GUIDING DIFFUSION SAMPLERS WITH BLACK-BOX OPTIMIZATION SOLVERS

Anonymous authors
Paper under double-blind review

ABSTRACT

From robot motion planning to scientific inverse problems, generative tasks often require optimizing target objectives while adhering to a complex data-driven prior. Diffusion models excel as such priors, but guiding them to optimize targets or obey constraints remains challenging, even with powerful optimization solvers available. When integrated into diffusion sampling, these solvers often distort the sampling trajectory with their hard projections. We propose Proximal Diffusion Guidance (ProxDG), a training-free method that cleanly integrates solvers as soft guidance during diffusion sampling. Specifically, ProxDG guides samples using the gradient of a time-varying Moreau envelope. This envelope smooths the original objective without shifting its optima, providing well-behaved guidance even under hard constraints. Importantly, computing this guidance reduces to a proximal subproblem that admits powerful, problem-specific solvers. With this proximal scheme, ProxDG samples along a continuous probability path to the target posterior, supporting both stochastic and deterministic sampling with exact constraint enforcement at the final step. We demonstrate in both synthetic and real-world experiments that ProxDG effectively leverages off-the-shelf solvers to significantly outperform training-free baselines.

1 INTRODUCTION

Across domains such as seismology (Virieux & Operto, 2010), optical tomography (Choi et al., 2007), computational imaging (Lustig et al., 2007; Chael et al., 2019; Ongie et al., 2020; Bertero et al., 2021), motion planning (Janner et al., 2022; Carvalho et al., 2023), and robotic manipulation (Chi et al., 2023), scientific and engineering problems are often under-specified: their objectives admit not just one optimal solution, but many. These problems call for sampling rather than mode-seeking optimization, so that we recover a diverse and representative set of solutions (Corso et al., 2023; Sun & Bouman, 2021; Florence et al., 2021). In practice, optimizing for target properties trades off with producing "natural" solutions. In motion planning, for example, trajectories must both avoid collisions and resemble expert behavior (Carvalho et al., 2023). In scientific inverse problems, reconstructed signals must fit noisy measurements while resembling natural data (Zheng et al., 2025).

"Naturalness" is difficult to capture in a hand-crafted formula, so we use a learned, data-driven prior instead. Pre-trained diffusion models are a natural choice for such a prior, since they excel at modeling the distribution of natural data and sampling from it (Pidstrigach, 2022; Stanczuk et al., 2024; Ho et al., 2020). Their iterative denoising process also provides a way to incorporate task objectives: guidance methods intervene at each step to steer the sampling trajectory toward desired solutions. We focus on training-free guidance (TFG), which allows us to adapt a pre-trained prior to new objectives and constraints at test time (Chung et al., 2023; Yu et al., 2023; Ye et al., 2024).

Yet, combining sampling with the optimization of task objectives remains challenging, especially under hard constraints. A standard TFG approach uses task gradients to directly update either the noisy sample or the model's estimate of the clean sample. While simple and efficient, these generic gradient updates largely neglect problem-specific structure. In constrained or ill-conditioned landscapes, they can stall or produce unstable updates, leaving samples far from task optima or causing sampling to diverge (Figure 1). Other TFG methods integrate powerful solvers directly into the sampling process (Zhu et al., 2023; Christopher et al., 2024; Zhang et al., 2025b). However, these methods often project samples off the noisy data manifold, which degrades sample quality

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

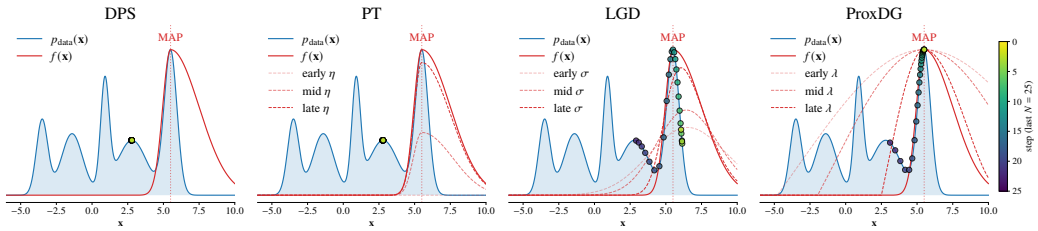

Figure 1: $\hat{\mathbf{x}}_0$ iterates under DPS (Chung et al., 2023), PT (Kim et al., 2025b), LGD (Song et al., 2023b), and ProxDG on a multimodal prior with objective maximized at MAP. Red lines show the smoothed objective at three scales. DPS and PT stall; LGD overshoots; ProxDG converges to MAP.

(Chung et al., 2022; Utkarsh et al., 2025; Yao et al., 2025). Their solver updates also disrupt the continuous-time dynamics of diffusion sampling, which can destabilize sampling (Section 4.3).

