNet		
SR 4×	0.05	0.32(11.70)	0.14(102.11)	0.17(39.81)
	0.30	0.24(32.95)	0.13(110.97)	0.14(46.24)
Center	0.05	0.83(31.75)	0.19(1221.75)	0.20(423.05)
	0.30	0.76(110.76)	0.20(1243.91)	0.21(455.36)
Half	0.05	0.42(37.38)	0.20(773.39)	0.22(279.55)
	0.30	0.47(91.57)	0.23(833.64)	0.26(323.64)

In conditional inverse problems, DCPS provides a complementary perspective on alignment by quantifying how the standard point-mass substitution $p_t(y|x_t) \approx p(y|\hat{x}_0(x_t, t))$ introduces a bias whose magnitude increases with the reverse-time interval (Janati et al., 2024). By shortening time gaps and relocating the likelihood term to intermediate noisy states, DCPS empirically reduces this approximation error, and MGPS further refines where in time the observation is injected (Moufad et al., 2025). These works primarily reason about *directional* alignment of scores and about where along the schedule likelihood information should be introduced.

Taken together, misalignment maps (Boutin et al., 2023), condition alignment scores (Hong et al., 2024), inversion-bias studies (Staniszewski et al., 2025), and scheduling methods (Janati et al., 2024; Moufad et al., 2025) support a unifying principle: alignment between conditional and unconditional scores, and between noise states and their intended conditioning, is predictive of downstream sample quality. Our work adds a complementary perspective: instead of designing new alignment metrics or new score estimators, we directly regularize the *norm* of the estimated likelihood term along the reverse trajectory, which in turn encourages the conditional and unconditional scores to stay close in regions that actually drive the sampling dynamics.

C.4. Positioning Our Work

Norm-Controlled Likelihood Guidance focuses on the magnitude of the likelihood term in Equation (6). In contrast to methods that derive new estimators for $\nabla_{x_t} \log p_t(y|x_t)$ or redesign the entire posterior sampler, we take the estimator provided by a host solver (DPS-type guidance, Tweedie projections, bridge potentials from DCPS/MGPS, or closed-form gradients from DDRM/DDNM/ΠGDM) as given, and build training-free modules that act *in noise space*.

Concretely, our modules (i) project the current noise toward states with smaller estimated likelihood-norm before performing a host step, and (ii) restart sampling from intermediate noise levels where the estimated likelihood term is typically smaller and more faithful. Because the only interface they require is a base reverse step and an estimate (or proxy) of $\nabla_{x_t} \log p_t(y|x_t)$, these modules are orthogonal to the choice of host solver and can be combined with DPS, DCPS, MGPS, ΠGDM, PnP-DM, RED-diff, FPS-SMC, MCGDiff, and related methods (Kawar et al., 2022; Wang et al., 2023; Song et al., 2023b; Chung et al., 2023; Janati et al., 2024; Moufad et al., 2025; Mardani et al., 2024; Dou & Song, 2024; Cardoso et al., 2024; Wu et al., 2024). Our empirical and theoretical results in later sections show that explicitly controlling the likelihood-norm in this way improves stability and reconstruction quality across a wide range of diffusion inverse problems, while preserving the flexibility and generality of existing guidance schemes.

D. Derivation of Score Function of High Dimensional MoG Dataset

We consider an MoG distribution in the following:

$$\mathbf{x}_0 \sim \frac{1}{K}\sum_{k=1}^{K}\mathcal{N}(\boldsymbol{\mu}_k, \sigma_k^2 \cdot),$$

where K is the number of Gaussian components, $\boldsymbol{\mu}_k$ and σ_k^2 are the means and variances of the Gaussian components, respectively. Suppose the solution of the diffusion process follows:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \cdot \xi \quad \text{where} \quad \xi \sim \mathcal{N}(0,).$$

Since $\mathbf{x}_0$ and ξ are both sampled from Gaussian distributions, their linear combination $\mathbf{x}_t$ also forms a Gaussian distribution, i.e.,

$$\mathbf{x}_t \sim \frac{1}{K}\sum_{k=1}^{K}\mathcal{N}(\alpha_t\boldsymbol{\mu}_k, (\sigma_k^2\alpha_t^2 + \sigma_t^2)\cdot).$$

Then, we have

$$\begin{aligned}
\nabla p(\mathbf{x}_t) &= \frac{1}{K}\sum_{i=1}^{K}\nabla_{\mathbf{x}_t}\left[\frac{1}{2}\left(\frac{1}{\sqrt{2\pi\sigma_i^2\alpha_t^2+\sigma_t^2}}\right)\cdot\exp\left(-\frac{1}{2}\left(\frac{\mathbf{x}_t-\boldsymbol{\mu}_i\alpha_t}{\sigma_i^2\alpha_t^2+\sigma_t^2}\right)^2\right)\right] \\
&= \frac{1}{K}\sum_{i=1}^{K} p_i(\mathbf{x}_t)\cdot\nabla_{\mathbf{x}_t}\left[-\frac{1}{2}\left(\frac{\mathbf{x}_t-\boldsymbol{\mu}_i\alpha_t}{\sigma_k^2\alpha_t^2+\sigma_t^2}\right)^2\right] \\
&= \frac{1}{K}\sum_{i=1}^{K} p_i(\mathbf{x}_t)\cdot\frac{-(\mathbf{x}_t-\boldsymbol{\mu}_i\alpha_t)}{\sigma_k^2\alpha_t^2+\sigma_t^2}.
\end{aligned}$$

We can also calculate the score of $\mathbf{x}_t$, i.e.,

$$\nabla\log p(\mathbf{x}_t) = \frac{\nabla p(\mathbf{x}_t)}{p(\mathbf{x}_t)} = \frac{1/K\cdot\sum_{i=1}^{K} p_i(\mathbf{x}_t)\cdot\left(\frac{-(\mathbf{x}_t-\boldsymbol{\mu}_i\alpha_t)}{\sigma_k^2\alpha_t^2+\sigma_t^2}\right)}{1/K\cdot\sum_{i=1}^{K} p_i(\mathbf{x}_t)}.$$

E. A multi-mode Tweedie DPS toy: likelihood-norm certificate

We work in a high-dimensional multi-mode Gaussian mixture setting and focus on the *likelihood term* in the Tweedie–DPS score. We show that, under a simple separation condition on the mixture components and for small measurement noise, a small Euclidean norm of the *likelihood part* of the DPS score is sufficient to guarantee that the full approximate conditional score is close to the true one. Crucially, we do *not* assume *a priori* that x lies near a particular mode; instead, we deduce this from the smallness of the likelihood term.

Throughout this section $d \geq 1$ is fixed and all vectors live in $\mathbb{R}^d$.

E.1. Multi-mode point-mass model

Discrete prior on the clean signal. Let $\{\mu^{(1)}, \ldots, \mu^{(K)}\} \subset \mathbb{R}^d$ be K distinct mode vectors and let $\pi = (\pi_1, \ldots, \pi_K)$ be prior weights with $\pi_k > 0$ and $\sum_{k=1}^{K}\pi_k = 1$. The clean signal x_0 takes values in this finite set:

$$x_0 \in \{\mu^{(1)}, \ldots, \mu^{(K)}\}, \qquad \mathbb{P}(x_0 = \mu^{(k)}) = \pi_k. \tag{8}$$

Gaussian diffusion and measurement. Given x_0, we observe a diffused state x and a noisy measurement y via

$$x = x_0 + \sigma_t z, \qquad z \sim \mathcal{N}(0, I_d), \tag{9}$$

$$y = x_0 + \sigma_y \eta, \qquad \eta \sim \mathcal{N}(0, I_d), \tag{10}$$

with $\sigma_t > 0$ and $\sigma_y > 0$. Conditioned on mode k we have

$$x \mid (x_0 = \mu^{(k)}) \sim \mathcal{N}(\mu^{(k)}, \sigma_t^2 I_d), \tag{11}$$

$$y \mid (x_0 = \mu^{(k)}) \sim \mathcal{N}(\mu^{(k)}, \sigma_y^2 I_d). \tag{12}$$

