NORM-CONTROLLED LIKELIHOOD GUIDANCE FOR DIFFUSION-BASED INVERSE SOLVER

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion-based inverse solvers approximate the posterior by combining a pretrained diffusion prior with an approximate likelihood guidance term. For Tweedie plug-in likelihood-guided solvers, which evaluate measurement likelihoods on a Tweedie posterior-mean estimate $\hat{x}_0(x_t)$, we identify a clear diagnostic signal in the evaluated settings: for the same image and inverse task, different noise realizations lead to different reconstructions, and larger low-noise likelihood-guidance norms correlate with worse perceptual quality across the tested hosts and runs. This makes the *likelihood-score norm* an actionable empirical diagnostic for this solver family. To leverage this signal, we propose **Norm-Controlled Likelihood Guidance**, a training-free wrapper composed of three modules that steer the sampling trajectory toward smaller likelihood norms while retaining the pretrained diffusion prior term. Theoretically, under posterior-concentration and measurement-consistency conditions in a tractable multi-modal Tweedie plug-in model, we establish a low-noise certificate bounding the difference between likelihood-score and score-approximation-error norms by an explicit residual. Experimentally, NCLG improves the host solvers across *all* reported rows in our FFHQ and ImageNet-256 evaluations, with ablation studies and cost analyses further supporting the design and quantifying representative 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 $\varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$ denotes additive Gaussian 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 $p(y|x)$. These two components give a target 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 difficult to model explicitly from data due to their 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., 2020; 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.

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

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., 2022). Related work studies Monte Carlo loss guidance, policy-gradient guidance, non-differentiable rule guidance, and PDE-conditioned generation (Song et al., 2023b; 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 $\hat{x}_0(x_t) = x_t + \sigma_t^2 \nabla_{x_t} \log p_t(x_t)$ is the Tweedie estimate of $\mathbb{E}[x_0|x_t]$. Throughout this paper, we refer to this family as *Tweedie plug-in likelihood guidance*: the solver first forms a Tweedie posterior-mean estimate $\hat{x}_0(x_t)$, typically implemented by the denoiser/score model, then evaluates a task likelihood or data-fidelity term at $\hat{x}_0(x_t)$ and backpropagates it to 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 still rely on approximate Tweedie plug-in likelihood guidance whose error can vary along a trajectory. This motivates a complementary diagnostic for the same solver family: if a quantity exposed by the guidance term tracks when the approximation is unreliable, it can be used to guide sampling without redesigning the likelihood estimator. In this work, we empirically identify the likelihood-score norm as such a diagnostic signal in our evaluated Tweedie plug-in setting, turning likelihood guidance from a fixed heuristic into a measurable quantity 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 whose low-noise aggregates correlate with reconstruction quality across the evaluated datasets, tasks, and Tweedie plug-in solvers. Furthermore, under posterior-concentration and measurement-consistency conditions in a tractable multi-modal model, we bound the difference between the likelihood-score and Tweedie plug-in score-error norms by an explicit residual.
- Guided by this diagnostic and certificate logic, we propose Norm-Controlled Likelihood Guidance (NCLG), which exploits the intrinsic stochasticity of diffusion sampling to *refine toward* lower-norm trajectories: Restart Sampling (RS) provides a coarse restart into a favorable basin, Likelihood-Norm Stochastic Selection (NSS) performs candidate selection among stochastic branches using predicted low-noise norm aggregates, and Likelihood-Norm Projection (LNP) applies a capped Langevin refinement that reduces instantaneous norms while retaining the diffusion prior term.
- We evaluate the framework on FFHQ and ImageNet inverse problems, where NCLG improves recent Tweedie plug-in host baselines across the reported rows while exposing representative trade-offs between quality and compute; ablations further isolate critical design choices, including the importance of low-noise norm control and the complementary roles of NCLG modules.

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., 2020; 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 likelihood term. Operator-based methods include DDRM, DDNM, and ΠGDM (Kawar et al., 2022; Wang et al., 2022; Song et al., 2023a); general guidance methods include DPS, LGD, DPG, and SCG (Chung et al., 2022; Song et al., 2023b; Tang et al., 2024; Huang et al., 2024b). Some works adopt 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 focuses on the Tweedie plug-in family, while borrowing MCMC/MC-style refinement and selection ideas.

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 Tweedie plug-in likelihood guidance, which reduces likelihood evaluation to the data domain through $\hat{x}_0(x_t, t)$ (Chung et al., 2022; Boys et al., 2023). A key limitation of such plug-ins is that

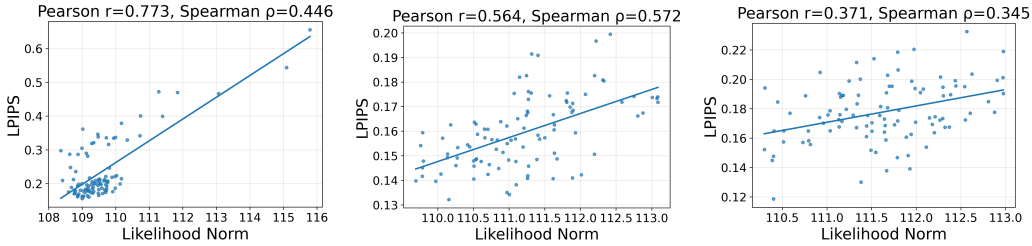

Figure 1: **Within-instance scatter: LPIPS vs. likelihood norm at $t = 0.3$.** Each point is one run with a different initial noise for the *same* image and task in FFHQ. Left: DPS (Chung et al., 2022); Middle: DCPS (Janati et al., 2024); Right: MGPS (MOUFAD et al., 2025). Larger likelihood norms are associated with larger LPIPS (worse), motivating norm-based candidate selection.

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 approximation error. Our work extends this line by proposing an external wrapper that controls the magnitude of the same likelihood proxy.

Diagnostics and alignment signals. Several works connect sample quality to score/noise alignment and provide post-hoc 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., 2026). While these are primarily post-hoc or sample-level, our approach studies a score-level diagnostic that is directly usable during sampling, as demonstrated next in Section 3.

3 Likelihood-Score Norm as a Diagnostic Proxy

In this section, we present the core methodological insight of our work: leveraging the norm of the estimated likelihood score as an empirical indicator of trajectory quality in Tweedie plug-in posterior sampling. Since the exact likelihood score $\nabla_{x_t} \log p(y|x_t)$ is typically intractable, practical diffusion inverse solvers often rely on differentiable Tweedie plug-in 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., 2022), $\hat{g}_t(x_t)$ 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 Gaussian measurement residual:

$$\hat{g}_t(x_t) = \nabla_{x_t} \log \widehat{p}_t(y \mid x_t) = \nabla_{x_t} \left[-\frac{1}{2\sigma_y^2} \left\| y - \mathcal{A}(\hat{x}_0(x_t)) \right\|_2^2 \right],$$

where $\widehat{p}_t(y \mid x_t)$ denotes the plug-in likelihood induced by the observation model $y = \mathcal{A}(x_0) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$. Such Tweedie plug-in likelihood-score proxy and its task-specific variants are 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 likelihood-score norm* $\mathrm{Norm}(\mathcal{W})$:

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

In what follows, we use this $\mathrm{Norm}(\mathcal{W})$ as a diagnostic proxy for Tweedie plug-in posterior diffusion solvers. We examine this proxy by measuring its correlation with reconstruction quality, identifying where it is most informative, and supporting the mechanism with a tractable 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 the evaluated Tweedie plug-in solvers (DPS, DCPS, MGPS). [1] 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 positively correlated with LPIPS across the tested solvers and runs. Since higher LPIPS values denote lower quality, higher norms are associated with degraded perceptual fidelity, supporting the use of the norm as a diagnostic signal in these settings.

Result 2: Low-Noise norms are more predictive. As shown in Table 1, while the aggregated likelihood-score norm still demonstrates positive correlation with LPIPS across various tasks and datasets, the reported correlations are substantially 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 suggests that the likelihood-score norm is more informative when measured in the low-noise regime.

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 10.

Task	σ_y	FFHQ		ImageNet	
		Low	High	Low	High
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

These findings suggest a design heuristic: *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 steps to improve reconstruction fidelity.

3.2 Theoretical Demonstration

To support the utility of the proposed likelihood score in subsection 3.1, we further analyze a 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 Tweedie plug-in approximation and the ground truth in Lemma 3.1. Theorem 3.2 then gives a bound relating its norm to the likelihood-score norm under low-noise, posterior-concentration, and measurement-consistency conditions.

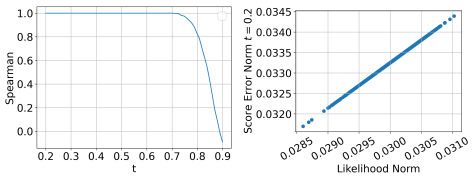

(a) (b)

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

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 $x_0 \in \{\mu^{(1)}, \ldots, \mu^{(K)}\} \subset \mathbb{R}^d$ 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, Tweedie plug-in estimated score and 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 Tweedie plug-in conditional score ("predict $\hat{x}_0$ via Tweedie, then*

[1] We adopt $t = 0.3$ as a practical low-noise boundary rather than a universal cutoff (see Appendix D.1 for sensitivity analysis to this choice).

apply likelihood") is

$$\tilde{s}(x) := \frac{1}{\sigma_t^2}\big(\bar{\mu}(x) - x\big) + s^{\mathrm{lik}}(x), \qquad s^{\mathrm{lik}}(x) := \nabla_x \log p(y|\bar{\mu}(x))\,. \tag{3}$$

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

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

Lemma 3.1 isolates a scaled mean-gap residual in addition to the likelihood-score proxy. Bounding this residual yields a relation between their norms. Figure 2 illustrates the corresponding empirical pattern: at low noise ($t < 0.3$), the likelihood-score norm and conditional-score error have a Spearman correlation near 1. The following theorem specifies conditions under which the residual can be bounded.

Theorem 3.2 (Likelihood-norm certificate; proof in Appendix F)**.** *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* $\Delta_{\max} := \max_{j\neq k^\star} \|\mu^{(j)} - \mu^\star\|_2$*, and the global separation* $\Delta_{\mathrm{sep}} := \min_{i\neq j} \|\mu^{(i)} - \mu^{(j)}\|_2 > 0$*. Assume* $\|y - \mu^\star\|_2 \le r_y \le \Delta_{\min}/4$*,* $\sigma_t \le \Delta_{\mathrm{sep}}/8$*, and* $\kappa_0 \le \sigma_t^2$*. For any state* x *satisfying both* $\operatorname{tr}\operatorname{Cov}(x_0 \mid x) \le \kappa_0$ *and* $\|y - \bar{\mu}(x)\|_2 \le r_y$*, let* $\ell(x) := \|s^{\mathrm{lik}}(x)\|_2$*. Then the conditional-score error satisfies*

$$\max\big\{0, \ell(x) - \delta^{\mathrm{lik}}_{\mathrm{floor}}\big\} \le \|e(x)\|_2 \le \ell(x) + \delta^{\mathrm{lik}}_{\mathrm{floor}}, \tag{5}$$

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

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

- **Vanishing Ambiguity Floor ($\delta^{\mathrm{lik}}_{\mathrm{floor}}$):** The floor term scales exponentially as $\mathcal{O}\big(\exp(-\Delta^2_{\min}/(4\sigma_y^2))\big)$. This exponential decay captures the rapid vanishing of mode ambiguity once the diffusion noise σ_t drops below the separation threshold.
- **Dominant Controllable Proxy ($\|s^{\mathrm{lik}}(x)\|_2$):** 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, when the floor is small, the bound tightens toward $\|e(x)\|_2 \approx \|s^{\mathrm{lik}}(x)\|_2$, supporting the likelihood norm as a diagnostic proxy for the intractable score error, matching the near-linear alignment observed at low t in Figure 2.