To address these challenges, we propose **Prox**imal **D**iffusion **G**uidance (**ProxDG**). ProxDG is a gradient guidance method that leverages black-box solvers while retaining the continuous-time dynamics of diffusion sampling. It is entirely training-free, supports diverse solvers, and handles both soft objectives and hard constraints within a unified framework. In proposing ProxDG, we make three key contributions:

- **Gradient guidance with black-box solvers.** ProxDG enables gradient guidance methods such as DPS (Chung et al., 2023) to employ black-box solvers for task-specific optimization. Our key insight is that the solver's output is exactly a gradient of the task objective's Moreau envelope. This insight lets us combine powerful solvers with advances in gradient guidance that do not readily transfer to prior solver-based methods.
- **On-manifold guidance with proximal solvers.** To compute guidance, ProxDG derives a gradient from the solver's output and backpropagates this gradient through the diffusion model. This operation approximately projects the guidance onto a tangent space of the data manifold, steering samples along the manifold. Other solver-based methods omit this projection, so they risk pushing samples off-manifold.
- **Novel probability path for constrained posterior sampling.** We construct a continuous probability path to the target posterior by annealing the objective's Moreau envelope through time. Our novel construction supports both stochastic and deterministic sampling, even under hard constraints (where the gradient of the original objective is undefined). A final projection guarantees exact feasibility without sacrificing distributional accuracy.

2 Related Work

Gradient guidance typically computes a loss gradient at the clean estimate, then backpropagates this gradient through the denoiser. MCG (Chung et al., 2022) introduces this approach for linear inverse problems and shows that denoiser backpropagation approximately projects the gradient onto a tangent space of the data manifold. DPS (Chung et al., 2023) extends the approach to arbitrary differentiable objectives. LGD (Song et al., 2023b) uses a Gaussian surrogate around the clean estimate, while FreeDoM (Yu et al., 2023) and Universal Guidance (Bansal et al., 2024) re-noise to bring samples back toward the data manifold. TFG (Ye et al., 2024) unifies these methods under one framework.

TFG's gradient formulation has enabled many advances in sample quality and sampling efficiency. DSG (Yang et al., 2024), GGDM (Guo et al., 2024), and Huang et al. (2024) rescale or project the gradient to keep iterates on the noisy data manifold. MPGD (He et al., 2024), SDO (Dou et al., 2026), and Graikos et al. (2025) reduce the per-step cost of denoiser backpropagation without compromising sample quality. DPM-Solver++ (Lu et al., 2025) uses multistep integration for fast, stable sampling under strong guidance. TDS (Wu et al., 2023) and DAS (Kim et al., 2025b) use gradient guidance to construct informed proposals for sequential Monte Carlo sampling. Beyond diffusion models, Feng et al. (2025) and FlowDPS (Kim et al., 2025a) extend gradient guidance to flow matching, while Sabour et al. (2025) extend it to test-time scaling with flow maps. Together, these advances contribute

to TFG's empirical success (Zhang et al., 2025c; Zheng et al., 2025), but do not readily transfer to methods that replace samples or clean estimates with solver outputs.

Gradient guidance can fail when the objective is stiff or discontinuous. Thus, an alternative line of work replaces gradient updates with powerful solvers, solving a proximal subproblem at each step. These *proximal guidance* methods fall into two classes. The first class substitutes the solver output in for the noisy sample. For example, some methods fully or partially project the noisy sample onto a constraint set (Christopher et al., 2024; Sharma et al., 2024; Liang et al., 2025; Luan et al., 2026; Römer et al., 2025; Zhang et al., 2025b), and others enforce safety with control-barrier functions (Xiao et al., 2025; Gadginmath et al., 2026). The second class instead solves a full or relaxed proximal subproblem anchored at the clean estimate (Zhu et al., 2023; Narasimhan et al., 2025; Cheng et al., 2025; Utkarsh et al., 2025; Yuan et al., 2023; Zampini et al., 2025; Rout et al., 2025; Christopher et al., 2026); this class can often reuse solvers designed for clean data. By offloading the optimization to a task-specific solver, both classes routinely outperform gradient guidance on constraint satisfaction and objective value. But substituting in the solver output can pull samples off the data manifold. Furthermore, these substitutions are not necessarily gradients, so none of the TFG advances transfer easily.