We use the shorthand

$$\varphi_k(x) := \mathcal{N}(x; \mu^{(k)}, \sigma_t^2 I_d), \qquad \psi_k(y) := \mathcal{N}(y; \mu^{(k)}, \sigma_y^2 I_d).$$

E.2. Responsibilities, posterior means, and Jacobian

Prior responsibilities and Tweedie mean. The unconditional diffusion density is the mixture

$$p_t(x) = \sum_{k=1}^{K} \pi_k \varphi_k(x). \tag{13}$$

The posterior responsibility of mode k given x is

$$r_k(x) := \mathbb{P}(x_0 = \mu^{(k)} \mid x) = \frac{\pi_k \varphi_k(x)}{\sum_{j=1}^{K} \pi_j \varphi_j(x)}. \tag{14}$$

The Tweedie prior mean (posterior mean of x_0 given x) is

$$\bar{\mu}(x) := \mathbb{E}[x_0 \mid x] = \sum_{k=1}^{K} r_k(x)\, \mu^{(k)}. \tag{15}$$

By the Tweedie identity,

$$\bar{\mu}(x) = x + \sigma_t^2 \nabla_x \log p_t(x). \tag{16}$$

Joint responsibilities and posterior mean given (x, y). The joint density of (x, y) is

$$p_t(x, y) = \sum_{k=1}^{K} \pi_k\, \varphi_k(x) \psi_k(y),$$

and the joint responsibility of mode k given (x, y) is

$$r_k(x, y) := \mathbb{P}(x_0 = \mu^{(k)} \mid x, y) = \frac{\pi_k \varphi_k(x) \psi_k(y)}{\sum_{j=1}^{K} \pi_j \varphi_j(x) \psi_j(y)}. \tag{17}$$

The full posterior mean of x_0 given (x, y) is

$$\bar{\mu}(x, y) := \mathbb{E}[x_0 \mid x, y] = \sum_{k=1}^{K} r_k(x, y)\, \mu^{(k)}. \tag{18}$$

A key identity relating $r_k(x)$ and $r_k(x, y)$ is obtained by dividing numerator and denominator in (17) by $p_t(x)$:

$$r_k(x, y) = \frac{r_k(x)\, \psi_k(y)}{\sum_{j=1}^{K} r_j(x)\, \psi_j(y)}. \tag{19}$$

Covariance and Jacobian of the Tweedie mean. Let the posterior covariance of x_0 given x be

$$\Sigma(x) := \mathrm{Cov}(x_0 \mid x). \tag{20}$$

A standard computation (or differentiating (16)) yields the Jacobian of the Tweedie mean:

$$J_{\bar{\mu}}(x) := \frac{\partial \bar{\mu}(x)}{\partial x} = \frac{1}{\sigma_t^2}\, \Sigma(x). \tag{21}$$

In particular $\Sigma(x)$ is symmetric positive semidefinite, hence so is $J_{\bar{\mu}}(x)$.

We also define the covariance trace

$$\kappa(x) := \operatorname{tr} \Sigma(x). \tag{22}$$

Since $\Sigma(x) \succeq 0$, we have $\|\Sigma(x)\|_{\mathrm{op}} \le \kappa(x)$ and hence

$$\|J_{\bar{\mu}}(x)\|_{\mathrm{op}} \le \frac{\kappa(x)}{\sigma_t^2}. \tag{23}$$

E.3. True conditional score, DPS score, and likelihood term

True conditional score. Conditioned on mode k, (x, y) is jointly Gaussian with $x \mid (x_0 = \mu^{(k)}) \sim \mathcal{N}(\mu^{(k)}, \sigma_t^2 I_d)$ and $y \mid (x_0 = \mu^{(k)}) \sim \mathcal{N}(\mu^{(k)}, \sigma_y^2 I_d)$. A standard mixture calculation gives the true conditional score

$$s^\star(x) := \nabla_x \log p_t(x \mid y) = \frac{1}{\sigma_t^2}\big(\bar{\mu}(x, y) - x\big), \tag{24}$$

with $\bar{\mu}(x, y)$ as in (18).

Tweedie DPS approximate score and likelihood term. Standard Tweedie–DPS uses $\bar{\mu}(x)$ as the Tweedie estimator of x_0 and approximates the conditional likelihood $p_t(y \mid x)$ by a point mass at $\bar{\mu}(x)$:

$$\log p_t(y \mid x) = \log \mathbb{E}_{x_0|x}[p(y \mid x_0)] \approx \log p\big(y \mid \bar{\mu}(x)\big),$$

where $p(y \mid x_0) = \mathcal{N}(y; x_0, \sigma_y^2 I_d)$.

We define the *likelihood term* of the DPS score as

$$s^{\mathrm{lik}}(x) := \nabla_x \log p\big(y \mid \bar{\mu}(x)\big). \tag{25}$$

Since $\nabla_\mu \log p(y \mid \mu) = (1/\sigma_y^2)(y - \mu)$, the chain rule together with (21) gives the explicit form

$$s^{\mathrm{lik}}(x) = \frac{1}{\sigma_y^2} J_{\bar{\mu}}(x)\,\big(y - \bar{\mu}(x)\big) = \frac{1}{\sigma_t^2 \sigma_y^2}\,\Sigma(x)\big(y - \bar{\mu}(x)\big). \tag{26}$$

We refer to $\|s^{\mathrm{lik}}(x)\|_2$ as the *likelihood norm* in the DPS score.

The full Tweedie–DPS approximate conditional score is

$$\tilde{s}(x) := \nabla_x \log p_t(x) + s^{\mathrm{lik}}(x) = \frac{1}{\sigma_t^2}\big(\bar{\mu}(x) - x\big) + \frac{1}{\sigma_y^2} J_{\bar{\mu}}(x)\big(y - \bar{\mu}(x)\big). \tag{27}$$

Score approximation error. The error between the approximate and true conditional scores is

$$e(x) := \tilde{s}(x) - s^\star(x). \tag{28}$$

Combining (24) and (27) and regrouping terms yields the exact decomposition

$$e(x) = \frac{1}{\sigma_t^2}\big(\bar{\mu}(x) - \bar{\mu}(x, y)\big) + s^{\mathrm{lik}}(x). \tag{29}$$

Our goal is to control $\|\bar{\mu}(x) - \bar{\mu}(x, y)\|_2$ in terms of the likelihood norm $\|s^{\mathrm{lik}}(x)\|_2$, without assuming *a priori* that x lies near a particular mode.

E.4. Geometric separation and measurement consistency

Mode separation. Fix a reference mode index $k^\star \in \{1, \dots, K\}$ and write $\mu^\star := \mu^{(k^\star)}$. Define the minimal and maximal distances to other modes as

$$\Delta_{\min} := \min_{j \neq k^\star} \|\mu^{(j)} - \mu^\star\|_2 > 0, \;\; \Delta_{\max} := \max_{j \neq k^\star} \|\mu^{(j)} - \mu^\star\|_2. \tag{30}$$

Measurement consistency event. Fix a radius $r_y > 0$ such that

$$r_y \le \frac{1}{4}\,\Delta_{\min}. \tag{31}$$

We say that y is *consistent with* $\mu^\star$ if

$$\|y - \mu^\star\|_2 \le r_y. \tag{32}$$

Under the measurement model (10) with $x_0 = \mu^\star$, this event holds with probability $F_{\chi^2_d}(r_y^2/\sigma_y^2)$ and has probability close to one when σ_y is small.