4 Norm-Controlled Likelihood Guidance

In this section, we introduce a plug-in framework for Tweedie plug-in solvers that expose a likelihood-score proxy through $\hat{x}_0(x_t)$, motivated by the model-specific certificate in Theorem 3.2 and the empirical correlations in Section 3. Specifically, *we exploit stochastic degrees of freedom in these solvers to search for or refine toward lower-norm sampling trajectories while leaving the host likelihood estimator intact.* 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. We treat the base inverse solver as a host that exposes $\textsc{BaseStep}(x_t, t \to t{-}1; \eta_t)$, where η_t is the injected Gaussian noise, $\textsc{PriorScoreEst}(x_t, t)$, and a differentiable likelihood proxy $\textsc{LikeEst}(x_t, t; y, \mathcal{A})$. We write $\hat{s}_t := \textsc{PriorScoreEst}(x_t, t)$ and $\hat{g}_t(x_t) := \nabla_{x_t} \log \textsc{LikeEst}(x_t, t; y, \mathcal{A})$ for the prior-score and likelihood-score estimates, respectively, and denote the first-pass host reconstruction by $x_0^{(1)}$. With this notation, RS uses a partial inverse encoder $\mathcal{S}^{-1}$ to construct a restart latent (Alg. 2); LNP uses short Langevin MCMC to refine x_t toward lower likelihood-score norm (Alg. 3); and NSS uses a probe runner $\mathcal{S}_{\mathrm{probe}}$ to score candidate noises by low-noise deterministic rollouts (Alg. 4).

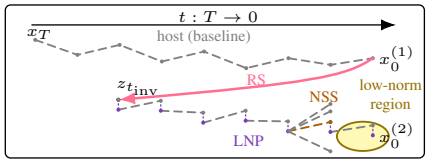

Figure 3: **Conceptual illustration.** RS restarts; LNP/NSS lower the likelihood norm.

4.1 LOCAL NORM CONTROL VIA REFINEMENT AND SELECTION

We first introduce two modules that 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.

Algorithm 1: **Norm-Controlled Likelihood Guidance**

Require: Timesteps $t = T, T-1, \dots, 1$.
Require: BASESTEP$(x_t, t \to t-1; \eta_t)$ and first-pass sample $x_0^{(1)}$.
Require: Observation y, operator $\mathcal{A}$; PRIORSCOREEST, LIKEEST.
Require: RS encoder $\mathcal{S}^{-1}$ with $\tau_{\text{inv}}, \eta_{\text{inv}}$; NSS probe $\mathcal{S}_{\text{probe}}$, $N, \mathcal{T}_{\text{sel}}, \mathcal{W}, \eta_{\text{sol}}$.
Require: LNP parameters $\gamma, \varepsilon, R_{\text{proj}}, \eta_{\text{mcmc}}, \tau_{\text{start}}$.
$t_{\text{inv}} \leftarrow \lceil \tau_{\text{inv}} T \rceil$
$x_{t_{\text{inv}}} \leftarrow \mathcal{S}^{-1}(x_0^{(1)}; \eta_{\text{inv}})$
$t_0 \leftarrow t_{\text{inv}}$
for $t = t_0, t_0-1, \dots, 1$ **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)
      $\hat{g}_t \leftarrow \nabla_{x_t} \log$ LIKEEST$(x_t, t;\, y, \mathcal{A})$
      $\alpha_t \leftarrow \gamma / (\|\hat{g}_t\|_2 + \varepsilon)$
      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 $t \in \mathcal{T}_{\text{sel}}$ **then**
    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

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 a corresponding Gibbs-style formulation (Geman & Hwang, 1986; Hajek, 1988; Erdogdu et al., 2018) that encourages diffusion-marginal consistency while favoring lower likelihood norm, and sample from it using Langevin dynamics. Specifically, at timestep t, 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 $\alpha_t := \gamma / (\|\hat{g}_t(x_t)\|_2 + \varepsilon)$ caps the likelihood drift by a user-defined projection gain γ, ensuring $\|\alpha_t \hat{g}_t(x_t)\|_2 \le \gamma$; see Appendix H.10 for details. To sample 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 $\eta_{\text{mcmc}} := 1/\gamma$. Here γ controls the trade-off between likelihood-norm reduction and preservation of the diffusion marginal, ε stabilizes the normalization, R_{proj} is the number of inner Langevin steps, and η_{mcmc} is the Langevin step size. In practice, we set γ to a moderate value (default $\gamma = 50$) to retain the influence of the prior score $\hat{s}_t$, which is intended to keep x_t close to the diffusion marginal $p(x_t)$ and support accurate downstream sampling. By default, τ_{start} restricts LNP to 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 noises that yield lower predicted low-noise norm aggregates. Specifically, at selected timesteps $t \in \mathcal{T}_{\text{sel}}$, NSS proposes N noise candidates $\{\eta^{(n)}\}_{n=1}^N$, uses $\mathcal{S}_{\text{probe}}$ to continue each candidate with probe noise scale η_{sol}, and assigns a score based on the *low-noise* norm aggregate:

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

where $\mathcal{W}$ is chosen in the low-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. We use deterministic probe rollouts by default ($\eta_{\text{sol}} = 0$), so the score reflects the candidate noise rather than additional rollout stochasticity.

LNP/NSS are local mechanisms that benefit from better global initialization. We highlight that LNP/NSS are *local* corrections 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 the complementary restart mechanism introduced next.

Table 2: **LPIPS results** (lower is better) for different hosts and NCLG on FFHQ and ImageNet256. Columns span tasks and $\sigma_y \in \{0.05, 0.3\}$; PSNR, SSIM, and FID show similar trends in Appendix Tables 6 and 7.

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%)

4.2 Global low-norm basin construction via restart

We now introduce Restart Sampling (RS), a *global* module for Tweedie plug-in solvers that aims to lower the likelihood norm by inverting a sampled x_0 back to a latent state x_t through the inverse of the denoising process, which then serves as a favorable initialization for subsequent local control. Prior work such as ReGuidance (Karan et al., 2025) empirically shows that reversing to the highest noise state x_T can yield more observation-consistent trajectories. Building on this idea, we design RS from the perspective of the likelihood-score norm, construct a partial restart state, and explain how norm-guided restarts can improve downstream control in this setting.

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 a deterministic inverse process (e.g., DDIM or ODE inversion) 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). The inverse step is written as $\mathcal{S}^{-1}(x_0^{(1)}; \eta_{\text{inv}})$ in the algorithms, where η_{inv} controls inverse-step stochasticity and is set to 0 by default for deterministic partial inversion; the same τ_{start} controls the late window in which local guidance/LNP is enabled after restart.

How RS enables effective local norm control. We formalize the theoretical motivation 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 the Tweedie plug-in model. For a deterministic inverse 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 the denoiser prediction for $\mathbb{E}[x_0|z_\tau]$, often via Tweedie's formula Efron (2011). Combining these yields:

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

where C is a Jacobian-scaled constant from the Tweedie plug-in. See Appendix. H.14 for the full derivation. Therefore, under this Tweedie plug-in approximation, a restart state yields a smaller likelihood-score norm, providing a favorable initialization for subsequent LNP/NSS refinement.

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 is because 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).

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 Tweedie plug-in inverse samplers. Our experiments address two questions: (i) whether norm control improves reconstruction quality across the evaluated solvers, and inverse tasks, and (ii) how the modules affect low-noise likelihood norms and reconstruction quality.

5.1 Experiment Setup

We report results on **FFHQ** and **ImageNet256** for three standard inverse tasks: *Center* and *Half* inpainting, and *SR* 4× super-resolution. Observations are corrupted by additive Gaussian noise with $\sigma_y \in \{0.05, 0.3\}$. We use **DPS** (Chung et al., 2022), **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), Likelihood-Norm Projection (LNP), and Likelihood-Norm Stochastic Selection (NSS). Table 2 reports LPIPS scores (lower is better), a perceptual metric that can serve as an empirical proxy for reconstruction quality, motivated by our observation that low-noise likelihood-score norms track perceptual degradation in the evaluated solvers. 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 H. Baseline configuration details and a DAPS tuning sanity check are in Appendix H.2.

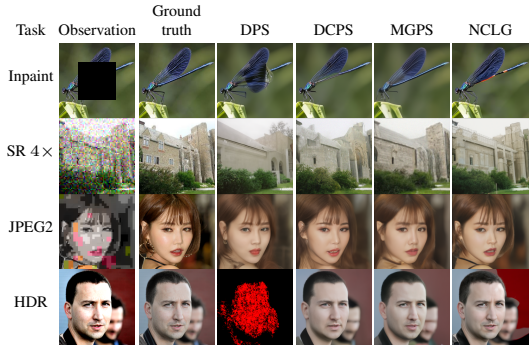

Figure 4: **Qualitative examples.** Rows show inpainting, SR 4×, JPEG2, and HDR; columns compare observation, ground truth, host solvers, and NCLG (DPS host).

5.2 Norm control improves the evaluated solvers

Table 2 summarizes LPIPS (lower is better) for four Tweedie plug-in solvers and their norm-controlled variants (+NCLG) across two datasets, three tasks, and two observation-noise levels. To broaden the task coverage beyond the standard inpainting/SR settings, Table 3 reports host-averaged LPIPS/PSNR/SSIM on ImageNet256 for Gaussian deblurring, motion deblurring, JPEG restoration (QF= 2), and HDR reconstruction, explicitly including nonlinear JPEG/HDR degradations; full per-host/FID results are in Appendix B, and settings are in Appendix H.1. Across the rows reported in Tables 2 and 3, NCLG improves the host, with relative LPIPS reductions ranging from 10% to 59% in Table 2. On the additional ImageNet degradations, NCLG reduces host-averaged LPIPS by 14.3%–26.6%, increases PSNR by 0.85–1.46 dB, and improves SSIM by 0.042–0.071.

Table 3: **Additional ImageNet256 tasks.** Host-averaged metrics; $\sigma_y = 0.05$, JPEG QF= 2.

Task	Variant	LPIPS↓	PSNR↑	SSIM↑
Gauss. Deblur	Host avg.	0.295	22.53	0.573
	+NCLG avg.	**0.253**	**23.76**	**0.627**
Motion Deblur	Host avg.	0.257	23.84	0.630
	+NCLG avg.	**0.216**	**24.70**	**0.672**
JPEG (QF=2)	Host avg.	0.347	21.73	0.549
	+NCLG avg.	**0.283**	**22.90**	**0.604**
HDR	Host avg.	0.276	22.39	0.616
	+NCLG avg.	**0.203**	**23.85**	**0.687**

Additional MRI/PnP-DM checks and DDRM/DDNM results are reported in Appendices G.3 and G.4, respectively. 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).

Table 4: **Computation cost and LPIPS.** FFHQ *Center*, $\sigma_y = 0.3$; time in s/sample, peak GPU in GB. Step-count rows increase the baseline sampling budget. RS/LNP/NSS rows are cumulative.

Host	Variant	Time [s] ↓	Peak GPU [GB] ↓	LPIPS ↓
DPS	baseline	**5.25**	**2.81**	0.316
	1200 steps	30.42	**2.81**	0.263
	RS	11.24	**2.81**	0.274
	RS+LNP	13.48	3.25	0.138
	RS+LNP+NSS	29.12	10.24	0.131
DCPS	baseline	14.86	**2.81**	0.287
	800 steps	63.25	**2.81**	0.241
	RS	28.51	**2.81**	0.225
	RS+LNP	30.27	3.25	0.166
	RS+LNP+NSS	48.28	10.24	0.153
MGPS	baseline	10.49	3.14	0.259
	800 steps	40.74	3.14	0.224
	RS	23.10	3.14	0.234
	RS+LNP	24.93	3.57	0.138
	RS+LNP+NSS	39.15	10.55	**0.128**

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 can provide larger perceptual gains in the evaluated settings. Finally, these results support that likelihood-norm control is *complementary* to schedule-level refinements such as DCPS/MGPS within the tested Tweedie plug-in guidance family: even when the host changes *when/how* guidance is injected, explicitly regulating the likelihood-score norm improves the reported reconstruction metrics.

Table 4 compares NCLG with extended-step baselines on FFHQ *Center* at $\sigma_y = 0.3$, measured on a 5090 laptop GPU. Full NCLG achieves 50.2%, 36.5%, and 42.9% lower LPIPS than 1,200-step DPS, 800-step DCPS, and 800-step MGPS, respectively, while taking less time in all three cases. RS+LNP

Table 5: **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 as a diagnostic. The **smallest** value in each column is bold; the second smallest is underlined. Module A = RS, B = LNP; NSS is evaluated separately in Table 21.

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**	**0.171**	**26.35**	**20.34**	**0.221**	**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

alone also achieves lower LPIPS in less time than these baselines; adding NSS further improves LPIPS but increases peak memory to 10.24–10.55 GB. ImageNet comparisons with 1,000- and 2,000-step hosts are in Appendix H.8. At comparable runtime, NCLG gives lower LPIPS than CAS (Hong et al., 2024), adapted misalignment-map selection (Boutin et al., 2023), Best-of-8, REGUIDANCE, random restarts, and data-consistency selection (Appendix H.9). Reducing NSS candidates or shortening its probe window lowers cost while retaining gains (Appendix Table 23).