3 Background

The Diffusion Process. Diffusion models generate data by reversing a noise-corruption process. In continuous-time, this forward process gradually corrupts clean data $\mathbf{x}_0 \sim p_{\mathrm{data}}$ into pure noise over time $t \in [0, T]$. We formulate this process with the following stochastic differential equation (SDE):

$$d\mathbf{x}_t = \mu(t)\mathbf{x}_t\, dt + g(t)\, d\mathbf{w}_t \tag{1}$$

where $\mu, g : [0, T] \to \mathbb{R}$ are the drift and diffusion coefficients and $\mathbf{w}_t \in \mathbb{R}^d$ denotes the standard d-dimensional Wiener process. Since Eq. 1 is linear in $\mathbf{x}_t$ with additive noise, its transition kernel is Gaussian, yielding the forward noising step:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{2}$$

where α_t, σ_t are determined by μ, g (Appendix B.2).

To generate data, we sample noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and simulate a reverse-time process that transports it back to p_{data}. By Anderson's theorem (Song et al., 2021b; Anderson, 1982), this process is the time-reversal of the forward SDE in Eq. 1:

$$d\mathbf{x}_t = [\mu(t)\mathbf{x}_t - g(t)^2 \nabla \log q_t(\mathbf{x}_t)]\, dt + g(t)\, d\overline{\mathbf{w}}_t, \qquad \mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{3}$$

where q_t denotes the time-marginal of Eq. 1 at time t and $\overline{\mathbf{w}}_t$ is a reverse-time Wiener process. Since the score $\nabla \log q_t$ is unavailable in closed-form, we approximate it with a trained neural network $s_\theta(\mathbf{x}_t, t)$. Substituting s_θ into Eq. 3 gives our learned generative process.

Guided Diffusion Sampling. In most applications, we want our samples to optimize a task objective while respecting the data distribution. Let $f(\mathbf{x}_0)$ denote this task objective. Our goal is then to sample from the target posterior

$$p_{\mathrm{tar}}(\mathbf{x}_0|\mathbf{y}) \propto q_0(\mathbf{x}_0) \exp(-f(\mathbf{x}_0)) \tag{4}$$

which concentrates mass on points that are both high-likelihood under f and high-probability under the diffusion prior q_0. To sample from p_{tar}, we construct a probability path $(p_t)_{t\in[0,T]}$ that interpolates between endpoints $p_T = \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $p_0 = p_{\mathrm{tar}}$ by tilting each marginal q_t with a time-dependent guidance function f_t:

$$p_t(\mathbf{x}_t|\mathbf{y}) \propto q_t(\mathbf{x}_t) \exp(-f_t(\mathbf{x}_t)) \tag{5}$$

with boundary conditions $f_T \equiv 0$ and $f_0 = f$. The ideal f_t matches each p_t to the true Bayesian posterior, but its closed form $f_t^*(\mathbf{x}_t) = -\log \mathbb{E}_{q_0(\mathbf{x}_0|\mathbf{x}_t)}[\exp(-f(\mathbf{x}_0))]$ is intractable without training (Denker et al., 2024). Diffusion Posterior Sampling (DPS) sidesteps the intractable expectation by approximating $q_0(\mathbf{x}_0|\mathbf{x}_t)$ with a Dirac mass at its mean, $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$. This reduces the guidance function to $f_t^{\mathrm{DPS}}(\mathbf{x}_t) = f(\hat{\mathbf{x}}_0(\mathbf{x}_t, t))$, where the denoiser $\hat{\mathbf{x}}_0(\mathbf{x}_t, t)$ is Tweedie's estimator of this mean. As shorthand, we write $\hat{\mathbf{x}}_0(\mathbf{x}_t, t)$ when the context is clear. By Tweedie's formula, the denoiser is determined by the score of q_t, which we approximate using s_θ:

$$\mathbb{E}_{q_0(\mathbf{x}_0|\mathbf{x}_t)}[\mathbf{x}_0|\mathbf{x}_t] = \frac{\mathbf{x}_t + \sigma_t^2 \nabla \log q_t(\mathbf{x}_t)}{\alpha_t}, \qquad \hat{\mathbf{x}}_0 \coloneqq \frac{\mathbf{x}_t + \sigma_t^2 s_\theta(\mathbf{x}_t, t)}{\alpha_t} \tag{6}$$

To sample along $(p_t)_t$, we simulate the reverse-time SDE in Eq. 3 with the guided score $s_\theta(\mathbf{x}_t, t) - \nabla f_t^{\mathrm{DPS}}(\mathbf{x}_t)$. By the chain rule, $\nabla f_t^{\mathrm{DPS}}(\mathbf{x}_t) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \nabla f(\hat{\mathbf{x}}_0)$, where $\mathbf{J}_{\hat{\mathbf{x}}_0}$ is the denoiser Jacobian. In practice, we compute $\nabla f_t^{\mathrm{DPS}}(\mathbf{x}_t)$ by backpropagating $\nabla f(\hat{\mathbf{x}}_0)$ through the denoiser.

4 Proximal Diffusion Guidance

Gradient guidance requires a well-defined and informative task gradient ∇f. This requirement can fail in two important settings, even when a solver provides an effective update.