On the event (32), for any $j \neq k^\star$ we have

$$\|y - \mu^{(j)}\|_2 \ge \|\mu^{(j)} - \mu^\star\|_2 - \|y - \mu^\star\|_2 \ge \Delta_{\min} - r_y \ge \frac{3}{4}\,\Delta_{\min}. \tag{33}$$

Since $\psi_k(y) \propto \exp(-\|y - \mu^{(k)}\|_2^2/(2\sigma_y^2))$, we obtain the likelihood ratio bound

$$\frac{\psi_j(y)}{\psi_{k^\star}(y)} = \exp\Big(-\frac{\|y - \mu^{(j)}\|_2^2 - \|y - \mu^\star\|_2^2}{2\sigma_y^2}\Big) \tag{34}$$

$$\le \exp\Big(-\frac{\Delta_{\min}^2}{4\sigma_y^2}\Big) =: \rho_y, \qquad j \neq k^\star, \tag{35}$$

where $\rho_y \in (0, 1)$ is exponentially small when the measurement noise is small.

E.5. Two ways small likelihood norm constrains the posterior

The DPS likelihood term admits the closed form

$$s^{\mathrm{lik}}(x) = \frac{1}{\sigma_t^2\sigma_y^2}\,\Sigma(x)\big(y - \bar\mu(x)\big), \tag{36}$$

where $\Sigma(x) = \mathrm{Cov}(x_0 \mid x)$ and $\bar\mu(x) = \mathbb{E}[x_0 \mid x]$. Taking norms and using $\|\Sigma(x)\|_{\mathrm{op}} \le \mathrm{tr}\,\Sigma(x) =: \kappa(x)$ yields

$$\|s^{\mathrm{lik}}(x)\|_2 \le \frac{\kappa(x)}{\sigma_t^2\sigma_y^2}\,\|y - \bar\mu(x)\|_2. \tag{37}$$

Equation (37) highlights two task-relevant quantities: (i) the *posterior dispersion* $\kappa(x)$ (how multi-modal $p(x_0 \mid x)$ remains), and (ii) the *measurement residual* $\|y - \bar\mu(x)\|_2$ (how well the Tweedie mean already matches the observation). Importantly, a small likelihood norm can occur for degenerate reasons (e.g., $\Sigma(x) \approx 0$ can suppress the term even when $\|y - \bar\mu(x)\|_2$ is large). For a certificate-style implication, we therefore work in a regime where **both** of the following conditions hold:

Condition A (covariance concentration). $\kappa(x) \le \kappa_0$ for a small threshold κ_0.

Condition B (measurement-consistent Tweedie mean). $\|y - \bar\mu(x)\|_2 \le r_y$ for a small radius r_y (the same r_y used in the measurement-consistency event).

We will show that under the measurement-consistency event $\|y - \mu^\star\|_2 \le r_y$ and in the *low-noise* regime (small σ_t relative to mode separation), Conditions A and B together force both $\bar\mu(x)$ and $\bar\mu(x, y)$ to concentrate near the *same* mode $\mu^\star$, which in turn makes the mean-gap term in (29) small.

For later use, we also define the global pairwise separation

$$\Delta_{\mathrm{sep}} := \min_{i \neq j} \|\mu^{(i)} - \mu^{(j)}\|_2 \;>\; 0. \tag{38}$$

E.6. Control of $\bar\mu(x)$ and $\bar\mu(x, y)$

We work throughout on the measurement-consistency event $\|y - \mu^\star\|_2 \le r_y$ (Definition (32)) with $r_y \le \Delta_{\min}/4$ (Eq. (31)), and assume that both Condition A and Condition B from Section E.5 hold at the queried x.

Condition B pins $\bar{\mu}(x)$ near $\mu^\star$. If $\|y - \bar{\mu}(x)\|_2 \le r_y$, then by the triangle inequality,

$$\|\bar{\mu}(x) - \mu^\star\|_2 \le \|\bar{\mu}(x) - y\|_2 + \|y - \mu^\star\|_2 \le 2r_y. \tag{39}$$

Condition A implies a dominant responsibility. Let $q_k := r_k(x)$ denote the posterior weights of x_0 given x, the trace admits the pairwise form

$$\kappa(x) = \frac{1}{2}\sum_{i=1}^{K}\sum_{j=1}^{K} q_i q_j \, \|\mu^{(i)} - \mu^{(j)}\|_2^2. \tag{40}$$

Using $\|\mu^{(i)} - \mu^{(j)}\|_2 \ge \Delta_{\text{sep}}$ for $i \ne j$,

$$\kappa(x) \ge \frac{\Delta_{\text{sep}}^2}{2}\sum_{i\ne j} q_i q_j = \frac{\Delta_{\text{sep}}^2}{2}\Big(1 - \sum_{k=1}^{K} q_k^2\Big). \tag{41}$$

Hence if $\kappa(x) \le \kappa_0$ and we define

$$\varepsilon_\kappa := \frac{2\kappa_0}{\Delta_{\text{sep}}^2}, \qquad (\text{assume } \varepsilon_\kappa < 1), \tag{42}$$

then $\sum_k q_k^2 \ge 1 - \varepsilon_\kappa$ and therefore there exists a dominant index

$$\hat{k} \in \arg\max_k q_k \quad \text{such that} \quad q_{\hat{k}} \ge 1 - \varepsilon_\kappa. \tag{43}$$

Moreover, since $\kappa(x) = \sum_k q_k \|\mu^{(k)} - \bar{\mu}(x)\|_2^2$ for a discrete posterior,

$$\|\mu^{(\hat{k})} - \bar{\mu}(x)\|_2^2 \le \frac{\kappa(x)}{q_{\hat{k}}} \le \frac{\kappa_0}{1 - \varepsilon_\kappa}. \tag{44}$$

Jointly: the dominant mode must be the measurement-consistent mode. Combining (39) and (44),

$$\|\mu^{(\hat{k})} - \mu^\star\|_2 \le \|\mu^{(\hat{k})} - \bar{\mu}(x)\|_2 + \|\bar{\mu}(x) - \mu^\star\|_2 \le \sqrt{\frac{\kappa_0}{1 - \varepsilon_\kappa}} + 2r_y.$$

Assume the thresholds satisfy the identification condition

$$\sqrt{\frac{\kappa_0}{1 - \varepsilon_\kappa}} + 2r_y < \Delta_{\min}. \tag{45}$$

Then necessarily $\hat{k} = k^\star$, hence

$$r_{k^\star}(x) = q_{k^\star} \ge 1 - \varepsilon_\kappa. \tag{46}$$

The joint posterior (x, y) is even more concentrated. Using (19) and the likelihood ratio bound (35),

$$\begin{aligned} 1 - r_{k^\star}(x, y) &= \sum_{j \ne k^\star} \frac{r_j(x)\psi_j(y)}{\sum_{\ell=1}^{K} r_\ell(x)\psi_\ell(y)} \\ &\le \frac{1}{r_{k^\star}(x)} \sum_{j \ne k^\star} r_j(x)\frac{\psi_j(y)}{\psi_{k^\star}(y)} \le \frac{\rho_y}{r_{k^\star}(x)}\big(1 - r_{k^\star}(x)\big) \\ &\le \frac{\rho_y}{1 - \varepsilon_\kappa}\,\varepsilon_\kappa. \end{aligned} \tag{47}$$

Therefore

$$\|\bar{\mu}(x, y) - \mu^\star\|_2 \le \big(1 - r_{k^\star}(x, y)\big)\Delta_{\max} \le \Delta_{\max}\frac{\rho_y}{1 - \varepsilon_\kappa}\,\varepsilon_\kappa. \tag{48}$$

A mean-gap bound under Conditions A & B. Finally, by the triangle inequality,

$$\|\bar{\mu}(x) - \bar{\mu}(x,y)\|_2 \le \|\bar{\mu}(x) - \mu^\star\|_2 + \|\bar{\mu}(x,y) - \mu^\star\|_2 \le 2r_y + \Delta_{\max}\frac{\rho_y}{1-\varepsilon_\kappa}\,\varepsilon_\kappa. \tag{49}$$

For completeness, we also record the mode-specific lower bound used later in the large-σ_t discussion: focusing on $k^\star$ in (40) and using $\|\mu^{(k^\star)} - \mu^{(j)}\|_2 \ge \Delta_{\min}$ for $j \ne k^\star$,

$$\kappa(x) \ge \frac{1}{2}\, r_{k^\star}(x)\big(1 - r_{k^\star}(x)\big)\Delta_{\min}^2. \tag{50}$$

E.7. Likelihood norm implies small score error

We now connect the likelihood norm $\|s^{\mathrm{lik}}(x)\|_2$ to the score approximation error $\|e(x)\|_2$ in the regime relevant to late (low-noise) timesteps.