5.3 RS AND LNP ARE COMPLEMENTARY

Table 5 ablates RS and LNP on FFHQ with a DPS host, and reports LPIPS/FID together with the low-t likelihood norm. We observe three trends linking the modules to these empirical diagnostics: **(i) RS reduces low-t likelihood norms and improves quality. (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 make reverse diffusion from x_0 effective. Injecting large inversion noise ($\eta_{\text{inv}} = 1.0$) or fully reversing to the initial timestep ($\tau_{\text{inv}} = 1.0$, REGUIDANCE (Karan et al., 2025)) violates this requirement, yielding negligible LPIPS gains and much higher likelihood-score norms; see Appendix H.13 for more discussion. **(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.

6 CONCLUSION

We introduced *Norm-Controlled Likelihood Guidance* (NCLG), a training-free wrapper for diffusion inverse-problem solvers that use Tweedie plug-in likelihood guidance. Empirically, we find that windowed *low-noise* aggregates of the likelihood-score norm are associated with perceptual reconstruction quality in the evaluated settings. To examine this diagnostic, we analyze a tractable multi-mode Gaussian-mixture toy with an exact conditional score. Under posterior-concentration and measurement-consistency conditions, our low-noise certificate bounds the difference between the likelihood-score and conditional-score-error norms by an explicit residual. 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, $4\times$ super-resolution, and additional degradations) and multiple Tweedie plug-in samplers, these modules show quality gains across the reported settings and representative quality–compute trade-offs in our experiments, 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) both the method design and certificate analysis are tied to Tweedie plug-in likelihood guidance, so extending norm-control ideas to alternative likelihood-score estimators remains an important direction; and (iii) while we observe useful signals beyond the low-noise window and gains on representative hosts and tasks, broader coverage and a sharper characterization of the high-noise regime remain open.

AI USE STATEMENT

We used a generative AI tool to polish the English prose and check citation wording and bibliographic metadata. The authors take responsibility for the final content.

REFERENCES

Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta numerica*, 28:1–174, 2019.

Jan-Hendrik Bastek, WaiChing Sun, and Dennis Kochmann. Physics-informed diffusion models. In *International Conference on Learning Representations*, volume 2025, pp. 3360–3385, 2025.

Victor Boutin, Thomas Fel, Lakshya Singhal, Rishav Mukherji, Akash Nagaraj, Julien Colin, and Thomas Serre. Diffusion models as artists: Are we closing the gap between humans and machines? *arXiv preprint arXiv:2301.11722*, 2023.

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

Gabriel Cardoso, Yazid Janati el idrissi, Sylvain Le Corff, and Eric Moulines. Monte carlo guided denoising diffusion models for bayesian linear inverse problems. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 44001–44037, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/bfb3c1db0deda449ff4027f77dd48704-Paper-Conference.pdf.

Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ode is provably fast. *Advances in Neural Information Processing Systems*, 36:68552–68575, 2023.

Sitan Chen, Vasilis Kontonis, and Kulin Shah. Learning general gaussian mixtures with efficient score matching. *arXiv preprint arXiv:2404.18893*, 2024.

Chaoran Cheng, Boran Han, Danielle Maddix, Abdul Fatir Ansari, Andrew Stuart, Michael W Mahoney, and Bernie Wang. Gradient-free generation for hard-constrained systems. In *International Conference on Learning Representations*, volume 2025, pp. 100510–100539, 2025.

Muthu Chidambaram, Khashayar Gatmiry, Sitan Chen, Holden Lee, and Jianfeng Lu. What does guidance do? a fine-grained analysis in a simple setting. *Advances in Neural Information Processing Systems*, 37:84968–85005, 2024.

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

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

Zehao Dou and Yang Song. Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *International Conference on Learning Representations*, volume 2024, pp. 12069–12095, 2024.

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

Murat A Erdogdu, Lester Mackey, and Ohad Shamir. Global non-convex optimization with discretized diffusions. *Advances in Neural Information Processing Systems*, 31, 2018.

Stuart Geman and Chii-Ruey Hwang. Diffusions for global optimization. *SIAM Journal on Control and Optimization*, 24(5):1031–1043, 1986.

Bruce Hajek. Cooling schedules for optimal annealing. *Mathematics of operations research*, 13(2): 311–329, 1988.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Chunsan Hong, ByungHee Cha, and Tae-Hyun Oh. Cas: A probability-based approach for universal condition alignment score. In *International Conference on Learning Representations*, volume 2024, pp. 32330–32351, 2024.

Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-solving under partial observation. *Advances in Neural Information Processing Systems*, 37: 130291–130323, 2024a.

Yujia Huang, Adishree Ghatare, Yuanzhe Liu, Ziniu Hu, Qinsheng Zhang, Chandramouli S Sastry, Siddharth Gururani, Sageev Oore, and Yisong Yue. Symbolic music generation with non-differentiable rule guided diffusion. *arXiv preprint arXiv:2402.14285*, 2024b.

Christian Jacobsen, Yilin Zhuang, and Karthik Duraisamy. Cocogen: Physically consistent and conditioned score-based generative models for forward and inverse problems. *SIAM Journal on Scientific Computing*, 47(2):C399–C425, 2025.

Yazid Janati, Badr Moufad, Alain Durmus, Eric Moulines, and Jimmy Olsson. Divide-and-conquer posterior sampling for denoising diffusion priors. *Advances in Neural Information Processing Systems*, 37:97408–97444, 2024.

Aayush Karan, Kulin Shah, and Sitan Chen. Reguidance: A simple diffusion wrapper for boosting sample quality on hard inverse problems. *arXiv preprint arXiv:2506.10955*, 2025.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in neural information processing systems*, 35:23593–23606, 2022.

Zeyu Li, Hongkun Dou, Shen Fang, Wang Han, Yue Deng, and Lijun Yang. Physics-aligned field reconstruction with diffusion bridge. In *International Conference on Learning Representations*, volume 2025, pp. 58864–58900, 2025.

Yingyu Liang, Zhizhou Sha, Zhenmei Shi, Zhao Song, Mingda Wan, and Yufa Zhou. Unraveling the smoothness properties of diffusion models: A gaussian mixture perspective. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 11436–11446. IEEE, 2025.

Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. In *International Conference on Learning Representations*, volume 2024, pp. 28027–28053, 2024.

Badr MOUFAD, Yazid Janati el idrissi, Lisa Bedin, Alain Oliviero Durmus, randal douc, Eric Moulines, and Jimmy Olsson. Variational diffusion posterior sampling with midpoint guidance. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 95170–95214, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/ed524bb14de1b52c8522b977ded241d3-Paper-Conference.pdf.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PmLR, 2021.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 10674–10685. ieee, 2022.

Kulin Shah, Sitan Chen, and Adam Klivans. Learning mixtures of gaussians using the ddpm objective. *Advances in Neural Information Processing Systems*, 36:19636–19649, 2023.

Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International conference on learning representations*, 2023a.

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

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

Łukasz Staniszewski, Łukasz Kuciński, and Kamil Deja. There and back again: On the relation between noise and image inversions in diffusion models. In *International Conference on Learning Representations*, volume 2026, pp. 20792–20838, 2026.

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

Haoyue Tang, Tian Xie, Aosong Feng, Hanyu Wang, Chenyang Zhang, and Yang Bai. Solving general noisy inverse problem via posterior sampling: A policy gradient viewpoint. In *International Conference on Artificial Intelligence and Statistics*, pp. 2116–2124. PMLR, 2024.

A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Math. Dokl.*, 5:1035–1038, 1963.

Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*, 2022.

Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine L Bouman. Principled probabilistic imaging using diffusion models as plug-and-play priors. *Advances in Neural Information Processing Systems*, 37:118389–118427, 2024.

Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. 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, 2025.

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

A Qualitative Results

We show two concrete examples comparing two solvers (DPS, DCPS) with our norm-controlled modules in Fig. 5. We also provide additional visualizations comparing other Tweedie plug-in solvers and their NCLG variants in Figs. 6 and 7.

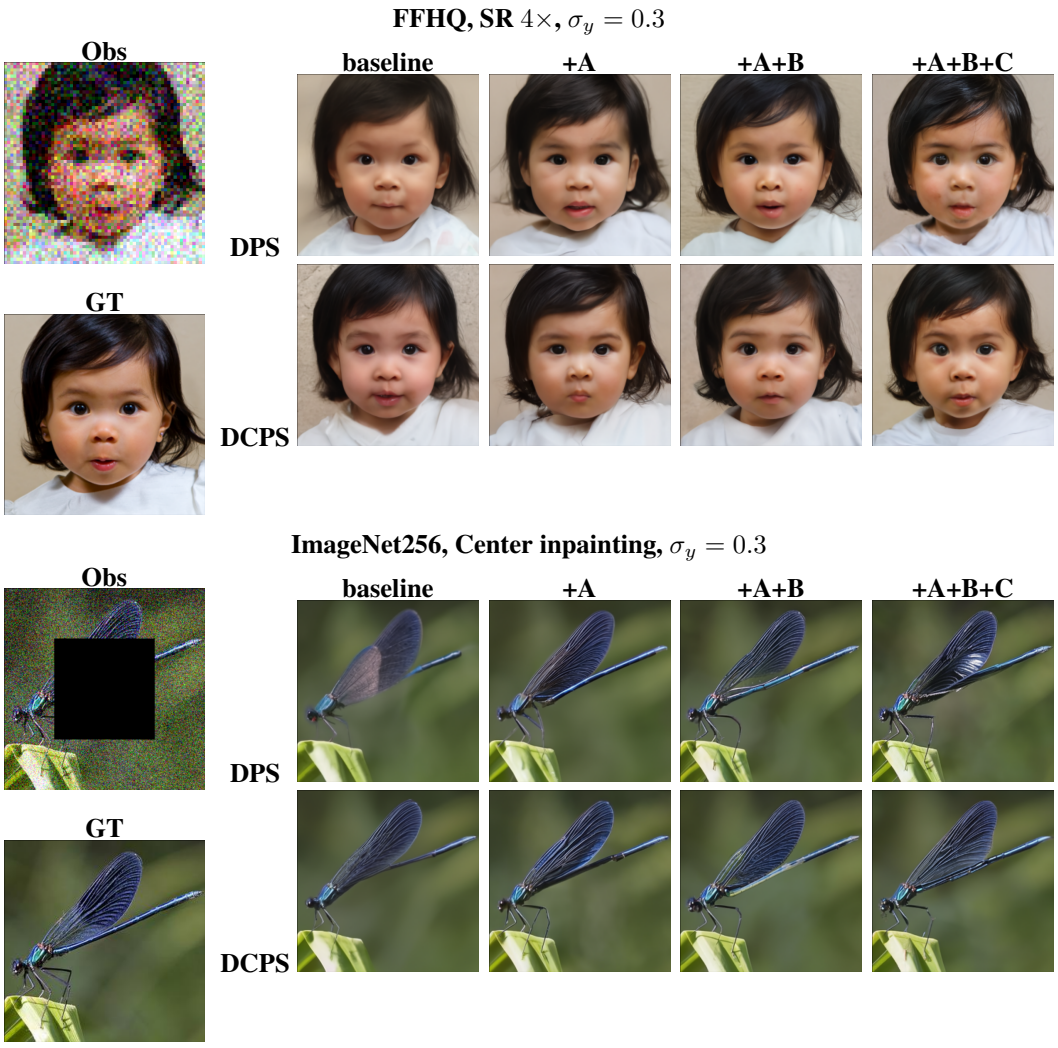

Figure 5: **Additional qualitative results (one example per task).** For each task, we show the noisy observation (Obs) and ground truth (GT) once (left). Reconstructions are shown for two hosts (rows: DPS, DCPS) under baseline and progressively adding our norm-control modules (columns: baseline, +A: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

B More Qualitative Results

In this section, we provide additional qualitative and quantitative results for NCLG. Tables 6 and 7 report metrics beyond LPIPS, including PSNR, SSIM, and FID. These results suggest that the observed gains are not limited to a single perceptual metric in the evaluated settings. We also evaluate Gaussian deblurring, motion deblurring, JPEG restoration (QF= 2), and HDR reconstruction; their settings are specified in Appendix H.1, with full metrics in Tables 8 and 9.