Setting 1: Constrained Sampling. To restrict samples to a constraint set $\mathcal{C}$, we employ the objective

$$f(\mathbf{x}) \coloneqq \chi_{\mathcal{C}}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in \mathcal{C} \\ \infty & \text{otherwise} \end{cases} \tag{7}$$

Outside $\mathcal{C}$, ∇f is undefined, so gradient guidance cannot use this objective to steer samples toward feasibility. A projection solver can nevertheless find a nearby feasible point.

Setting 2: Ill-Conditioned Objectives. Gradient updates become erratic in high-curvature regions and vanish in flat regions, leaving samples far from the optimum (Figure 1). A task-specific solver can use information beyond the gradient, such as curvature, to navigate these regions.

We address these limitations by constructing a *smoothed* guidance function f_t^λ, whose gradients are derived from solver outputs. These gradients steer toward feasibility under hard constraints and use task-specific solvers to navigate ill-conditioned objectives.

4.1 Smoothing Guidance with Moreau Envelopes

We replace the DPS guidance function with a proxy that smooths f without shifting its minimizers. Specifically, we employ the Moreau envelope (Moreau, 1965), a classical smoothing operator from variational analysis. For any $\lambda > 0$, the envelope of f at a point $\mathbf{u}$ is $\Psi_f(\mathbf{u}, \lambda) \coloneqq \inf_{\mathbf{z} \in \mathbb{R}^d} \{f(\mathbf{z}) + \frac{1}{2\lambda}\|\mathbf{u} - \mathbf{z}\|^2\}$. Intuitively, $\Psi_f(\mathbf{u}, \lambda)$ is the minimum value of f attainable in a λ-sized neighborhood of $\mathbf{u}$. We provide more geometric intuition for the Moreau envelope in Appendix B.3. Even for discontinuous or ∞-valued f, the envelope enjoys three properties that make it an attractive proxy for optimization (Parikh & Boyd, 2014; Beck & Teboulle, 2009; Lee et al., 2014; Becker et al., 2019) and sampling (Pereyra, 2016; Durmus et al., 2018a; Habring et al., 2025) alike:

1. Ψ_f is finite and jointly continuous in $(\mathbf{u}, \lambda)$.
2. Ψ_f converges to f pointwise as $\lambda \to 0$.
3. Where f and Ψ_f are both differentiable, every critical point of Ψ_f is a critical point of f.

Together, these properties are uncommon among smoothing operators. Gaussian convolution, for example, smooths f but can shift its minimizers (violating Property 3), causing its gradients to steer toward the wrong points. Power transformations preserve the minimizers but cannot smooth a discontinuous or ∞-valued f (violating Property 1). The envelope makes neither tradeoff. We adapt it for diffusion guidance by evaluating it at Tweedie's estimate of the clean sample, and annealing the regularization parameter λ with diffusion time:

$$f_t^\lambda(\mathbf{x}_t) \coloneqq \Psi_f\left(\hat{\mathbf{x}}_0(\mathbf{x}_t, t), \lambda_t\right) = \inf_{\mathbf{z}_t \in \mathbb{R}^d} \left\{ f(\mathbf{z}_t) + \tfrac{1}{2\lambda_t}\|\hat{\mathbf{x}}_0(\mathbf{x}_t, t) - \mathbf{z}_t\|^2 \right\} \tag{8}$$

for $t \in (0, T]$, with $\lambda_t \to 0$ as $t \to 0$. Appendix C shows that f_t^λ inherits properties 1–3 from the envelope, so it anneals smoothly to the task objective. Figure 1 depicts this on a multimodal landscape. f_t^λ guides samples to the MAP, while LGD's (Song et al., 2023b) Gaussian-smoothed guidance overshoots a peak that drifts with the smoothing scale, and DAS's (Kim et al., 2025b) power-transformed guidance stalls in flat regions.

Despite these failure modes, DPS, LGD, and DAS remain widely used and empirically successful, owing to a simple gradient formula that supports diverse objectives. We show that f_t^λ uses the same formula with a different loss. A task-specific solver first computes a proximal point $\mathbf{z}_t^*$ of f at $\hat{\mathbf{x}}_0$. We then hold $\mathbf{z}_t^*$ fixed and backpropagate the squared error $\|\hat{\mathbf{x}}_0 - \mathbf{z}_t^*\|^2$ through the denoiser. The resulting gradient is exactly ∇f_t^λ (Proposition 4.1), so standard tools of gradient guidance still apply.