Recall the exact decomposition (29):

$$e(x) = \frac{1}{\sigma_t^2}\big(\bar{\mu}(x) - \bar{\mu}(x,y)\big) + s^{\mathrm{lik}}(x).$$

Thus

$$\|e(x)\|_2 \le \frac{1}{\sigma_t^2}\|\bar{\mu}(x) - \bar{\mu}(x,y)\|_2 + \|s^{\mathrm{lik}}(x)\|_2. \tag{51}$$

The remaining work is to control the mean gap, which is achieved by Section E.6 under **both** Conditions A & B.

Theorem E.1 (Low-noise likelihood-norm certificate in the multi-mode toy)**.** *Consider the multi-mode point-mass model (8)–(10) and fix a reference mode $\mu^\star = \mu^{(k^\star)}$. Let $\Delta_{\min}, \Delta_{\max}$ be as in (30) and let Δ_{sep} be the global separation (38). Work on the measurement-consistency event $\|y - \mu^\star\|_2 \le r_y$ with $r_y \le \Delta_{\min}/4$.*

(Low-noise regime.) *Assume the diffusion noise level satisfies*

$$\sigma_t \le \frac{1}{8}\,\Delta_{\mathrm{sep}}. \tag{52}$$

(This assumption reflects that the certificate is intended for late/small-t states where posteriors are more mode-concentrated.)

Let $x \in \mathbb{R}^d$ be such that both Conditions A & B hold:

$$\kappa(x) = \operatorname{tr}\Sigma(x) \le \kappa_0, \qquad \|y - \bar{\mu}(x)\|_2 \le r_y,$$

and define $\varepsilon_\kappa := 2\kappa_0/\Delta_{\mathrm{sep}}^2$ as in (42). Assume in addition that

$$\kappa_0 \le \sigma_t^2. \tag{53}$$

(Under (52)–(53), we have $\varepsilon_\kappa \le 1/32 < 1$ and, since $\Delta_{\mathrm{sep}} \le \Delta_{\min}$ and $r_y \le \Delta_{\min}/4$, the identification condition (45) holds automatically.)

Then the score error satisfies

$$\|e(x)\|_2 \le \delta_{\mathrm{floor}}^{\mathrm{lik}} + \|s^{\mathrm{lik}}(x)\|_2, \tag{54}$$

where the (low-noise) floor term is

$$\delta_{\mathrm{floor}}^{\mathrm{lik}} := \frac{1}{\sigma_t^2}\left(2r_y + \Delta_{\max}\frac{\rho_y}{1-\varepsilon_\kappa}\,\varepsilon_\kappa\right), \qquad \rho_y := \exp\!\Big(-\frac{\Delta_{\min}^2}{4\sigma_y^2}\Big). \tag{55}$$

In particular, when σ_t is small and the thresholds (κ_0, r_y) are chosen so that ε_κ and r_y are small (as is typical in the late/noise-decayed regime), the floor $\delta_{\mathrm{floor}}^{\mathrm{lik}}$ is small and the bound becomes informative: $\|e(x)\|_2 \approx \|s^{\mathrm{lik}}(x)\|_2$.

Proof. Under (52) and (53),

$$\varepsilon_\kappa = \frac{2\kappa_0}{\Delta_{\text{sep}}^2} \le \frac{2\sigma_t^2}{\Delta_{\text{sep}}^2} \le \frac{2}{64} = \frac{1}{32},$$

so $\varepsilon_\kappa < 1$. Moreover, since $\Delta_{\text{sep}} \le \Delta_{\min}$, we also have $\sigma_t \le \Delta_{\min}/8$. Using $r_y \le \Delta_{\min}/4$ and $\varepsilon_\kappa \le 1/32$,

$$\sqrt{\frac{\kappa_0}{1-\varepsilon_\kappa}} + 2r_y \le \frac{\sigma_t}{\sqrt{1-\varepsilon_\kappa}} + \frac{\Delta_{\min}}{2} \le \sigma_t\sqrt{\frac{32}{31}} + \frac{\Delta_{\min}}{2} < \Delta_{\min},$$

so the identification condition (45) holds and hence the mean-gap bound (49) applies. Combining (51) with (49) yields (54). □

E.8. Implications for norm control and scope

Why norm control is a principled intervention. Theorem 3.2 provides a certificate depending only on an *already-computed* quantity $\|s^{\text{lik}}(x)\|_2$. This justifies host-agnostic mechanisms that either (i) shrink this norm locally before a host step, or (ii) select among stochastic branches those with smaller windowed aggregates. The intervention is solver-agnostic: it only assumes access to a host likelihood-score estimator.

Why the low-noise window is the right place to measure. The certificate tightens when responsibilities concentrate and the floor becomes small, which is most plausible when the trajectory is already near a single mode (more common at later / smaller t in separated mixtures). This matches the empirical observation that low-noise aggregates correlate more strongly with LPIPS than high-noise aggregates, without requiring any additional qualitative claims.

Scope Our guarantee is mechanistic: it is proved under Gaussian diffusion/observation and a stylized multi-mode prior, which is a standard tractable setup for diffusion-style posterior identities (Ho et al., 2020; Song et al., 2021). More importantly, the *certificate mechanism uses Tweedie structure* through Equation 3 and enables the implication "small norm ⇒ concentrated responsibilities ⇒ small mean gap". This matches a broad and practically important family of inverse solvers that approximate $p_t(y|x_t)$ by a *plug-in at a denoiser/Tweedie posterior-mean estimate* $\hat{x}_0(x_t, t)$, and backpropagate a data-fidelity term through it, including DPS and its scheduling refinements DCPS/MGPS (Chung et al., 2023; Janati et al., 2024; Moufad et al., 2025), as well as several restoration/projection solvers that rely on $\hat{x}_0$ as a posterior-mean surrogate (e.g., DDRM/DDNM/ΠGDM) (Kawar et al., 2022; Wang et al., 2023; Song et al., 2023b), and Tweedie-moment refinements such as TMPD (Boys et al., 2024). Recent physics/constrained conditional generation frameworks apply similar conditioning interfaces (Cheng et al., 2025; Huang et al., 2024a; Bastek et al., 2025; Li et al.; Jacobsen et al., 2024).

There exist alternative ways to estimate or incorporate $p_t(y|x_t)$ without this Tweedie plug-in, e.g., particle/SMC-based posterior samplers and Monte Carlo moves (Dou & Song, 2024; Cardoso et al., 2024), MCMC formulations that treat diffusion priors as plug-and-play priors (Wu et al., 2024). For such methods, the likelihood-score proxy need not admit a covariance–residual form, so Theorem 3.2 does not apply verbatim. Nevertheless, our *modules* remain well-defined whenever a host exposes a differentiable likelihood-score proxy: they only consume its norm, not its derivation. Extending certificate-type implications beyond Tweedie-style estimators (non-Gaussian likelihoods, nonlinear operators, particle-based estimators, and neural-score approximation error) is an important direction for future work.

F. Discussion, Limitations, and Future Work

What this paper establishes. This work studies a simple but actionable diagnostic quantity in diffusion inverse sampling: the norm of the host-provided likelihood-score estimator $\|\hat{g}_t\|_2 \approx \|\nabla_{x_t} \log p_t(y \mid x_t)\|_2$. Our empirical results show that low-noise (late-t) likelihood norms are a particularly reliable predictor of downstream reconstruction quality, and that explicitly *controlling* these norms via training-free wrappers (RS/LNP/NSS) improves performance across multiple hosts and tasks. On the theory side, we provide a concrete certificate mechanism (Section 3.2) showing that, in a tractable multi-mode model, a *small* likelihood norm implies *small conditional-score error* up to an explicit floor (Theorem 3.2), which motivates norm-control as a principled, solver-agnostic intervention rather than a heuristic add-on.