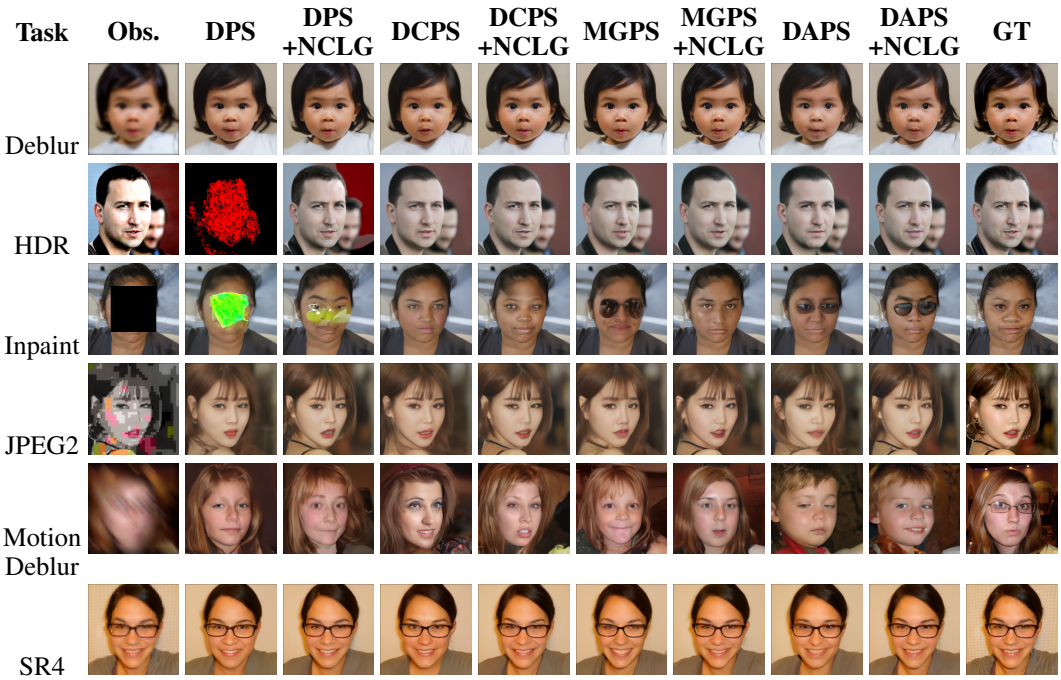

Figure 6: Additional qualitative comparisons on FFHQ. Columns show the observation, each host baseline, the corresponding host with NCLG, and the ground truth.

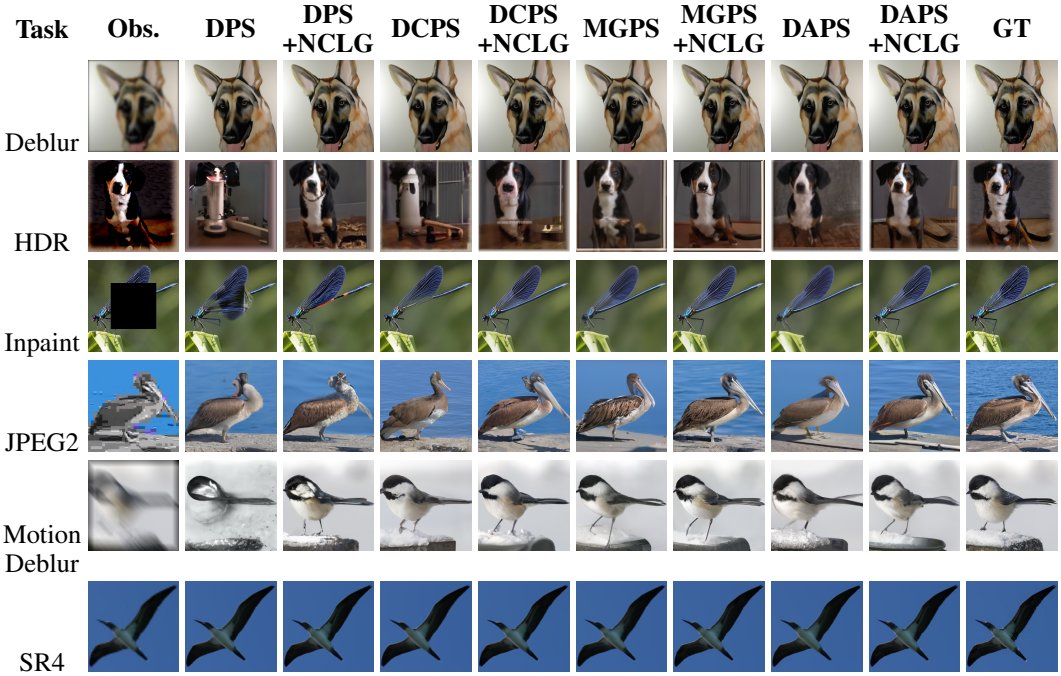

Figure 7: Additional qualitative comparisons on ImageNet. Columns show the observation, each host baseline, the corresponding host with NCLG, and the ground truth.

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., 2020; Nichol & Dhariwal,

Table 6: **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	16.20/0.55/31.55	15.72/0.66/44.63	13.84/0.52/55.72	22.46/0.65/29.51	18.72/0.58/32.45
+NCLG	23.17/0.79/26.92	20.55/0.74/28.46	17.05/0.75/31.10	16.62/0.70/36.84	24.98/0.75/27.24	21.08/0.69/30.75
DCPS	22.86/0.77/28.11	16.95/0.58/30.68	16.43/0.71/35.87	14.55/0.58/49.18	23.38/0.69/27.83	17.42/0.52/30.71
+NCLG	23.10/0.79/27.10	19.62/0.70/28.20	17.62/0.78/27.80	15.90/0.65/41.02	26.18/0.79/26.31	19.54/0.63/28.49
MGPS	22.48/0.76/29.23	17.38/0.61/30.94	15.98/0.68/39.68	14.28/0.56/46.35	23.09/0.68/28.90	19.22/0.61/29.78
+NCLG	23.05/0.79/27.24	20.72/0.75/28.85	17.58/0.79/28.10	16.95/0.72/34.96	25.70/0.77/26.85	22.02/0.73/27.53
DAPS	22.35/0.76/29.54	17.10/0.59/30.63	16.27/0.70/37.12	14.64/0.59/47.28	23.27/0.69/27.97	19.40/0.62/31.02
+NCLG	22.90/0.78/27.88	20.81/0.75/28.70	17.36/0.77/29.20	16.74/0.71/36.08	25.56/0.77/27.54	22.18/0.74/29.28

Table 7: **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

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) provides a common starting point for guidance-based diffusion inverse solvers, and many methods can be viewed as different approximations or parametrizations of the conditional likelihood term.

Existing solvers can be grouped by how they handle $\nabla_{x_t} \log p_t(y|x_t)$. *Operator-based methods* exploit measurement structure. DDRM (Kawar et al., 2022) and DDNM (Wang et al., 2022) target linear inverse problems; ΠGDM (Song et al., 2023a) uses pseudoinverse guidance and also handles nonlinear measurements. More recent moment-projected 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., 2023).

A second family, which we refer to as *general guidance*, uses data-fidelity or other conditioning objectives to steer a pretrained diffusion prior. Diffusion Posterior Sampling (DPS) differentiates an observation loss through the model's predictor $\hat{x}_0(x_t, t)$ (Chung et al., 2022). Loss-Guided Diffusion (LGD) uses Monte Carlo samples around a denoiser prediction to estimate an intractable loss-guidance term (Song et al., 2023b). Diffusion Policy Gradient (DPG) estimates guidance through a policy-gradient formulation (Tang et al., 2024); Stochastic Control Guidance (SCG) evaluates non-differentiable rules on sampled candidates (Huang et al., 2024b).

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

Table 8: **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 9: **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

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 focuses on the Tweedie plug-in guidance family in which the solver exposes (i) a base reverse step and (ii) a likelihood-score estimator built by evaluating the task likelihood at a Tweedie posterior-mean estimate $\hat{x}_0(x_t, t)$ (or a proxy thereto). Our modules then adjust the *norm* of this term without changing how it is estimated. Extending the same idea to closed-form projection or probabilistic sampler families requires checking their specific likelihood objects and transition kernels.

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:

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

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., 2022). LGD estimates a smoothed loss-guidance term, DPG uses a policy-gradient estimator, and SCG evaluates non-differentiable rules on sampled candidates (Song et al., 2023b; Tang et al., 2024; Huang et al., 2024b); their estimators need not follow the chain-rule form in (7).

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. Moment-projected 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., 2023). 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 & 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 are consistent with this view in a limited sense: for a fixed image and task, the aggregated likelihood-norm along a trajectory can be dominated by early high-noise steps and can correlate with downstream metrics such as LPIPS, SSIM, and PSNR. In this sense, large-t segments can be influential for guidance and 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, pixelwise absolute differences 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.

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., 2026). 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.

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.

Table 10: **Pearson correlations between LPIPS and windowed likelihood-norm means (appendix).** Each entry reports Pearson correlation ρ_P followed by a representative per-step mean in parentheses, $\overline{\mathrm{Norm}}(\mathcal{W}) := |\mathcal{W}|^{-1} \sum_{t \in \mathcal{W}} \|\hat{g}_t(x_t)\|_2^2$. Low-noise and High-noise use timestep windows $\mathcal{W}_{\mathrm{low}}$ and $\mathcal{W}_{\mathrm{high}}$ (defined in text); All (Full) averages over the entire reverse trajectory. For a fixed window shared by all runs, dividing the sum by its number of timesteps leaves the correlation unchanged.

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)

Misalignment maps (Boutin et al., 2023), condition alignment scores (Hong et al., 2024), inversion-bias studies (Staniszewski et al., 2026), and scheduling methods (Janati et al., 2024; MOUFAD et al., 2025) examine different aspects of conditioning and sampling. Our work studies the *norm* of the estimated likelihood term along the reverse trajectory and tests its use in guiding inverse-problem sampling.

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 Tweedie plug-in likelihood estimator provided by DPS and related refinements 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 empirically smaller in the evaluated settings. This interface matches DPS and Tweedie plug-in refinements such as DCPS/MGPS, which expose a differentiable plug-in likelihood-score proxy. Our experiments in later sections show that explicitly controlling the likelihood-norm in this way can improve reported reconstruction quality across the evaluated diffusion inverse problems, while preserving the core likelihood estimator used by these guidance schemes.

D Additional Empirical Diagnostics

D.1 Late/low-noise boundary sensitivity

We use $t = 0.3$ as a practical boundary for the late/low-noise regime, rather than a universal or theorem-derived optimum. This choice is supported by three observations: low-noise aggregates are more correlated with LPIPS than high-noise aggregates in Table 1; the MoG visualization shows stronger late-time norm–error alignment, while the certificate additionally requires posterior concentration and measurement consistency; and the NSS ablation in Table 21 favors low-noise scoring windows over middle windows. As an additional sensitivity check, Table 11 reports Pearson correlations between LPIPS and the instantaneous likelihood-score norm at different fixed timesteps in the same DPS diagnostic setting used for the main correlation study. The correlation remains positive over the tested timesteps but decreases as the measurement is taken earlier in the reverse trajectory, which supports using $t = 0.3$ as a practical cutoff for the low-noise diagnostic.