Proposition 4.1. *(Informal) Fix $t \in (0, T]$. Given $\mathbf{x}_t \in \mathbb{R}^d$, let $\mathbf{z}_t^*$ be any proximal point of f at the clean estimate $\hat{\mathbf{x}}_0$,*

$$\mathbf{z}_t^* \in \arg\min_{\mathbf{z}\in\mathbb{R}^d} \left\{ f(\mathbf{z}) + \tfrac{1}{2\lambda_t} \|\hat{\mathbf{x}}_0(\mathbf{x}_t, t) - \mathbf{z}\|^2 \right\}. \tag{9}$$

Apply a stop-gradient to $\mathbf{z}_t^*$. Then for almost every $\mathbf{x}_t$, ∇f_t^λ is the gradient of a squared error,

$$\nabla f_t^\lambda(\mathbf{x}_t) = \nabla_{\mathbf{x}_t} \tfrac{1}{2\lambda_t} \left\| \hat{\mathbf{x}}_0(\mathbf{x}_t, t) - \mathrm{stopgrad}(\mathbf{z}_t^*) \right\|^2. \tag{10}$$

Thus, computing ∇f_t^λ takes one proximal solve and one backward pass through the denoiser, and never differentiates through the solver.

We call guidance with the gradient ∇f_t^λ **Prox**imal **D**iffusion **G**uidance (**ProxDG**).

4.2 Computing the Guidance Update

ProxDG replaces a single line of DPS, which is the standard template for gradient guidance (Figure 2). Instead of backpropagating the task loss $f(\hat{\mathbf{x}}_0)$, ProxDG backpropagates the squared error in Eq. 10. We discuss three design aspects of ProxDG below: the annealing schedule for λ_t, the proximal solver, and the denoiser backpropagation.

```
x0 = denoiser(xt, t)
# loss = f(x0)
z = stopgrad(prox(f, x0, lam_t))
loss = ||x0 - z||^2 / (2*lam_t)
g = backprop(loss, xt)
xt = guided_step(xt, t, g)
```

Figure 2: ProxDG update step. ProxDG replaces the DPS loss (red, commented) with a squared error to the proximal point (blue).

Annealing schedule. In Appendix B.4.2, we show that the optimal annealing corresponds to the posterior covariance $\mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t]$, which is expensive to approximate (Peng et al., 2024; Rissanen et al., 2025). We instead set $\lambda_t \coloneqq \beta\sigma_t^2$ for some $\beta > 0$. This schedule keeps the guidance scale constant in t, so the guidance neither diverges nor vanishes as $t \to 0$ (Appendix B.4.1). DiffPIR (Zhu et al., 2023) uses this schedule successfully.

Proximal solver. We can employ any ℓ_2-regularized solver to compute the proximal subproblem in Eq. 9, so ProxDG applies to a wide range of objectives. We make a few minor adjustments to the subproblem for numerical stability, details in Appendix B.4.1. We also show that varying the solver recovers a large class of prior guidance methods. For example, a single gradient step at $\hat{\mathbf{x}}_0$ reduces ProxDG to DPS, and a weighted average of Gaussian samples around $\hat{\mathbf{x}}_0$ reproduces TFG (Song et al., 2023b; Ye et al., 2024; Zhang et al., 2025c) in expectation. DPS and TFG are weak solvers, yet their empirical success suggests that even rough proximal solves produce useful guidance (Zhang et al., 2025c; Zheng et al., 2025).

ProxDG also retains the guarantees of Section 4.3 under rough solves (Appendix H). A solver that converges to the wrong stationary point still yields the exact gradient of a valid guidance function (Corollary H.9), and early stopping adds a path-space KL error that shrinks quadratically near a stationary point (Theorem H.12).

Denoiser backpropagation. Backpropagating guidance through the denoiser does more than apply the chain rule. The backward pass approximately projects the guidance onto a tangent space of the data manifold. MCG (Chung et al., 2022) derives this result for locally linear manifolds, and Kharitenko et al. (2026) extend it to compact C^3 submanifolds in the small-noise limit (Appendix B.3.2). Ventura et al. (2025) confirm that, in trained diffusion models, backpropagation suppresses directions normal to the manifold. ProxDG therefore steers samples along the manifold. Prior proximal guidance methods skip denoiser backpropagation, substituting the proximal point directly for the clean estimate or noisy sample. Such substitutions can push samples into low-density regions of the diffusion prior, degrading sample quality (Figure 3).

4.3 Sampling in Continuous Time

Equipped with ∇f_t^λ, we aim to sample along the probability path $(p_t^\lambda)_t$:

$$p_t^\lambda(\mathbf{x}_t|\mathbf{y}) \propto q_t(\mathbf{x}_t) \exp(-f_t^\lambda(\mathbf{x}_t)) \tag{11}$$

This path is justified because it is continuous and terminates at the target posterior (Appendix D).