Computation–quality trade-off. Modules A/B/C introduce additional computation beyond the host sampler: RS performs an extra sampling pass, LNP adds a small number of inner correction steps per timestep, and NSS runs probe rollouts to score noise candidates. This trade-off is inherent to any method that performs explicit trajectory-level diagnostics or

selection. That said, the modules are *continuously tunable*: RS can be made cheaper by reducing the inverse depth and using late-start guidance; LNP typically works with $R_{\text{proj}}=1$ and a late-start threshold; NSS can be used sparsely (few selection times, short scoring windows, small candidate pools). Hence, the practical limitation is not that norm control requires excessive compute, but that the user must choose an operating point on a smooth compute–quality curve (Section 4, Table 11).

Why we emphasize low-noise norms, and what remains open at high noise. Our design and analysis prioritize low-noise (late-t) windows for two complementary reasons. First, empirically the correlation between likelihood norms and perceptual/accuracy metrics is strongest in this regime (Section 3), so norm control there is the most reliable signal-to-improvement pathway. Second, the certificate perspective (Eq. (5)) is most informative precisely when the likelihood norm meaningfully reflects conditional-score mismatch. Importantly, we do *not* claim that high-noise (early-t) likelihood norms are useless: we consistently observe weaker but non-trivial relationships at higher t, and RS ablations indicate that fully traversing or inverting the very high-noise region can change outcomes (Table 3). A principled treatment of how to exploit high-noise norms (e.g., adaptive schedules that decide *where* to spend computation and *which* timesteps to diagnose) is a clear direction for future work.

Scope of the analysis: Tweedie-type likelihood handling and beyond. Our theoretical development focuses on a Tweedie/denoiser-based likelihood construction, which yields closed-form identities and makes the norm–error mechanism explicit. This choice is defensible for two reasons. First, Tweedie-style substitutions are a central design pattern in widely used inverse solvers, including DPS-style guidance and its refinements, as well as moment-projected variants for linear inverse problems (Chung et al., 2023; Janati et al., 2024; Moufad et al., 2025; Boys et al., 2024). Second, Tweedie-type denoiser-to-posterior surrogates also appear broadly in conditioned generation and scientific inverse problems, where they serve as the computationally feasible bridge between learned priors and task likelihoods (Cheng et al., 2025; Huang et al., 2024a; Bastek et al., 2025; Li et al.; Jacobsen et al., 2024). Nevertheless, this is still a limitation of the *formal certificate*: the current theorem is proved for a Tweedie-style likelihood proxy under Gaussian diffusion/observation in a multi-mode toy model.

Crucially, the *method* does not inherit this restriction. RS/LNP/NSS only require that the host exposes a likelihood-score proxy (or a differentiable likelihood surrogate whose gradient plays that role), which is also available in non-Tweedie families such as projection/closed-form linear-Gaussian solvers (e.g., DDRM/DDNM/ΠGDM) (Kawar et al., 2022; Wang et al., 2023; Song et al., 2023b) and in probabilistic particle/MCMC formulations (Dou & Song, 2024; Cardoso et al., 2024; Wu et al., 2024). Extending the *certificate* to these alternative likelihood constructions (and clarifying when the norm remains a reliable proxy for score error) is a natural theoretical next step.

Host guarantees and what our wrappers do (and do not) claim. Our modules are deliberately *non-invasive* in that they do not retrain the diffusion prior, and they treat the host solver as a black box whose likelihood estimator is left intact. However, it is important to distinguish two notions of "not altering the solver": (1) At the implementation level, we preserve the host's core update rule BASESTEP and only add (i) restarts (RS), (ii) pre-step local corrections (LNP), or (iii) noise-level selection (NSS). (2) At the Markov-chain level, these interventions *do* change the overall transition kernel relative to the unwrapped host, so we do not claim that any host-specific asymptotic correctness guarantees automatically transfer without additional assumptions. This is not a weakness of the approach but a statement of scope: our primary target is improved reconstruction quality under fixed compute budgets in practical inverse problems, where even baseline solvers typically rely on approximations to the likelihood score. A rigorous characterization of how norm-control wrappers interact with exact samplers (e.g., SMC/MCMC methods) is an interesting and concrete future-work direction.

Breadth of evaluation and remaining coverage. We demonstrate consistent gains across multiple hosts and tasks (Table 2) and validate key design choices via ablations (Tables 3–11). Still, diffusion inverse solvers cover a wide design space (different operators, noise models, guidance parametrizations, and latent spaces), and we have not exhaustively evaluated every combination. Because our interface is minimal (Section 4), broader validation is primarily an engineering and benchmarking effort rather than a conceptual blocker; we view it as high-impact future work.

G. Additional Experimental Details

G.1. Implementation Details

Baselines and codebase. All baseline inverse solvers are run using their *public* implementations. We use the **DCPS** (Janati et al., 2024) codebase[1] as a common scaffold (it already contains reference implementations of DPS and DCPS), and we integrate additional hosts (Moufad et al., 2025; Zhang et al., 2025) based on their original codebases[2] on top of this scaffold to ensure consistent data preprocessing, operators, and evaluation. Unless otherwise stated, we keep the default hyper-parameters provided by the corresponding repositories and/or the DCPS "playground" configuration, and only modify task-specific parameters (e.g., operator $\mathcal{A}$ and observation noise σ_y).

Pretrained diffusion priors. For **FFHQ**, we use the pretrained diffusion prior released with DPS (Chung et al., 2023). For **ImageNet256**, we use the pretrained unconditional diffusion model from (Dhariwal & Nichol, 2021).

Likelihood-score estimator (Tweedie plug-in). Across all experiments in this paper, the likelihood guidance exposed to our wrappers is based on the standard *Tweedie-style* plug-in used by DPS-type inverse solvers: given (x_t, t), the denoiser predicts $\hat{x}_0(x_t, t)$ and we evaluate the measurement likelihood on $\hat{x}_0$. For the default Gaussian observation model $y = \mathcal{A}(x_0) + \varepsilon,\ \varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$, we use the surrogate log-likelihood

$$\log \widehat{p}_t(y \mid x_t) \ := \ -\frac{1}{2\sigma_y^2} \left\| y - \mathcal{A}\hat{x}_0(x_t, t) \right\|_2^2 + \text{const}, \tag{56}$$

and obtain the likelihood-score estimator by automatic differentiation,

$$\hat{g}_t(x_t; y, \mathcal{A}) \ := \ \nabla_{x_t} \log \widehat{p}_t(y \mid x_t) \ \approx \ \nabla_{x_t} \log p_t(y \mid x_t). \tag{57}$$

This is the same interface quantity used to form our windowed likelihood-norm aggregates and to drive LNP/NSS.

G.2. Evaluation Protocol

LPIPS. For each reported configuration (dataset, task, σ_y, method), we report the **mean LPIPS** computed over **100** evaluation instances. For stochastic samplers, each instance corresponds to an independent run with fresh injected noise. LPIPS is computed using the standard implementation of Zhang et al. (2018).

FID. When we report FID in ablations, it is computed on the corresponding set of generated reconstructions for that setting using the standard definition of Heusel et al. (2017).