Table 11: **Sensitivity of the fixed-time norm diagnostic.** Pearson correlation between LPIPS and $\ell_t = \|\hat{g}_t(x_t)\|_2$ at different timesteps in the DPS diagnostic setting.

t	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Pearson	0.773	0.684	0.652	0.551	0.432	0.215	0.174

E SCORE OF A DIFFUSED GAUSSIAN MIXTURE

Let $x_0 \in \mathbb{R}^d$ follow an equally weighted mixture with K components,

$$x_0 \sim \frac{1}{K}\sum_{k=1}^{K}\mathcal{N}(\mu_k, \sigma_k^2 I_d), \qquad x_t = \alpha_t x_0 + \sigma_t \xi, \qquad \xi \sim \mathcal{N}(0, I_d),$$

where ξ is independent of x_0, $\sigma_k^2 > 0$, and $\sigma_t > 0$. Conditioned on component k, x_t is Gaussian with mean $\alpha_t \mu_k$ and variance $v_{k,t} := \alpha_t^2 \sigma_k^2 + \sigma_t^2$ in each coordinate. Its marginal remains a mixture:

$$p_t(x) = \frac{1}{K}\sum_{k=1}^{K}\phi_{k,t}(x), \qquad \phi_{k,t}(x) = (2\pi v_{k,t})^{-d/2}\exp\left(-\frac{\|x - \alpha_t\mu_k\|_2^2}{2v_{k,t}}\right).$$

Differentiating each component gives $\nabla_x \phi_{k,t}(x) = \phi_{k,t}(x)(\alpha_t\mu_k - x)/v_{k,t}$. Consequently, defining the component responsibilities $r_{k,t}(x) := \phi_{k,t}(x)/\sum_{j=1}^{K}\phi_{j,t}(x)$, we obtain

$$\nabla_x \log p_t(x) = \frac{\sum_k \nabla_x \phi_{k,t}(x)}{\sum_k \phi_{k,t}(x)} = \sum_{k=1}^{K} r_{k,t}(x)\frac{\alpha_t\mu_k - x}{v_{k,t}}.$$

F A MULTI-MODE TWEEDIE PLUG-IN TOY: LIKELIHOOD-NORM CERTIFICATE

We work in a high-dimensional multi-mode Gaussian mixture setting and focus on the *likelihood term* in the Tweedie plug-in score. Under separation, posterior-concentration, and measurement-consistency conditions, we derive a relation between the likelihood-score norm and conditional-score error. We first present the identity-observation case and then extend the argument to general linear observations in Section F.8.

Throughout, signal vectors live in $\mathbb{R}^d$ with $d \geq 1$; the linear extension allows measurements in $\mathbb{R}^m$.

F.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

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

F.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

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

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

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

F.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 plug-in approximate score and likelihood term. Standard Tweedie plug-in guidance uses $\bar{\mu}(x)$ as the Tweedie estimator of x_0 and approximates the conditional distribution $p(x_0 \mid x)$ 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 plug-in score.

The full Tweedie plug-in 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.

F.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

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

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

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

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

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

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

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

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

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

F.5 Concentration and measurement-consistency conditions

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 $\|\Sigma(x)\|_{\mathrm{op}} \leq \mathrm{tr}\,\Sigma(x) =: \kappa(x)$ yields

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

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). $\kappa(x) \leq \kappa_0$ for a small threshold κ_0.

Condition B (measurement-consistent Tweedie mean). $\|y-\bar{\mu}(x)\|_2 \leq r_y$ 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 $\|y-\mu^\star\|_2 \leq r_y$ 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

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

F.6 Control of $\bar{\mu}(x)$ and $\bar{\mu}(x,y)$

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

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

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

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 $\|\mu^{(i)}-\mu^{(j)}\|_2 \geq \Delta_{\mathrm{sep}}$ for $i \neq j$,

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

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,

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

F.7 Relation between likelihood norm and score error

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

Recall the exact decomposition (29):

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

The reverse triangle inequality gives

$$\big|\|e(x)\|_2 - \|s^{\text{lik}}(x)\|_2\big| \le \frac{1}{\sigma_t^2}\|\bar{\mu}(x) - \bar{\mu}(x, y)\|_2. \tag{50}$$

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

Theorem F.1 (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*

$$\sigma_t \le \frac{1}{8}\,\Delta_{\text{sep}}. \tag{51}$$

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

$$\kappa_0 \le \sigma_t^2. \tag{52}$$

(Under (51)–(52)*, we have* $\varepsilon_\kappa \le 1/32 < 1$ *and, since* $\Delta_{\text{sep}} \le \Delta_{\min}$ *and* $r_y \le \Delta_{\min}/4$*, the identification condition* (45) *holds automatically.)*

Let $\ell(x) := \|s^{\text{lik}}(x)\|_2$*. Then the score error satisfies*

$$\max\{0, \ell(x) - \delta_{\text{floor}}^{\text{lik}}\} \le \|e(x)\|_2 \le \ell(x) + \delta_{\text{floor}}^{\text{lik}}, \tag{53}$$

where the explicit residual bound is

$$\delta_{\text{floor}}^{\text{lik}} := \frac{1}{\sigma_t^2}\left(2r_y + \Delta_{\max}\,\frac{\rho_y}{1-\varepsilon_\kappa}\,\varepsilon_\kappa\right), \qquad \rho_y := \exp\!\Big(-\frac{\Delta_{\min}^2}{4\sigma_y^2}\Big). \tag{54}$$

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* $\delta_{\text{floor}}^{\text{lik}}$ *is small and the bound becomes informative:* $\|e(x)\|_2 \approx \|s^{\text{lik}}(x)\|_2$*.*

Proof. Under (51) and (52),

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

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 the reverse triangle inequality (50) with the mean-gap bound (49) shows that the two norms differ by at most $\delta_{\text{floor}}^{\text{lik}}$. Together with nonnegativity of $\|e(x)\|_2$, this yields the bound (53). □

F.8 Extension to general linear observation operators

Retain the finite-support prior and Gaussian diffusion above, and let

$$y = Ax_0 + \sigma_y\eta, \qquad A \in \mathbb{R}^{m\times d}, \quad \eta \sim \mathcal{N}(0, I_m), \tag{55}$$

with observation noise independent of the diffusion noise. The matrix A may be rank-deficient. The Tweedie mean $\bar\mu(x)$ and covariance $\Sigma(x)$ are unchanged, and the plug-in likelihood term becomes

$$s_A^{\text{lik}}(x) := \nabla_x \log p\big(y \mid \bar\mu(x)\big) = \frac{\Sigma(x)A^\top\big(y - A\bar\mu(x)\big)}{\sigma_t^2\sigma_y^2}. \tag{56}$$

Write $\bar\mu_A(x,y) := \mathbb{E}[x_0 \mid x, y]$ for the posterior mean under (55). If $\tilde s_A(x) := \nabla_x \log p_t(x) + s_A^{\text{lik}}(x)$ and $e_A(x) := \tilde s_A(x) - \nabla_x \log p_t(x \mid y)$, Gaussian diffusion gives the same exact decomposition:

$$e_A(x) = s_A^{\text{lik}}(x) + \frac{\bar\mu(x) - \bar\mu_A(x,y)}{\sigma_t^2}. \tag{57}$$

Observation classes preserve unresolved modes. Partition the support into classes $\mathcal{C}_1, \dots, \mathcal{C}_G$ such that $A\mu^{(k)} = a_g$ for every $k \in \mathcal{C}_g$, with distinct a_g across classes. At a fixed state x, define

$$w_g := \sum_{k \in \mathcal{C}_g} r_k(x), \qquad m_g := \frac{1}{w_g} \sum_{k \in \mathcal{C}_g} r_k(x)\mu^{(k)},$$
$$\psi_g(y) := \exp\left(-\frac{\|y - a_g\|_2^2}{2\sigma_y^2}\right).$$

All w_g are positive. Conditioning on y updates only the class weights:

$$w_g^y = \frac{w_g \psi_g(y)}{\sum_h w_h \psi_h(y)}, \qquad \bar{\mu}(x) = \sum_g w_g m_g, \qquad \bar{\mu}_A(x, y) = \sum_g w_g^y m_g. \tag{58}$$

Within each class, the relative probabilities $r_k(x)/w_g$ remain unchanged after conditioning on y. Thus, the argument permits ambiguity along $\ker(A)$ and does not require concentration on a single clean-signal mode.

For $G \geq 2$, fix the generating observation class $g^\star$ and write $a_\star := a_{g^\star}$. Define

$$\Delta_{A,\min} := \min_{g \neq g^\star} \|a_g - a_\star\|_2, \qquad \Delta_{A,\text{sep}} := \min_{g \neq h} \|a_g - a_h\|_2 > 0,$$

and let $D := \max_{i,j} \|\mu^{(i)} - \mu^{(j)}\|_2$ be the diameter of the clean support. Concentration is measured in observation space by $\kappa_A(x) := \operatorname{tr}\big(A\Sigma(x)A^\top\big)$.

Proposition F.2 (Certificate for linear observations)**.** *Assume $\|y - a_\star\|_2 \leq r_y \leq \Delta_{A,\min}/4$ and choose $0 \leq \kappa_{A,0} \leq \Delta_{A,\text{sep}}^2/64$. For any state x satisfying*

$$\kappa_A(x) \leq \kappa_{A,0}, \qquad \|y - A\bar{\mu}(x)\|_2 \leq r_y,$$

let $\ell_A(x) := \|s_A^{\text{lik}}(x)\|_2$ and define

$$\varepsilon_A := \frac{2\kappa_{A,0}}{\Delta_{A,\text{sep}}^2} \leq \frac{1}{32}, \qquad \rho_A := \exp\left(-\frac{\Delta_{A,\min}^2}{4\sigma_y^2}\right),$$
$$\delta_{A,\text{floor}} := \frac{D\varepsilon_A}{\sigma_t^2}\left(1 + \frac{\rho_A}{1 - \varepsilon_A}\right).$$

Then

$$\max\{0, \ell_A(x) - \delta_{A,\text{floor}}\} \leq \|e_A(x)\|_2 \leq \ell_A(x) + \delta_{A,\text{floor}}. \tag{59}$$

Proof. Since $Am_g = a_g$, the observation-space covariance has the pairwise form

$$\kappa_A(x) = \frac{1}{2}\sum_{g,h} w_g w_h \|a_g - a_h\|_2^2 \geq \frac{\Delta_{A,\text{sep}}^2}{2}\left(1 - \sum_g w_g^2\right).$$

Thus, some dominant class $\hat{g}$ has $w_{\hat{g}} \geq 1 - \varepsilon_A$. Also, $\|a_{\hat{g}} - A\bar{\mu}(x)\|_2^2 \leq \kappa_{A,0}/(1 - \varepsilon_A)$. Measurement consistency therefore gives

$$\begin{aligned}\|a_{\hat{g}} - a_\star\|_2 &\leq \sqrt{\frac{\kappa_{A,0}}{1 - \varepsilon_A}} + 2r_y \\ &\leq \frac{\Delta_{A,\text{sep}}}{8\sqrt{31/32}} + \frac{\Delta_{A,\min}}{2} < \Delta_{A,\min}.\end{aligned}$$

Hence $\hat{g} = g^\star$ and $w_{g^\star} \geq 1 - \varepsilon_A$. The separation and measurement event imply $\psi_g(y)/\psi_{g^\star}(y) \leq \rho_A$ for $g \neq g^\star$, exactly as in (35). Applying (58) yields

$$1 - w_{g^\star}^y \leq \rho_A \frac{1 - w_{g^\star}}{w_{g^\star}} \leq \frac{\rho_A \varepsilon_A}{1 - \varepsilon_A}.$$

Every class mean lies in the convex hull of the clean support, so $\|m_g - m_{g^\star}\|_2 \leq D$. Consequently,

$$\|\bar{\mu}(x) - m_{g^\star}\|_2 \leq D\varepsilon_A,$$
$$\|\bar{\mu}_A(x, y) - m_{g^\star}\|_2 \leq \frac{D\rho_A\varepsilon_A}{1 - \varepsilon_A}.$$

The triangle inequality bounds the mean gap by $D\varepsilon_A(1 + \rho_A/(1 - \varepsilon_A))$. Using (57) and the reverse triangle inequality, and clipping the lower endpoint at zero, proves (59). □

Relation to the low-noise conditions and rank deficiency. The covariance condition follows, for example, from $\operatorname{tr}\Sigma(x) \le \kappa_0 \le \sigma_t^2$ and $\|A\|_{\mathrm{op}}\sigma_t \le \Delta_{A,\mathrm{sep}}/8$, since $\kappa_A(x) \le \|A\|_{\mathrm{op}}^2 \operatorname{tr}\Sigma(x)$. Separation is required only between distinct observation classes; modes in the same class may remain unresolved. If $G = 1$, the likelihood is constant on the entire clean support, so $\bar{\mu}_A(x, y) = \bar{\mu}(x)$ and $s_A^{\mathrm{lik}}(x) = e_A(x) = 0$ directly. This extension covers linear masking and blur/downsampling operators under the stated model and state conditions. For instance, with $A = [1\ 0]$, modes sharing the first coordinate belong to the same class, while their distinct second coordinates retain their relative conditional weights. The proposition concerns the local score error; it does not establish sampling diversity or reconstruction accuracy for an approximate sampler. Nonlinear JPEG/HDR operators require additional analysis of the differentiable surrogate and its Jacobian; their reconstruction results here are supported by the experiments.