Theorem 4.2. *The probability path $(p_t^\lambda)_t$ is continuous on $[0, T]$ in W_p, for every $p \in [1, \infty)$.*

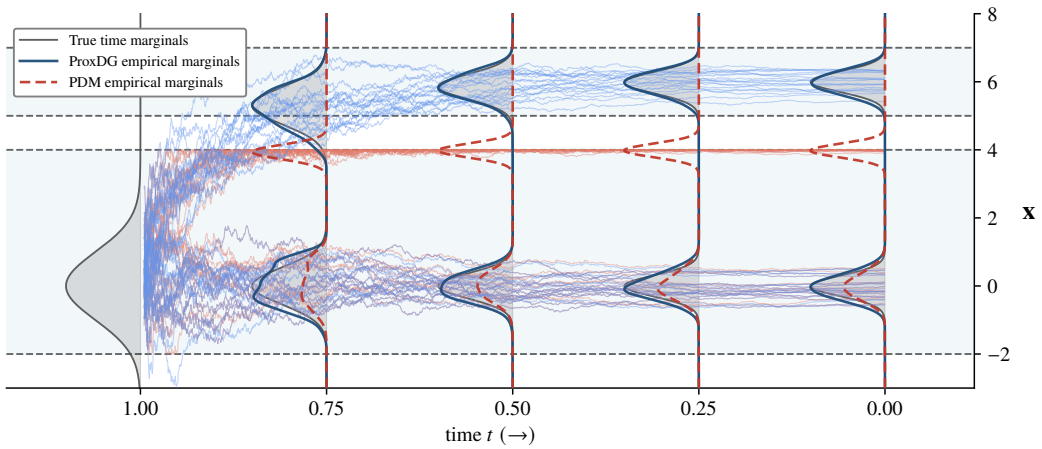

Figure 3: Constrained sampling from a nonconvex constraint set (blue shading). Sampling trajectories and empirical time marginals are shown for PDM (red) and ProxDG (blue), with the ideal time marginals as gray fill. ProxDG softly anneals to the target posterior, whereas PDM hard-projects samples into low-density boundary regions.

To sample toward $(p_t^\lambda)_t$, ProxDG uses standard gradient guidance. For a tilted path $p_t \propto q_t \exp(-f_t)$, gradient guidance plugs the score of p_t into the reverse-time SDE (Eq. 3), yielding

$$d\mathbf{x}_t = [\mu(t)\mathbf{x}_t - g(t)^2(\nabla \log q_t(\mathbf{x}_t) - w\nabla f_t(\mathbf{x}_t))]dt + g(t)d\overline{\mathbf{w}}_t, \qquad \mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{12}$$

Here, $w \geq 0$ is the guidance scale, DPS sets $f_t = f_t^{\mathrm{DPS}}$, and ProxDG sets $f_t = f_t^\lambda$ (Algorithm 1). Even with exact scores, integrating Eq. 12 is biased unless $f_t = f_t^*$ (Chidambaram et al., 2024; Du et al., 2023; Lai et al., 2025); ProxDG can remove this bias asymptotically with MCMC or sequential Monte Carlo corrections (Du et al., 2023; Wu et al., 2023; Kim et al., 2025b). Either way, the sampler integrates Eq. 12 or its probability-flow ODE, whose trajectories neither blow up nor split.

Trajectories stay well-defined only if the guidance is well-behaved. DPS differentiates f directly, so kinks, jumps, and large gradients in f destabilize sampling (Delgadino et al., 2026; Lu et al., 2025), while hard constraints make the guidance undefined. Proximal guidance methods avoid differentiating f, but each projection makes the trajectory jump (Kurtz & Davydov, 2026); to recover guarantees, they must re-noise and restrict to convex or prox-regular constraint sets (Christopher et al., 2024; Zampini et al., 2025). ProxDG is well-behaved without re-noising, for discontinuous f and nonconvex constraints.

Theorem 4.3 (Informal, see Theorems E.6 and F.3)**.** *The ProxDG SDE and ODE are well-posed away from $t = 0$, with the following guarantees.*

- ***SDE.*** *For almost every noise draw, the sampling trajectory exists, is unique, and never splits.*
- ***ODE.*** *For almost every initial noise, the sampling trajectory exists, is unique, and never splits. No two trajectories ever merge, so sampling is invertible.*

The ODE guarantee relies on the one-sided smoothness (semiconcavity) of the Moreau envelope (Lions & Seeger, 2024). This guarantee extends ProxDG to flow matching priors and deterministic solvers (Kerrigan et al., 2024; Song et al., 2021a; Lu et al., 2025), even under hard constraints. Without re-noising, proximal guidance methods have no such guarantee, and repeated projection traps samples at the constraint boundary (Rochman-Sharabi & Louppe, 2026; Yang et al., 2026).