Pearson correlation and Spearman correlation. Given paired samples $\{(\mathrm{Norm}_i(\mathcal{W}),\ \mathrm{LPIPS}_i)\}_{i=1}^R$ across R repeated runs (or, equivalently, paired samples $(\ell_{t_0,i}, \mathrm{LPIPS}_i)$ at a fixed timestep t_0), Pearson correlation measures linear association and is computed as the normalized covariance

$$\rho_{\mathrm{P}}(\mathcal{W}) := \mathrm{Corr}\big(\{\mathrm{Norm}_i(\mathcal{W})\}_{i=1}^R, \{\mathrm{LPIPS}_i\}_{i=1}^R\big) = \frac{\sum_{i=1}^R (\mathrm{Norm}_i(\mathcal{W}) - \overline{\mathrm{Norm}(\mathcal{W})})(\mathrm{LPIPS}_i - \overline{\mathrm{LPIPS}})}{\sqrt{\sum_{i=1}^R (\mathrm{Norm}_i(\mathcal{W}) - \overline{\mathrm{Norm}(\mathcal{W})})^2}\sqrt{\sum_{i=1}^R (\mathrm{LPIPS}_i - \overline{\mathrm{LPIPS}})^2}}, \tag{58}$$

where bars denote empirical means over i. Spearman correlation measures monotonic association and is computed as Pearson correlation of the ranks:

$$\rho_{\mathrm{S}}(\mathcal{W}) := \mathrm{Corr}\big(\mathrm{rank}(\{\mathrm{Norm}_i(\mathcal{W})\}),\ \mathrm{rank}(\{\mathrm{LPIPS}_i\})\big) = \rho_{\mathrm{P}}\big(\mathrm{rank}(\{\mathrm{Norm}_i(\mathcal{W})\}),\ \mathrm{rank}(\{\mathrm{LPIPS}_i\})\big). \tag{59}$$

G.3. Synthetic Mixture-of-Gaussians Experiment

This subsection specifies the 2D synthetic experiment used to validate the likelihood-norm/score-error relationship in a setting with analytic ground truth.

[1] https://github.com/Badr-MOUFAD/dcps

[2] https://github.com/YazidJanati/mgps, https://github.com/zhangbingliang2019/DAPS

Prior. We consider a two-component, 2D distribution where the second coordinate is discrete and the first coordinate is Gaussian:

$$p(x) = \frac{1}{2}\mathcal{N}\big(x_1;\, -0.1, \sigma^2\big)\,\delta\big(x_2 + 1\big) + \frac{1}{2}\mathcal{N}\big(x_1;\, +0.1, \sigma^2\big)\,\delta\big(x_2 - 1\big), \tag{60}$$

where $\delta(\cdot)$ denotes the Dirac delta, so $(x_0)_2 \in \{-1, +1\}$ almost surely.

Observation. The observation reveals the second coordinate,

$$y \;=\; (x_0)_2 \in \{-1, +1\}, \tag{61}$$

i.e., $\mathcal{A}$ selects the second coordinate (a 1D linear operator).[3]

Solver and guidance scale. We implement a DPS-style guided sampler on this toy problem and set the DPS guidance scale to

$$\gamma \;=\; 1 \times 10^{-2}. \tag{62}$$

Since the full distribution is known, we can compute the relevant analytic quantities needed in the paper: (i) the true conditional score and (ii) the DPS-style conditional score estimate. We then record the likelihood-score norm and the conditional-score error norm along trajectories to produce the plots reported in the empirical section.

G.4. Default Hyper-parameters for Norm-Control Modules

This subsection summarizes the default hyper-parameters used by our training-free wrappers (unless explicitly ablated).

Module A (RS). Restart Sampling with Partial Inverse Encoding uses:

- inverse depth ratio $\tau_{\text{inv}} = 0.7$ (truncate inverse encoding at $t_{\text{inv}} = \lceil \tau_{\text{inv}} T \rceil$);
- inverse stochasticity $\eta_{\text{inv}} = 0$ (deterministic inverse encoding);
- late-start ratio $\tau_{\text{start}} = 0.7$ (enable guidance, and optionally LNP, primarily in the late/low-noise regime).

Module B (LNP). Likelihood-Norm Projection (LNP-MCMC) uses:

- projection gain $\gamma = 50$;
- inner rounds $R_{\text{proj}} = 1$ per timestep (one projection round is typically sufficient);
- Langevin step size $\eta_{\text{mcmc}} = 1/\gamma$;
- the same late-start threshold $\tau_{\text{start}} = 0.7$.

Module C (NSS). Likelihood-Norm Self-Rejection (NSS) is configured as a lightweight Monte Carlo selector:

- global selection by default, i.e., $\mathcal{T}_{\text{sel}} = \{T\}$;
- deterministic probe rollouts ($\eta_{\text{sol}} = 0$) for stable scoring;
- a low-noise scoring window $\mathcal{W}_{\text{low}}$ near the end of sampling (as defined in the main text 3.1);
- candidate count N set to 5 by default and ablated when relevant (larger N improves selection but increases cost).

G.5. Algorithms

The detailed algorithms of each modules (RS, LNP, NSS) are shown in Alg. 2, 3, and 4.

G.6. Additional Experiment Results

We report results on **FFHQ** and **ImageNet256** for three standard inverse tasks: *Center* and *Half* inpainting, and *SR* $4\times$ super-resolution. Observations are corrupted by additive Gaussian noise with $\sigma_y \in \{0.05,\, 0.3\}$. We use **DPS** (Chung et al., 2023), **DCPS** (Janati et al., 2024), **MGPS** (Moufad et al., 2025), and **DAPS** (Zhang et al., 2025) in our codebase as black-box samplers (Section 4). For each host, we report results for the base solver with three modules: A. Restart Sampling (RS), B. Noise Selection (NSS), and C. Langevin Projection (LNP).

[3] If additive observation noise is used in a particular run, we follow the same Gaussian likelihood as in (56).

Algorithm 2 RS (restart sampling via *partial* inverse encoding)

Require: observation y, operator $\mathcal{A}$ (passed to the host); host sampler $\mathcal{S}$; unconditional inverse encoder $\mathcal{S}^{-1}$; inverse depth ratio τ_{inv}; inverse noise η_{inv}; late-start ratio τ_{start} (enable guidance/LNP on a late window)

Round 1 (host): $x_0^{(1)} \leftarrow \mathcal{S}(y, \mathcal{A})$

Partial inverse encoding (unconditional): $z_{t_{\text{inv}}} \leftarrow \mathcal{S}^{-1}\big(x_0^{(1)}; \eta = \eta_{\text{inv}},\ t_{\text{inv}} = \lceil \tau_{\text{inv}} T \rceil\big)$

Round 2 (restart with late guidance): $x_0^{(2)} \leftarrow \mathcal{S}\big(y, \mathcal{A};\ \text{init} = z_{t_{\text{inv}}},\ \text{enable guidance/LNP for } t/T \in [0, \tau_{\text{start}}]\big)$

return $x_0^{(2)}$

Algorithm 3 LNP-MCMC (normalized Langevin projection toward bounded likelihood norms)

Require: observation y, operator $\mathcal{A}$; host BASESTEP
Require: LIKEEST, PRIORSCOREEST; timesteps $t = T, T-1, \ldots, 1$
Require: projection gain $\gamma > 0$, stability $\varepsilon > 0$
Require: late-start ratio $\tau_{\text{start}} \in (0, 1]$, rounds $R_{\text{proj}} \in \mathbb{N}$
Require: Langevin step size $\eta_{\text{mcmc}} > 0$

Initialize $x_T \sim \mathcal{N}(0, I)$
for $t = T, T-1, \ldots, 1$ **do**
  if $1 - t/T \geq \tau_{\text{start}}$ **then**
    for $r = 1, \ldots, R_{\text{proj}}$ **do**
      $\hat{s}_t \leftarrow \text{PRIORSCOREEST}(x_t, t)$
      $\ell_t \leftarrow \text{LIKEEST}(x_t, t;\ y, \mathcal{A})$
      $\hat{g}_t \leftarrow \nabla_{x_t} \log \ell_t$ ▷ $\approx \nabla_{x_t} \log p_t(y \mid x_t)$
      $\alpha_t \leftarrow \dfrac{\gamma}{\|\hat{g}_t\|_2 + \varepsilon}$
      $\tilde{s}_t \leftarrow \hat{s}_t + \alpha_t \hat{g}_t$
      Sample $\xi \sim \mathcal{N}(0, I)$
      $x_t \leftarrow x_t + \eta_{\text{mcmc}}\, \tilde{s}_t + \sqrt{2\eta_{\text{mcmc}}}\, \xi$
    end for
  end if
  Sample $\eta_t \sim \mathcal{N}(0, I)$
  $x_{t-1} \leftarrow \text{BASESTEP}(x_t,\ t \to t-1; \eta_t)$
  $x_t \leftarrow x_{t-1}$
end for
return x_0

G.7. LNP as Langevin refinement of a tempered diffusion posterior

This subsection gives a distributional interpretation of LNP and explains why it promotes likelihood-norm reduction.