F.9 Implications for norm control and scope

Why norm control is a model-motivated intervention. Theorem F.1 places the realized conditional-score error in an explicit interval around the likelihood-score norm on states satisfying its assumptions. At fixed t, y, κ_0, r_y, two such states whose likelihood norms differ by more than $2\delta_{\mathrm{floor}}^{\mathrm{lik}}$ have the same ordering in score-error norm. Differences within this tolerance do not determine the ordering. This pointwise relation, together with the empirical correlations in Section 3, motivates local refinement and selection using low-noise norm aggregates; the theorem itself does not establish a trajectory-level or reconstruction-quality guarantee.

When the low-noise diagnostic is informative. The certificate is informative when concentration and measurement consistency control the scaled mean-gap residual. These state conditions are additional to the low-noise assumption and do not follow from a small likelihood norm alone. The empirical preference for low-noise windows is supported by their stronger correlations with LPIPS in the evaluated settings.

Scope The guarantee is mechanistic and model-specific: it is proved under Gaussian diffusion, linear observation with Gaussian noise (including rank-deficient operators in Section F.8), and a stylized multi-mode prior, which is a standard tractable setup for diffusion-style posterior identities (Ho et al., 2020; Song et al., 2020). The *certificate mechanism uses Tweedie structure* through Equation 3: assumed concentration and measurement consistency control the mean-gap residual, yielding the relation. This provides a model-specific motivation for methods that approximate $p_t(y|x_t)$ by a *plug-in at a 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., 2022; Janati et al., 2024; MOUFAD et al., 2025); it does not establish the theorem's state conditions or residual width for their real-image trajectories. Related methods use operator-based restoration or pseudoinverse guidance (e.g., DDRM/DDNM/ΠGDM) (Kawar et al., 2022; Wang et al., 2022; Song et al., 2023a), or moment projection such as TMPD (Boys et al., 2023). Recent physics and constrained conditional generation frameworks use other conditioning mechanisms (Cheng et al., 2025; Huang et al., 2024a; Bastek et al., 2025; Li et al., 2025; Jacobsen et al., 2025); the theorem here does not automatically cover these approaches.

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. Adapting norm-control modules to these settings would require checking how their likelihood objects interact with the transition kernel. Extending certificate-type implications beyond Tweedie plug-in estimators (non-Gaussian likelihoods, nonlinear operators, particle-based estimators, and neural-score approximation error) remains an important direction for future work.

G Discussion, Limitations, and Future Work

Diagnostic and theoretical scope. This work studies the norm of the host-provided likelihood-score estimator in Tweedie plug-in diffusion inverse sampling: $\|\hat{g}_t\|_2 \approx \|\nabla_{x_t} \log p_t(y \mid x_t)\|_2$. Section 3 reports the empirical relation between low-noise likelihood norms and reconstruction quality.

The norm–error bound in Theorem 3.2 applies under posterior-concentration and measurement-consistency conditions in a tractable multi-mode model; it does not give a generic guarantee for the evaluated solvers.

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. The computational cost can be adjusted through the module parameters: RS can be made cheaper by reducing the inverse depth and using late-start guidance; LNP typically works with $R_{\mathrm{proj}}{=}1$ and a late-start threshold; NSS can be used sparsely (few selection times, short scoring windows, small candidate pools). The host-step scaling results in Table 19 and the comparable-budget alternatives in Table 20 quantify this trade-off: the reported NCLG configurations improve the quality–compute operating point on these tasks, while NSS adds a tunable runtime cost.

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 supported operating regime in our study. Second, the model-specific certificate requires low noise together with posterior concentration and measurement consistency (Theorem 3.2). We 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 5). A more systematic 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 plug-in likelihood handling. Our theoretical development and current method design focus on Tweedie plug-in likelihood guidance, where a likelihood or data-fidelity term is evaluated at a Tweedie posterior-mean estimate $\hat{x}_0(x_t)$. This construction yields closed-form identities and makes the norm–error mechanism explicit. Tweedie plug-in substitutions are a central design pattern in widely used inverse solvers such as DPS and its scheduling refinements; moment-projected methods provide a related approach for linear inverse problems (Chung et al., 2022; Janati et al., 2024; MOUFAD et al., 2025; Boys et al., 2023). Conditioned generation and scientific inverse methods also use learned generative models with data- or physics-consistency signals, though their conditioning mechanisms differ (Cheng et al., 2025; Huang et al., 2024a; Bastek et al., 2025; Li et al., 2025; Jacobsen et al., 2025). Our theory covers a Tweedie plug-in likelihood proxy with Gaussian diffusion, linear observations with Gaussian noise (including rank-deficient operators; Appendix F.8), and a point-mass multi-mode prior under concentration and measurement-consistency conditions; our implemented RS/LNP/NSS modules use the same kind of differentiable plug-in likelihood interface. Other families such as operator-based restoration and pseudoinverse guidance (e.g., DDRM/DDNM/ΠGDM) (Kawar et al., 2022; Wang et al., 2022; Song et al., 2023a) and probabilistic particle/MCMC formulations (Dou & Song, 2024; Cardoso et al., 2024; Wu et al., 2024) expose different likelihood objects and transition kernels. We evaluate an adaptation based on native observation-correction norms for DDRM/DDNM in Appendix G.4. Establishing corresponding guarantees for their native transitions and other likelihood constructions remains a theoretical direction.

G.1 When likelihood-norm control can be unreliable

The likelihood-score norm is a proxy, not a standalone guarantee. It is most informative when the host likelihood proxy is reasonably specified and the current state x_t remains typical under the diffusion marginal p_t. The ablation in Table 5 already shows one such failure mode: full inversion $+\mathrm{A}(\tau_{\mathrm{inv}} = 1.0) + \mathrm{B}$ can reduce the low-t norm relative to the DPS baseline while producing worse LPIPS, suggesting that inverse consistency and marginal typicality matter.

Table 12 reports a controlled misspecification study with three proxy families: an incorrect observation-noise variance (IN), Poisson corruption scored by a Gaussian proxy (WN), and a Huber data term. RS and NSS remain robust in these settings, but aggressive LNP can become misleading under IN/WN: the proxy norm decreases after RS+LNP, while LPIPS worsens. This supports our

design choice of partial RS, capped late-time LNP, and optional NSS rather than unconstrained norm minimization.

Table 12: **Proxy misspecification study.** LPIPS: lower is better; the norm is the low-noise likelihood-score aggregate; z-score is the typicality diagnostic from Appendix G.2.

Proxy	Variant	LPIPS ↓	Norm ↓	z-score ↓
IN	baseline	0.452	1046.5	2.14
	+RS	0.394	238.5	0.93
	+RS+LNP	0.437	204.7	1.41
	+RS+NSS	0.384	224.8	0.76
WN	baseline	0.427	965.7	1.97
	+RS	0.395	127.6	0.88
	+RS+LNP	0.456	123.2	1.56
	+RS+NSS	0.387	124.3	0.72
Huber	baseline	0.304	45.7	0.48
	+RS	0.147	13.5	0.16
	+RS+LNP	0.122	11.2	0.09
	+RS+NSS	0.114	10.7	0.05

G.2 Typicality diagnostic for proxy failures

The failure pattern above can be detected by checking whether the current latent remains typical under the diffusion marginal p_t. We use a time-t version of a multi-noise reconstruction-error diagnostic: perturb x_t to a slightly noisier level $u > t$, denoise it back to t, and measure the round-trip error

$$r_t(x_t) := \mathbb{E}_{u,\epsilon}\left[\|\Pi_t(F_{t\to u}(x_t,\epsilon)) - x_t\|_2^2\right], \tag{60}$$

where $F_{t\to u}$ is the forward noising map and Π_t denotes denoising back to time t with the same pretrained diffusion prior. We standardize this quantity using on-marginal statistics:

$$z_t(x_t) := \frac{r_t(x_t) - \mu_t}{\sigma_t + \delta}, \qquad \mu_t = \mathbb{E}_{x_t\sim p_t}[r_t(x_t)], \quad \sigma_t^2 = \mathrm{Var}_{x_t\sim p_t}[r_t(x_t)]. \tag{61}$$

If the likelihood-score norm decreases while $z_t(x_t)$ increases, the sampler is likely improving a biased proxy off the diffusion marginal. In practice, this suggests attenuating or disabling LNP, for example by reducing γ, stopping local projection, or falling back to RS-only/RS+NSS.

Host guarantees. Our modules do not retrain the diffusion prior and leave the host's Tweedie plug-in likelihood estimator intact. 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). 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. Characterizing how norm-control wrappers interact with exact samplers (e.g., SMC/MCMC methods) remains an open problem.

Breadth of evaluation and remaining coverage. We report gains across multiple Tweedie plug-in solvers and tasks in the evaluated settings (Table 2) and probe key design choices via ablations (Tables 5–21). 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. The DDRM/DDNM study in Appendix G.4 provides additional empirical evidence using native observation corrections; broader validation across operators, noise models, and host families remains an important direction.

G.3 Preliminary operator and host checks

As preliminary evidence beyond the main natural-image restoration benchmark, we ran two small additional checks. First, on MRI reconstruction with an MRI subsampling operator and the public AmbientDiffusion $R = 4$ checkpoint, adding NCLG improves PSNR/SSIM from $36.45/0.95$ to $38.94/0.97$. Second, we tested PnP-DM (Wu et al., 2024) with an external DPS-style likelihood proxy on FFHQ inpainting, deblurring, and HDR. Because PnP-DM does not natively expose the

same Tweedie plug-in likelihood-score interface, these results are not used as main evidence for our method; they only suggest that likelihood-norm control may be worth studying beyond the hosts evaluated in the main paper.

Table 13: **Preliminary MRI reconstruction check** with an MRI subsampling operator.

Method	PSNR ↑	SSIM ↑
AmbientDiffusion	36.45	0.95
+NCLG	38.94	0.97

Table 14: **Preliminary PnP-DM check** on FFHQ. Each cell is PSNR/SSIM/LPIPS.

Method	Inpaint	Deblur	HDR
PnP-DM	18.43/0.723/0.314	20.34/0.785/0.289	22.05/0.753/0.264
NCLG (+RS+LNP)	16.84/0.688/0.293	18.85/0.713/0.250	20.24/0.737/0.241

G.4 Extension to DDRM and DDNM

DDRM (Kawar et al., 2022) and DDNM (Wang et al., 2022) condition denoised estimates on the observation through native correction steps. To adapt NCLG to these hosts, we use the magnitude of their observation correction as the diagnostic. For DDNM, the correction to the denoised estimate is

$$\delta_t^{\mathrm{DDNM}}(x_t) := A^\dagger\big(y - A\hat{x}_0(x_t, t)\big),$$

where $A^\dagger$ is the pseudoinverse. For DDRM, δ_t^{DDRM} is the noise- and singular-value-weighted observation correction in its SVD coordinates. In both cases, the diagnostic is $\|\delta_t(x_t)\|_2$. RS constructs restart states requiring smaller subsequent observation corrections, and NSS selects branches using cumulative squared correction norms. For the LNP update, define

$$E_t(x_t) := \tfrac{1}{2}\|\delta_t(x_t)\|_2^2, \qquad g_t^{\mathrm{corr}} := -\nabla_{x_t} E_t(x_t), \qquad \alpha_t := \frac{\gamma}{\|g_t^{\mathrm{corr}}\|_2 + \varepsilon}.$$

The normalized Langevin update uses $\hat{s}_t + \alpha_t g_t^{\mathrm{corr}}$ in place of the plug-in guided score.

Table 15 reports the resulting LPIPS on ImageNet256 center inpainting at $\sigma_y = 0.3$. NCLG reduces LPIPS from 0.418 to 0.347 for DDRM and from 0.405 to 0.352 for DDNM, corresponding to relative reductions of 17.0% and 13.1%. These results support the empirical extension to native observation corrections. The Tweedie plug-in certificate, including its linear operator extension, does not establish a guarantee for these native transitions.

Table 15: **Native-correction adaptation to DDRM/DDNM.** ImageNet256 center inpainting with $\sigma_y = 0.3$. LPIPS is lower for better reconstruction quality.

Host	Baseline LPIPS ↓	+NCLG LPIPS ↓	Relative reduction
DDRM	0.418	**0.347**	17.0%
DDNM	0.405	**0.352**	13.1%

H Additional Experimental Details

H.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[2] 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[3] 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).