Explicit integration. ProxDG integrates the SDE or ODE explicitly, with one proximal solve per timestep (Algorithms 2 and 3).

Proximal correction. Constrained sampling requires $\mathbf{x}_0 \in \mathcal{C}$, but explicit integration cannot guarantee feasibility. ProxDG therefore applies a proximal correction only at the end, $\mathbf{x}_0 = \mathrm{prox}_{\lambda_{\Delta t} f}(\tilde{\mathbf{x}}_0)$, leaving the sampling trajectory continuous beforehand. For hard constraints $f = \chi_{\mathcal{C}}$, the correction reduces to projection onto $\mathcal{C}$ and never increases the sampling error.

Table 1: Zero-shot constrained generation of PDE solutions. Values $\times 10^{-2}$ except FPD. Best in bold, second-best underlined, and "div." marks diverged runs.

Dataset	Method	CE (IC/BC)	CE (CL)	MMSE	SMSE	FPD
Navier–Stokes	PCFM	**0**	**0**	17.8	12.7	2.47
	PDM	**0**	**0**	12.2	6.01	1.96
	FunDPS	<u>579</u>	9.32	<u>2.98</u>	<u>1.57</u>	1.91
	D-Flow	1020	<u>3.02</u>	5.94	4.06	<u>0.823</u>
	ProxDG (Ours)	**0**	**0**	**2.02**	**0.821**	**0.451**
Burgers (IC)	PCFM	**0**	**0**	3.74	2.67	20.2
	PDM	div.	div.	div.	div.	div.
	FunDPS	<u>35.3</u>	<u>14.1</u>	<u>0.397</u>	<u>0.651</u>	**0.583**
	D-Flow	68.2	16.9	0.693	1.75	3.00
	ProxDG (Ours)	**0**	**0**	**0.180**	**0.268**	<u>0.897</u>

Theorem 4.4 (Informal, see Proposition G.3 and Theorem G.5)**.** *For $f = \chi_{\mathcal{C}}$, the projected samples converge to the target posterior up to integration error, even under equality constraints.*

5 Experiments

We evaluate ProxDG on three tasks that span hard constraints and soft objectives. PDE solution generation enforces hard constraints, scientific inverse problems optimize soft objectives, and robot motion planning does both. We compare against the gradient and proximal guidance baselines of Section 2. For each task, we follow the setup of the paper that introduced it and run all experiments on a single NVIDIA A100 GPU (80 GB). See Appendix I for experiment details.

5.1 PDE Solution Generation

Setup. We evaluate zero-shot constrained generation of PDE solutions (Utkarsh et al., 2025). Given a PDE and a fixed initial or boundary condition (IC/BC), the goal is to sample the distribution of solutions under that condition. We guide a Functional Flow Matching (FFM) prior (Kerrigan et al., 2024) to satisfy both the condition and the PDE's conservation law. Five problems span four PDE families: Heat, 2D Navier–Stokes, Reaction–Diffusion, and Burgers with an IC or a BC (Appendix I). We compare against the proximal guidance methods PDM (Christopher et al., 2024) and PCFM (Utkarsh et al., 2025), which project the noisy sample and the clean estimate respectively. Our gradient guidance baseline is FunDPS (Yao et al., 2025). D-Flow (Ben-Hamu et al., 2024) serves as an optimization-based baseline that differentiates through the flow. All proximal guidance methods, including ProxDG, share one Gauss–Newton solver, so performance differences reflect the guidance mechanism alone. Following ECI (Cheng et al., 2025), we report the constraint residual (CE), the errors of the ensemble mean and standard deviation (MMSE, SMSE), and the Fréchet distance in Poseidon-B embeddings (Herde et al., 2024) (FPD). Heat, Reaction–Diffusion, and Burgers (BC) results appear in Appendix I.1.3.

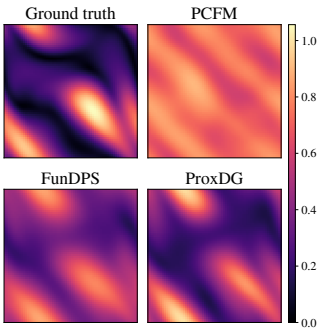

Figure 4: Ensemble variability on Navier–Stokes, shown as per-pixel standard deviation at the last frame. Samples differ only in external forcing.

Results. Tables 1 and 8 report feasibility (CE) and fidelity (MMSE, SMSE, FPD). The proximal guidance methods (PCFM, PDM, ProxDG) achieve zero constraint residual on every task, while FunDPS and D-Flow do not. Unlike PCFM and PDM, ProxDG also matches or exceeds the fidelity of the unconstrained baselines. It attains the best MMSE on four of five tasks, the best SMSE on Navier–Stokes and Burgers (IC), and the best FPD on Navier–Stokes. We attribute ProxDG's higher fidelity to its on-manifold guidance, whereas PCFM and PDM project samples off-manifold. In the Navier–Stokes ground truth (Figure 4), ensemble variability concentrates in a few regions. ProxDG reproduces this concentration, whereas PCFM spreads it out and FunDPS misjudges its magnitude.