A tempered (power) posterior at a fixed diffusion time. At diffusion time t, let $p_t(x_t)$ denote the unconditional diffusion marginal. Given a differentiable proxy $\text{LIKEEST}(x_t, t;\ y, \mathcal{A}) \approx p_t(y \mid x_t)$, LNP considers the tempered density

$$q_{t,\alpha}(x_t) \ \propto\ p_t(x_t)\, \text{LIKEEST}(x_t, t;\ y, \mathcal{A})^{\alpha}, \tag{63}$$

where $\alpha \geq 1$ is an inverse-temperature on the likelihood proxy (a Gibbs/power-posterior construction) (Geman & Hwang, 1986; Hajek, 1988). Increasing α sharpens the density around states that maximize LIKEEST (i.e., improve data consistency), while keeping $q_{t,\alpha}$ absolutely continuous w.r.t. p_t and therefore biased toward the diffusion manifold.

Gaussian likelihood proxy and why larger α reduces the likelihood-score norm. In the common linear-Gaussian setting $y = \mathcal{A}(x) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$, a standard residual-based proxy has the form

$$\text{LIKEEST}(x_t, t;\ y, \mathcal{A}) \ \propto\ \exp\Big(-\frac{\|\mathcal{A}(x_t) - y\|_2^2}{2\sigma_y^2}\Big) \qquad (\text{or } \mathcal{A}(\hat{x}_0(x_t, t)) \text{ in plug-in solvers}), \tag{64}$$

so the corresponding likelihood score is

$$\hat{g}_t(x_t) \ \approx\ \nabla_{x_t} \log \text{LIKEEST}(x_t, t;\ y, \mathcal{A}) = -\frac{1}{\sigma_y^2}\, J_{\mathcal{A}}(x_t)^{\top}\big(\mathcal{A}(x_t) - y\big). \tag{65}$$

Thus $\|\hat{g}_t(x_t)\|_2$ scales with the measurement residual $\|\mathcal{A}(x_t) - y\|_2$ (up to the operator/Jacobian):

$$\|\hat{g}_t(x_t)\|_2 \leq \frac{\|J_{\mathcal{A}}(x_t)\|_{\text{op}}}{\sigma_y^2}\, \|\mathcal{A}(x_t) - y\|_2. \tag{66}$$

Algorithm 4 NSS: ODE-style likelihood-norm self-rejection at selected timesteps

Require: observation y, operator $\mathcal{A}$, timesteps $t = T, T-1, \dots, 1$
Require: $\textsc{BaseStep}(x_t,\ t \to t-1; \eta)$ (host reverse step)
Require: probe runner $\mathcal{S}$ (batched rollout scoring over window $\mathcal{W}$)
Require: candidate count N, selection set $\mathcal{T}_{\text{sel}}$, scoring window $\mathcal{W}$, probe noise scale $\eta_{\text{sol}} \in [0, 1]$
Initialize $x_T \sim \mathcal{N}(0, I)$
for $t = T, T-1, \dots, 1$ **do**
  if $t \in \mathcal{T}_{\text{sel}}$ **then**
    $\mathcal{W}_t \leftarrow \{\, s \in \mathcal{W} :\ s \le t \,\}$
    Sample candidate noises $\xi := \{\eta^{(n)}\}_{n=1}^{N} \sim \mathcal{N}(0, I)$
    Set batch states $\chi := \{X_t^{(n)}\}_{n=1}^{N}$ with $X_t^{(n)} \leftarrow x_t$
    $\{S^{(n)}\}_{n=1}^{N} \leftarrow \mathcal{S}\big(\chi,\ t,\ \mathcal{W}_t,\ \eta_{\text{sol}},\ \xi\big)$
    $n^\star \leftarrow \arg\min_n S^{(n)}, \quad \eta_t \leftarrow \eta^{(n^\star)}$
  else
    Sample $\eta_t \sim \mathcal{N}(0, I)$
  end if
  $x_{t-1} \leftarrow \textsc{BaseStep}(x_t,\ t \to t-1; \eta_t)$
  $x_t \leftarrow x_{t-1}$
end for
return x_0

Raising LikeEst to a larger power α in (63) concentrates $q_{t,\alpha}$ more sharply on smaller-residual regions; in the consistent/noiseless limit it concentrates near the constraint set $\{x : \mathcal{A}(x) = y\}$, where the likelihood score vanishes. Consequently, sampling (even approximately) from $q_{t,\alpha}$ with larger α tends to produce states with smaller $\|\hat{g}_t(x_t)\|_2$, i.e., a smaller likelihood-norm signal.

Score and Langevin refinement. The score of (63) is

$$\nabla_{x_t} \log q_{t,\alpha}(x_t) = \nabla_{x_t} \log p_t(x_t) + \alpha\, \nabla_{x_t} \log \textsc{LikeEst}(x_t, t;\ y, \mathcal{A}). \tag{67}$$

Approximating $\nabla_{x_t} \log p_t(x_t)$ by the prior score estimator $\hat{s}_t(x_t)$ and $\nabla_{x_t} \log \textsc{LikeEst}$ by $\hat{g}_t(x_t)$ yields the practical score $\hat{s}_t(x_t) + \alpha\, \hat{g}_t(x_t)$. LNP then applies a short unadjusted Langevin refinement toward this tempered target (Erdogdu et al., 2018):

$$x_t \leftarrow x_t + \eta_{\text{mcmc}}\big(\hat{s}_t(x_t) + \alpha\, \hat{g}_t(x_t)\big) + \sqrt{2\eta_{\text{mcmc}}}\,\xi, \qquad \xi \sim \mathcal{N}(0, I). \tag{68}$$

Norm capping and the design $\eta_{\text{mcmc}} = 1/\gamma$. Rather than fixing α, LNP uses a state-dependent scaling α_t so that the likelihood drift is norm-capped:

$$\tilde{g}_t(x_t) := \alpha_t\, \hat{g}_t(x_t), \qquad \|\tilde{g}_t(x_t)\|_2 \le \gamma. \tag{69}$$

We couple this cap with $\eta_{\text{mcmc}} := 1/\gamma$. This has two direct consequences: (i) the deterministic likelihood displacement per inner step is uniformly bounded, $\|\eta_{\text{mcmc}}\tilde{g}_t(x_t)\|_2 \le 1$; and (ii) the injected noise has variance $2\eta_{\text{mcmc}} = 2/\gamma$, which decreases as γ increases. In particular, increasing γ allows a larger effective likelihood temperature (via the larger admissible α_t in (69)) while simultaneously reducing both discretization aggressiveness (through η_{mcmc}) and stochasticity (through $\sqrt{2\eta_{\text{mcmc}}}$), yielding a stable but more "projection-like" refinement toward high-likelihood (low-residual) regions.

Large-γ behavior and descent on likelihood norms. In the Gaussian proxy (64), ignoring the $O(1/\gamma)$ prior drift and the vanishing noise as γ grows, the refinement becomes a deterministic update along the likelihood-score direction, $x_t \leftarrow x_t + \eta_{\text{mcmc}}\tilde{g}_t(x_t)$. Since $\hat{g}_t(x_t) \approx \nabla_{x_t} \log \textsc{LikeEst}(x_t, t;\ y, \mathcal{A})$, this is gradient ascent on LikeEst (equivalently, gradient descent on the quadratic data-fidelity $\|\mathcal{A}(x_t) - y\|_2^2$), which directly reduces the residual and therefore reduces $\|\hat{g}_t(x_t)\|_2$ by (66). Moreover, when $\mathcal{A}$ is linear, $\hat{g}_t(x_t) = -(1/\sigma_y^2)\mathcal{A}^\top(\mathcal{A}x_t - y)$ and one can verify

$$\nabla_{x_t} \frac{1}{2}\|\hat{g}_t(x_t)\|_2^2 = -\frac{1}{\sigma_y^2}\,(\mathcal{A}^\top\mathcal{A})\,\hat{g}_t(x_t), \tag{70}$$

so stepping along $\hat{g}_t$ performs a preconditioned descent on the squared likelihood-score norm. This explains why, in practice, using a larger gain γ (hence allowing larger effective α_t while keeping each inner move stable via $\eta_{\text{mcmc}} = 1/\gamma$) can make the refinement aggressively drive down the likelihood norm, while the retained prior term $\hat{s}_t$ continues to regularize toward regions where the diffusion marginal p_t is reliable for downstream sampling.