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

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

Pretrained diffusion priors. For **FFHQ**, we use the pretrained diffusion prior released with DPS (Chung et al., 2022). 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 plug-in* likelihood guidance used by DPS and related 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

$$\log \widehat{p}_t(y \mid x_t) \ := \ -\frac{1}{2\sigma_y^2} \left\| y - \mathcal{A}\hat{x}_0(x_t, t) \right\|_2^2 + \text{const}, \tag{62}$$

and obtain the likelihood-score estimator by automatic differentiation,

$$\hat{g}_t(x_t; y, \mathcal{A}) \ := \ \nabla_{x_t} \log \widehat{p}_t(y \mid x_t) \ \approx \ \nabla_{x_t} \log p_t(y \mid x_t). \tag{63}$$

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

Additional degradation settings. For the additional degradation tasks in Tables 3, 8, and 9, we use the same measurement noise as in the table captions, $\sigma_y = 0.05$, and follow common degradation settings from prior diffusion inverse-problem benchmarks (Chung et al., 2022; Mardani et al., 2024; Zhang et al., 2025). Gaussian deblurring uses a 61×61 isotropic Gaussian blur kernel with standard deviation 3.0. Motion deblurring uses a 61×61 motion-blur kernel with kernel standard deviation 3.0 and intensity 0.2. JPEG restoration uses a differentiable JPEG forward operator with quality factor QF$=$ 2. HDR reconstruction uses the benchmark low-dynamic-range/HDR forward operator with severity factor 2; we apply the same Tweedie plug-in likelihood interface in Eq. (62) through this differentiable operator.

H.2 Baseline configuration protocol and tuning sanity check

We keep each baseline within its official public configuration family rather than using a single shared hand-tuned configuration file across all hosts. DCPS uses the public playground-style configuration from its released codebase, and the DCPS paper reports the same core parameters across tasks and datasets with 300 DDPM steps. MGPS uses the official dataset-dependent presets from its released sampler configurations. DAPS uses its official task/model-dependent Hydra configurations, combining the EDM-DAPS sampler configuration with the corresponding task configuration. Across our benchmark, the unavoidable changes are the benchmark-defining quantities: dataset/model selection, forward-operator geometry, and observation noise σ_y. When an exact official DAPS preset is unavailable, most notably for half-mask/outpainting, we keep the official scheduler and pixel-space MCMC hyper-parameters and change only the forward-operator geometry.

As a sanity check, Table 16 reports an additional DAPS-tuned setting on ImageNet half-mask/outpainting. The tuned DAPS setting uses $\mathrm{lr} = 7{\times}10^{-5}$, $\tau = 0.012$, and $\mathrm{lr_min_ratio} = 0.008$. NCLG remains beneficial when applied to this tuned host configuration.

Table 16: **DAPS tuning sanity check** on ImageNet half-mask/outpainting.

Method	PSNR ↑	SSIM ↑	LPIPS ↓
DAPS	18.13	0.650	0.305
DAPS tuned	18.42	0.670	0.300
DAPS tuned (+RS+LNP)	19.52	0.730	0.285

H.3 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 (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 (Heusel et al., 2017).

Correlation metrics. 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), let $n_i = \mathrm{Norm}_i(\mathcal{W})$ and $\ell_i = \mathrm{LPIPS}_i$. Pearson correlation measures linear association and is computed as the normalized covariance

$$\begin{aligned}\rho_{\mathrm{P}}(\mathcal{W}) &:= \mathrm{Corr}\big(\{n_i\}_{i=1}^R, \{\ell_i\}_{i=1}^R\big) \\ &= \frac{\sum_{i=1}^R (n_i - \bar{n})(\ell_i - \bar{\ell})}{\sqrt{\sum_{i=1}^R (n_i - \bar{n})^2}\sqrt{\sum_{i=1}^R (\ell_i - \bar{\ell})^2}},\end{aligned} \tag{64}$$

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

$$\begin{aligned}\rho_{\mathrm{S}}(\mathcal{W}) &:= \mathrm{Corr}\big(\mathrm{rank}(\{n_i\}),\ \mathrm{rank}(\{\ell_i\})\big) \\ &= \rho_{\mathrm{P}}\big(\mathrm{rank}(\{n_i\}),\ \mathrm{rank}(\{\ell_i\})\big).\end{aligned} \tag{65}$$

H.4 Synthetic Mixture-of-Gaussians Experiment

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

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

$$p(x) = \frac{1}{2}\,\mathcal{N}\big(x_1;\,-0.1, \sigma^2\big)\,\delta\big(x_2 + 1\big) + \frac{1}{2}\,\mathcal{N}\big(x_1;\,+0.1, \sigma^2\big)\,\delta\big(x_2 - 1\big), \tag{66}$$

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

Observation. The observation reveals the second coordinate,

$$y \;=\; (x_0)_2 \in \{-1, +1\}, \tag{67}$$

i.e., $\mathcal{A}$ selects the second coordinate (a 1D linear operator).[4]

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

$$\gamma \;=\; 1 \times 10^{-2}. \tag{68}$$

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 Tweedie plug-in 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.

H.5 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 $\tau_{\mathrm{inv}} = 0.7$ (truncate inverse encoding at $t_{\mathrm{inv}} = \lceil \tau_{\mathrm{inv}} T \rceil$);
- inverse stochasticity $\eta_{\mathrm{inv}} = 0$ (deterministic inverse encoding);
- late-start ratio $\tau_{\mathrm{start}} = 0.7$ (enable guidance, and optionally LNP, primarily in the late/low-noise regime).

[4] If additive observation noise is used in a particular run, we follow the same Gaussian likelihood as in (62).

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 step INVSTEP; inverse depth ratio τ_{inv}; inverse noise η_{inv}; late-start ratio τ_{start} (enable guidance/LNP on a late window)
InvStep is the inverse of the host prior step: it moves one step from lower to higher diffusion time, without using $(y, \mathcal{A})$ or likelihood guidance.

1: **Round 1 (host):** $x_0^{(1)} \leftarrow \mathcal{S}(y, \mathcal{A})$
2: $t_{\text{inv}} \leftarrow \lceil \tau_{\text{inv}} T \rceil, \quad z_0 \leftarrow x_0^{(1)}$
3: **for** $k = 1, \ldots, t_{\text{inv}}$ **do**
4: $\quad z_k \leftarrow \text{INVSTEP}(z_{k-1},\ k{-}1 \to k; \eta_{\text{inv}})$
5: **end for**
The partial inverse encoder in Alg. 1 is therefore $\mathcal{S}^{-1}(x_0^{(1)}; \eta_{\text{inv}}) := z_{t_{\text{inv}}}$.
6: **Round 2 (restart with late guidance):** $x_0^{(2)} \leftarrow \mathcal{S}\big(y, \mathcal{A};\ \text{init} = z_{t_{\text{inv}}},\ \text{enable guidance/LNP for } t/T \in [0, \tau_{\text{start}}]\big)$
7: **return** $x_0^{(2)}$

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

- projection gain $\gamma = 50$;
- stabilizer ε for the normalized likelihood drift;
- 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 Stochastic Selection (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 $\mathcal{W}$ 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).

H.6 ALGORITHMS

The detailed algorithms of the three modules (RS, LNP, NSS) are shown in Alg. 2, 3, and 4. Algorithm 2 expands the partial inverse encoder $\mathcal{S}^{-1}$, and Algorithm 4 expands the probe runner $\mathcal{S}_{\text{probe}}$ used in Algorithm 1.

H.7 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., 2022), **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. Likelihood-Norm Projection (LNP), and C. Likelihood-Norm Stochastic Selection (NSS).

H.8 IMAGENET COMPUTATION COST AND HOST SAMPLING BUDGETS

Table 18 complements the FFHQ computation table in the main text with ImageNet256 center inpainting at $\sigma_y = 0.3$. All entries in Table 18 use 200 reverse steps and are measured using multiple NVIDIA RTX 4090 GPUs.

Increasing the compute available to the same host. Table 19 increases the sampling budget of DPS, MGPS, and DCPS from 200 to 1,000 and 2,000 host steps on the same ImageNet256 center-inpainting task at $\sigma_y = 0.3$. We compare these operating points with RS+LNP and RS+LNP+NSS applied to a 200-step host, reporting the full runtime including the wrappers. These measurements compare quality across compute budgets; the runtimes are not exactly equal across configurations. Increasing the unwrapped host from 1,000 to 2,000 steps approximately doubles runtime but reduces

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

Table 17: **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	0.324	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	0.143	0.269	0.304	0.337	0.317	0.335	0.350	0.379
+A+B+C	0.108	0.153	**0.143**	0.197	**0.135**	0.257	**0.290**	0.328	**0.310**	0.328	0.341	0.358
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	0.108	0.138	0.146	0.172	0.152	**0.164**	0.304	0.340	0.328	0.332	0.352	0.374
+A+B+C	0.109	0.128	0.145	**0.168**	0.147	0.168	0.295	0.331	0.315	**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	0.165	0.306	**0.301**	0.316	0.331	0.347	0.360

LPIPS by only 0.008, 0.011, and 0.006 for DPS, MGPS, and DCPS, respectively. For each of these hosts, RS+LNP obtains lower LPIPS than both higher-step baselines while taking less time than the 1,000-step baseline. Adding NSS further reduces LPIPS, with the additional runtime shown in the table.

H.9 MATCHED-BUDGET DIAGNOSTIC COMPARISONS

Table 20 compares norm control with alternative uses of computation on FFHQ center inpainting at $\sigma_y = 0.3$, with DPS as the NCLG host. The lower-budget group includes the 12-candidate CAS (Hong et al., 2024) and adapted misalignment-map (Boutin et al., 2023) comparisons. The approximately 30-second group compares full NCLG with 24-candidate versions of these selectors, Best-of-8, full-inversion ODE REGUIDANCE (Karan et al., 2025), multiple random restarts, and data-consistency selection. Their measured runtimes range from 28.13 to 32.56 seconds, so this group compares approximately matched wall-clock budgets. RS+LNP attains LPIPS 0.138 in 13.48 seconds, and adding NSS attains 0.131 in 29.12 seconds, below the 0.173–0.197 range of the other methods in the approximately 30-second group. The separate PnP-DM (Wu et al., 2024) result

Algorithm 4 NSS: Likelihood-Norm Stochastic Selection 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: $\textsc{LikeEst}$; candidate count N, selection set $\mathcal{T}_{\text{sel}}$, scoring window $\mathcal{W}$, probe noise scale $\eta_{\text{sol}} \in [0, 1]$
The probe runner $\mathcal{S}_{\text{probe}}$ below scores each candidate by continuing the host reverse chain to lower timesteps and accumulating low-noise likelihood norms.

Initialize $x_T \sim \mathcal{N}(0, I)$
for $t = T, T-1, \dots, 1$ **do**
  if $t \in \mathcal{T}_{\text{sel}}$ **then**
    $\mathcal{W}_t \leftarrow \{\, s \in \mathcal{W} : s \le t \,\}, \quad s_{\min} \leftarrow \min \mathcal{W}_t$
    Sample candidate noises $\{\eta^{(n)}\}_{n=1}^{N} \sim \mathcal{N}(0, I)$
    for $n = 1, \dots, N$ **do**
      $X_t^{(n)} \leftarrow x_t, \quad S^{(n)} \leftarrow 0$
      $X_{t-1}^{(n)} \leftarrow \textsc{BaseStep}(X_t^{(n)},\ t \to t-1; \eta^{(n)})$
      for $s = t-1, t-2, \dots, s_{\min}$ **do**
        if $s \in \mathcal{W}_t$ **then**
          $\ell_s^{(n)} \leftarrow \textsc{LikeEst}(X_s^{(n)}, s;\ y, \mathcal{A})$
          $\hat{g}_s^{(n)} \leftarrow \nabla_{X_s^{(n)}} \log \ell_s^{(n)}, \quad S^{(n)} \leftarrow S^{(n)} + \|\hat{g}_s^{(n)}\|_2^2$
        end if
        if $s > s_{\min}$ **then**
          Sample probe noise $\tilde{\eta}_s^{(n)} \sim \mathcal{N}(0, \eta_{\text{sol}}^2 I)$
          $X_{s-1}^{(n)} \leftarrow \textsc{BaseStep}(X_s^{(n)},\ s \to s-1; \tilde{\eta}_s^{(n)})$
        end if
      end for
    end for
  Equivalently, the loop above is $\{S^{(n)}\}_{n=1}^{N} = \mathcal{S}_{\text{probe}}(x_t, t, \mathcal{W}_t, \eta_{\text{sol}}, \{\eta^{(n)}\})$.
    $n^\star \leftarrow \arg\min_n S^{(n)}, \quad \eta_t \leftarrow \eta^{(n^\star)}$
  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

Table 18: **Computation cost and LPIPS** on ImageNet *Center* inpainting with 200 reverse steps and $\sigma_y = 0.3$. Time is wall-clock per sample (seconds), Peak GPU is peak memory usage, LPIPS is perceptual error (lower is better). Configurations cumulatively add RS, LNP, and NSS.

Host	Variant	Time [s] ↓	Peak GPU [GB] ↓	LPIPS ↓
DPS	baseline	84.71	8.1	0.483
	RS	181.36	8.1	0.437
	RS+LNP	253.73	10.3	0.337
	RS+LNP+NSS	548.12	32.5	0.324
DCPS	baseline	239.77	8.1	0.401
	RS	460.02	8.1	0.363
	RS+LNP	569.76	10.3	0.337
	RS+LNP+NSS	908.76	32.5	0.328
MGPS	baseline	169.26	9.1	0.431
	RS	372.72	9.1	0.388
	RS+LNP	469.25	11.3	0.340
	RS+LNP+NSS	736.91	33.4	0.331

reports the best of 20 samples at 34.52 seconds per sample. Its LPIPS and timing therefore describe a different evaluation budget: generating 20 samples sequentially would take approximately 690.4 seconds before selection overhead, based on the reported per-sample time.

Table 19: **Increasing the host sampling budget on ImageNet256 center inpainting**, $\sigma_y = 0.3$. Time is wall-clock seconds per reconstruction. The NCLG configurations use 200 host steps plus the indicated modules; their additional computation is included in the reported time.

Host	Configuration	Host steps	Time [s] ↓	LPIPS ↓
DPS	Baseline	200	84.7	0.483
	Increased steps	1,000	394.2	0.425
	Increased steps	2,000	786.3	0.417
	RS+LNP	200	253.7	0.337
	RS+LNP+NSS	200	548.1	**0.324**
MGPS	Baseline	200	169.3	0.431
	Increased steps	1,000	953.5	0.403
	Increased steps	2,000	1882.4	0.392
	RS+LNP	200	469.3	0.340
	RS+LNP+NSS	200	736.9	**0.331**
DCPS	Baseline	200	239.8	0.401
	Increased steps	1,000	1184.3	0.374
	Increased steps	2,000	2346.1	0.368
	RS+LNP	200	569.8	0.337
	RS+LNP+NSS	200	908.8	**0.328**

Table 20: **Quality under comparable computation budgets** on FFHQ center inpainting at $\sigma_y = 0.3$. Time is in seconds. PnP-DM uses the separate per-sample and best-of-20 reporting convention specified below.

Method	Configuration	Time [s] ↓	LPIPS ↓
Lower-budget comparison			
NCLG	RS+LNP	13.48	**0.138**
CAS (Hong et al., 2024)	12 candidates	16.43	0.212
Adapted map (Boutin et al., 2023)	12 candidates	15.28	0.253
Approximately 30-second comparison			
NCLG	RS+LNP+NSS	29.12	**0.131**
CAS (Hong et al., 2024)	24 candidates	30.04	0.173
Adapted map (Boutin et al., 2023)	24 candidates	32.56	0.174
Best-of-8	8 noise candidates	31.25	0.197
REGUIDANCE	Full inversion + ODE, 300 steps	28.13	0.183
Multiple random restarts	45 restarts	31.62	0.177
Data-consistency selection	45 candidates	31.62	0.182
Separate sampling protocol			
PnP-DM (Wu et al., 2024)	100 outer steps × 100 Langevin steps; 20 samples	$34.52^{\dagger}$	$0.153^{\dagger}$

† PnP-DM time is per sample, whereas its LPIPS is the best of 20 samples; this row does not use the same total computation budget as the approximately 30-second group.

H.10 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 $\textsc{LikeEst}(x_t, t;\, y, \mathcal{A}) \approx p_t(y \mid x_t)$, LNP considers the tempered density

$$q_{t,\alpha}(x_t) \;\propto\; p_t(x_t)\,\textsc{LikeEst}(x_t, t;\, y, \mathcal{A})^{\alpha}, \tag{69}$$

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 $y = \mathcal{A}(x_0) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_y^2 I)$, the Tweedie plug-in proxy used in the main text has the form

$$\textsc{LikeEst}(x_t, t;\, y, \mathcal{A}) \;\propto\; \exp\Big(-\frac{\|y - \mathcal{A}(\hat{x}_0(x_t, t))\|_2^2}{2\sigma_y^2}\Big), \tag{70}$$

so the corresponding likelihood score is

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

where $J_t(x_t) := \partial \hat{x}_0(x_t, t)/\partial x_t$. Thus $\|\hat{g}_t(x_t)\|_2$ scales with the plug-in measurement residual $\|y - \mathcal{A}(\hat{x}_0(x_t, t))\|_2$ (up to the operator/Jacobian):

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

Raising LikeEst to a larger power α in (69) 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 (69) is

$$\nabla_{x_t} \log q_{t,\alpha}(x_t) = \nabla_{x_t} \log p_t(x_t) + \alpha\, \nabla_{x_t} \log \textsc{LikeEst}(x_t, t;\, y, \mathcal{A}). \tag{73}$$

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):

$$x_t \leftarrow x_t + \eta_{\mathrm{mcmc}} \big(\hat{s}_t(x_t) + \alpha\, \hat{g}_t(x_t)\big) + \sqrt{2\eta_{\mathrm{mcmc}}}\, \xi, \qquad \xi \sim \mathcal{N}(0, I). \tag{74}$$

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

$$\tilde{g}_t(x_t) := \alpha_t\, \hat{g}_t(x_t), \qquad \|\tilde{g}_t(x_t)\|_2 \le \gamma. \tag{75}$$

We couple this cap with $\eta_{\mathrm{mcmc}} := 1/\gamma$. This has two direct consequences: (i) the deterministic likelihood displacement per inner step is uniformly bounded, $\|\eta_{\mathrm{mcmc}} \tilde{g}_t(x_t)\|_2 \le 1$; and (ii) the injected noise has variance $2\eta_{\mathrm{mcmc}} = 2/\gamma$, which decreases as γ increases. In particular, increasing γ allows a larger effective likelihood temperature (via the larger admissible α_t in (75)) while simultaneously reducing both discretization aggressiveness (through η_{mcmc}) and stochasticity (through $\sqrt{2\eta_{\mathrm{mcmc}}}$), 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 (70), 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, $x_t \leftarrow x_t + \eta_{\mathrm{mcmc}} \tilde{g}_t(x_t)$. 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 plug-in quadratic data-fidelity $\|y - \mathcal{A}(\hat{x}_0(x_t, t))\|_2^2$), which directly reduces the residual and therefore reduces $\|\hat{g}_t(x_t)\|_2$ by (72). Moreover, in the direct linear special case $\hat{x}_0(x_t, t) = x_t$, $\hat{g}_t(x_t) = -(1/\sigma_y^2)\mathcal{A}^\top(\mathcal{A}x_t - y)$ and one can verify

$$\nabla_{x_t} \frac{1}{2} \|\hat{g}_t(x_t)\|_2^2 = -\frac{1}{\sigma_y^2}\, (\mathcal{A}^\top \mathcal{A})\, \hat{g}_t(x_t), \tag{76}$$

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 $\eta_{\mathrm{mcmc}} = 1/\gamma$) 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.

H.11 Ablating NSS design choices: low-noise scoring and deterministic rollouts are critical

Table 21 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.

Table 21: **Ablation of NSS scoring window, solver noise, and candidate count on FFHQ with DPS host**. We report LPIPS (lower is better). Module C = NSS. Baseline NSS scores at the single timestep $\mathcal{W} = \{0.3\}$; subsequent rows vary the indicated parameter.

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 ($\mathcal{W} = \{0.3\}$)	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

(i) Scoring in the low-noise regime is most effective. Scoring windows restricted to the low-noise regime consistently yield the best LPIPS improvements, supporting 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}) degrade LPIPS. We conjecture 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 selection mechanism operating on noisy proposals.

H.12 NSS SEARCH AND LIGHTWEIGHT VARIANTS

NSS is a sparse in-trajectory selector rather than a full best-of-N resampling procedure. Table 22 compares NSS with a naive Best-of-5 baseline on FFHQ center inpainting at $\sigma_y = 0.3$. NSS alone is comparable to Best-of-5 while being cheaper, and RS+NSS substantially improves over both by combining a better restart basin with norm-based branch selection.

Table 22: **NSS versus Best-of-5** on FFHQ center inpainting at $\sigma_y = 0.3$.

Method	Time [s] ↓	LPIPS ↓	Candidates
DPS	5.25	0.316	1
NSS	19.47	0.268	5
RS+NSS	27.95	0.173	5
Best-of-5	24.52	0.265	5

Table 23 reports lighter NSS variants on FFHQ SR 4× at $\sigma_y = 0.3$. The results indicate a quality-compute trade-off: reducing the candidate count or stopping probe rollouts earlier lowers cost while preserving improvement over the DPS host. Here, the baseline NSS setting uses a single initial selection time $\mathcal{T}_{\text{sel}} = \{T\}$, deterministic probe rollouts $\eta_{\text{sol}} = 0$, $N = 5$ candidates, and the single scoring timestep $\mathcal{W} = \{0.3\}$. The single-factor rows change only one of N or $\mathcal{W}$. The "Lightweight NSS" row keeps $\mathcal{T}_{\text{sel}} = \{T\}$ and $\eta_{\text{sol}} = 0$, but jointly uses $N = 2$ and the earlier-stopping probe window $\mathcal{W} = [0.5, 0.7]$; in the corresponding runtime measurement, this reduces NSS time from about 15s/sample to about 5s/sample while retaining an LPIPS improvement over the DPS host.

Table 23: **Lightweight NSS variants** on FFHQ SR 4× with DPS host at $\sigma_y = 0.3$.

Configuration	Candidates	Window	LPIPS ↓
DPS baseline	–	–	0.282
Baseline NSS	5	$\{0.3\}$	0.255
$N = 2$	2	$\{0.3\}$	0.280
Middle narrow window	5	$[0.5, 0.7]$	0.261
Lightweight NSS	2	$[0.5, 0.7]$	0.257

H.13 COMPARISON TO REGUIDANCE

We also compare to the REGUIDANCE wrapper (Karan et al., 2025), which corresponds to the ablation variant "+ A (τ_{inv}=1.0 ODE)" in Table 5. 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 constraints emphasized by REGUIDANCE. **(i) Deterministic requirement.** REGUIDANCE requires deterministic second-pass sampling and reports that its benefit can diminish when combined with stochastic sampling (Karan et al., 2025). This restriction can hurt distribution-level metrics such as FID, whereas our RS remains compatible with stochastic second-pass sampling. **(ii) Full reversion to high noise.** Full reversion traverses the highest-noise regime and may break the inverse-consistency condition in Eq. (4.2), which RS exploits. Partial inversion avoids this regime by design, and therefore yields smaller low-t likelihood norms and better downstream refinement.

H.14 RS DETAILS: RESIDUAL TRANSFER AND LIKELIHOOD-NORM CONTROL

This appendix makes explicit why RS tends to reduce the likelihood-score norm for common Tweedie plug-in solvers. 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

$$\hat{g}_t(x_t) := \nabla_{x_t} \ell_t(x_t) \;\approx\; \frac{1}{\sigma_y^2}\, J_t(x_t)^\top \mathcal{A}^\top \Big(y - \mathcal{A}\hat{x}_0(x_t, t)\Big), \tag{77}$$

and therefore the norm bound

$$\|\hat{g}_t(x_t)\|_2 \;\le\; \frac{\|J_t(x_t)\|_{\text{op}}\, \|\mathcal{A}\|_{\text{op}}}{\sigma_y^2}\, \big\|y - \mathcal{A}\hat{x}_0(x_t, t)\big\|_2. \tag{78}$$

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

$$\big\|y - \mathcal{A}\hat{x}_0(z_{t_{\text{inv}}}, t_{\text{inv}})\big\|_2 \le \big\|y - \mathcal{A}x_0^{(1)}\big\|_2 + \|\mathcal{A}\|_{\text{op}}\, \big\|\hat{x}_0(z_{t_{\text{inv}}}, t_{\text{inv}}) - x_0^{(1)}\big\|_2, \tag{79}$$

which is small when the first pass is measurement-consistent and the truncated inversion is accurate. Combining (79) with (78) explains why RS produces a low-norm restart and why full inversion (high-noise) can break this mechanism by degrading inverse-consistency.