Table 10. **LPIPS results** (lower is better) for different hosts and our norm-controlled modules A/B/C on FFHQ and ImageNet256. Columns correspond to task (*Center*, *Half*, *SR* 4×) and observation noise levels $\sigma_y \in \{0.05, 0.3\}$. Module A = RS, B = LNP, C = NSS.

Method	**FFHQ** (LPIPS ↓)						**ImageNet256** (LPIPS ↓)					
	Center		Half		SR 4×		Center		Half		SR 4×	
	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$
DPS	0.153	0.316	0.246	0.325	0.243	0.282	0.378	0.483	0.443	0.495	0.453	0.537
+A	0.149	0.274	0.219	0.288	0.218	0.259	0.343	0.437	0.408	0.461	0.390	0.450
+A+B	0.114	0.138	0.165	0.207	0.171	0.221	0.315	0.337	0.377	0.425	0.361	0.388
+A+B+C	**0.106**	0.131	0.159	0.174	0.156	0.208	0.300	<u>0.324</u>	0.361	0.397	0.342	0.371
DCPS	0.122	0.287	0.207	0.268	0.227	0.319	0.326	0.401	0.359	0.391	0.441	0.472
+A	0.113	0.225	0.177	0.231	0.182	0.284	0.301	0.363	0.322	0.356	0.404	0.410
+A+B	0.117	0.166	0.152	0.205	<u>0.143</u>	0.269	0.304	0.337	0.317	0.335	0.350	0.379
+A+B+C	<u>0.108</u>	0.153	**0.143**	0.197	**0.135**	0.257	**0.290**	0.328	**0.310**	<u>0.328</u>	<u>0.341</u>	<u>0.358</u>
MGPS	0.138	0.259	0.232	0.279	0.235	0.261	0.341	0.431	0.366	0.387	0.436	0.485
+A	0.125	0.234	0.180	0.247	0.193	0.217	0.326	0.388	0.330	0.342	0.387	0.428
+A+B	<u>0.108</u>	0.138	0.146	<u>0.172</u>	0.152	**0.164**	0.304	0.340	0.328	0.332	0.352	0.374
+A+B+C	0.109	<u>0.128</u>	<u>0.145</u>	**0.168**	0.147	0.168	<u>0.295</u>	0.331	<u>0.315</u>	**0.320**	**0.331**	**0.352**
DAPS	0.142	0.274	0.218	0.264	0.229	0.257	0.339	0.425	0.357	0.420	0.396	0.455
+A	0.137	0.248	0.189	0.233	0.204	0.229	0.323	0.401	0.338	0.396	0.392	0.418
+A+B	0.124	0.167	0.167	0.194	0.167	0.174	0.311	0.352	0.329	0.348	0.368	0.387
+A+B+C	0.115	**0.126**	0.151	0.181	0.150	<u>0.165</u>	0.306	**0.301**	0.316	0.331	0.347	0.360

Table 11. **Ablation of NSS scoring window, solver noise, and candidate count on FFHQ with DPS host**. We report LPIPS (lower is better). Module C = NSS.

Method	Center		SR 4×	
	$\sigma_y=0.05$	$\sigma_y=0.3$	$\sigma_y=0.05$	$\sigma_y=0.3$
DPS (baseline)	0.153	0.316	0.243	0.282
Baseline NSS	0.127	0.268	0.212	0.255
$\mathcal{W} = [0.5, 0.7]$	0.135	0.290	0.232	0.261
$\mathcal{W} = [0.3, 0.5]$	0.121	0.272	0.205	0.247
$\mathcal{W} = [0, 0.3]$	0.118	0.266	0.197	0.221
$\eta_{\text{sol}} = 0.1$	0.132	0.275	0.220	0.263
$\eta_{\text{sol}} = 0.3$	0.148	0.307	0.239	0.273
$N = 2$	0.142	0.305	0.238	0.280
$N = 16$	0.120	0.258	0.204	0.243

G.8. Ablating NSS design choices: low-noise scoring and deterministic rollouts are critical

Table 11 ablates key design choices for NSS on FFHQ with a DPS host. We evaluate how the scoring window, rollout stochasticity, and candidate count impact performance.

(i) Scoring in the low-noise regime is most effective. Scoring windows restricted to the low-noise regime consistently yield the best LPIPS improvements, validating our diagnostics in Section 3 where low-noise likelihood norms better correlate with final reconstruction quality.

(ii) Deterministic rollouts enable more reliable candidate evaluation. To score each noise candidate, NSS rolls out the corresponding future trajectory using a deterministic sampler and evaluates its aggregated likelihood norm. We observe that stochastic rollouts with larger noise (η_{sol}) degrades LPIPS. We congest that the stochastic rollouts weaken the correspondence between a noise candidate and its measured low-noise likelihood norm.

(iii) More candidates improve selection quality. Performance improves with larger candidate counts N, consistent with the expected behavior of a select mechanism operating on noisy proposals.

G.9. RS details: residual transfer and likelihood-norm control

This appendix makes explicit why RS tends to reduce the likelihood-score norm for common Tweedie-style hosts. Consider a linear-Gaussian observation model $y = \mathcal{A}x_0 + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$. Many guidance-based inverse solvers implement

(or approximate) a likelihood score through a denoiser prediction $\hat{x}_0(x_t, t)$, e.g., by differentiating a quadratic measurement loss $\ell_t(x_t) \propto -\frac{1}{2\sigma_y^2}\|y - \mathcal{A}\hat{x}_0(x_t, t)\|_2^2$. Let $J_t(x_t) := \partial \hat{x}_0(x_t, t)/\partial x_t$. A chain-rule calculation yields the proxy

$$\hat{g}_t(x_t) := \nabla_{x_t} \ell_t(x_t) \;\approx\; \frac{1}{\sigma_y^2}\, J_t(x_t)^\top \mathcal{A}^\top \Big(y - \mathcal{A}\hat{x}_0(x_t, t)\Big), \tag{71}$$

and therefore the norm bound

$$\|\hat{g}_t(x_t)\|_2 \;\le\; \frac{\|J_t(x_t)\|_{\mathrm{op}}\, \|\mathcal{A}\|_{\mathrm{op}}}{\sigma_y^2}\, \big\|y - \mathcal{A}\hat{x}_0(x_t, t)\big\|_2. \tag{72}$$

Now let $x_0^{(1)}$ be the first-pass host reconstruction and let $z_{t_{\mathrm{inv}}}$ be the truncated unconditional inverse encoding. Under inverse-consistency $\hat{x}_0(z_{t_{\mathrm{inv}}}, t_{\mathrm{inv}}) \approx x_0^{(1)}$, we obtain

$$\big\|y - \mathcal{A}\hat{x}_0(z_{t_{\mathrm{inv}}}, t_{\mathrm{inv}})\big\|_2 \le \big\|y - \mathcal{A}x_0^{(1)}\big\|_2 + \|\mathcal{A}\|_{\mathrm{op}}\, \big\|\hat{x}_0(z_{t_{\mathrm{inv}}}, t_{\mathrm{inv}}) - x_0^{(1)}\big\|_2, \tag{73}$$

which is small when the first pass is measurement-consistent and the truncated inversion is accurate. Combining (73) with (72) explains why RS produces a low-norm restart and why full inversion (high-noise) can break this mechanism by degrading inverse-consistency.