Table 2: Black hole imaging at the two sparsest EHT observation-time ratios (3% and 10%), mean (std) over the 100-image InverseBench test set. Higher is better for PSNR and Blur PSNR, while $\tilde{\chi}^2_{\text{cp}}$ and $\tilde{\chi}^2_{\text{logca}}$ are ideal near 1. Best per column in bold, second-best underlined.

Observation time ratio	3%				10%			
Method	PSNR↑	Blur PSNR↑	$\tilde{\chi}^2_{\text{cp}}$	$\tilde{\chi}^2_{\text{logca}}$	PSNR↑	Blur PSNR↑	$\tilde{\chi}^2_{\text{cp}}$	$\tilde{\chi}^2_{\text{logca}}$
DPS	24.20 (3.72)	30.83 (5.58)	8.024 (24.336)	5.007 (5.750)	24.36 (3.72)	30.79 (5.75)	13.052 (43.087)	6.614 (26.789)
LGD	22.51 (3.76)	28.50 (5.49)	15.825 (16.838)	12.862 (12.663)	22.08 (3.75)	27.48 (5.09)	10.775 (21.684)	13.375 (56.397)
RED-diff	20.74 (2.62)	26.10 (3.35)	6.713 (6.925)	9.128 (19.052)	22.53 (3.02)	27.67 (4.53)	2.488 (2.925)	4.916 (13.221)
PnP-DM	24.25 (3.45)	30.49 (4.93)	2.201 (1.352)	<u>1.668</u> (0.551)	24.57 (3.47)	30.80 (5.22)	1.433 (0.417)	<u>1.336</u> (0.478)
DAPS	23.54 (3.28)	29.48 (4.88)	3.647 (3.287)	2.329 (1.354)	23.99 (3.56)	30.10 (5.13)	1.545 (0.705)	2.253 (9.903)
DiffPIR	24.12 (3.25)	30.45 (4.88)	14.085 (14.105)	10.545 (8.860)	23.84 (3.59)	30.04 (5.03)	5.374 (3.733)	5.205 (5.556)
SITCOM	25.12 (3.19)	<u>32.20</u> (4.76)	<u>1.648</u> (0.682)	1.918 (0.864)	25.49 (3.69)	32.56 (5.28)	1.415 (0.477)	2.179 (2.055)
DMPlug	25.11 (3.32)	31.65 (5.33)	2.201 (1.597)	1.955 (0.893)	25.40 (3.97)	32.16 (6.17)	1.392 (0.537)	1.348 (0.299)
ProxDG (DS)	<u>25.19</u> (3.59)	**32.28** (5.27)	**1.555** (0.689)	**1.595** (0.599)	**26.16** (3.63)	**33.52** (5.48)	**1.207** (0.265)	**1.324** (0.243)
ProxDG (PP)	**25.59** (3.55)	32.09 (5.28)	1.686 (1.295)	1.891 (0.812)	<u>26.10</u> (3.24)	<u>33.32</u> (5.28)	<u>1.268</u> (0.255)	1.557 (0.314)

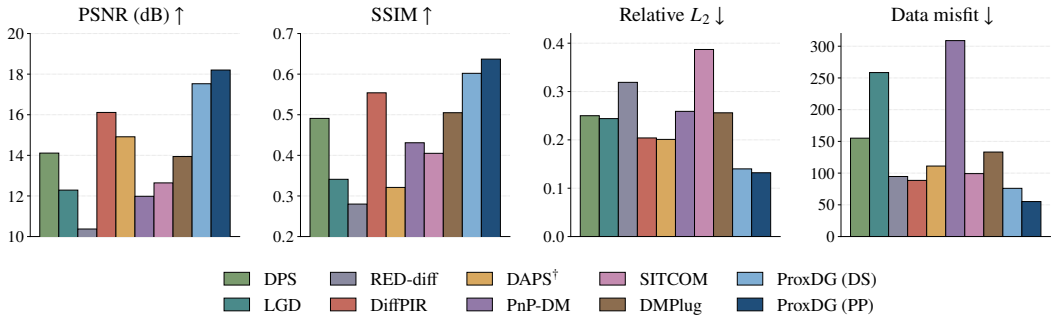

Figure 5: Full waveform inversion results. Mean PSNR, SSIM, relative ℓ_2, and data misfit over the 10-case InverseBench test set. †: initialized from a Gaussian-blurred ground truth, per InverseBench.

Unlike PCFM and PDM, ProxDG also runs with DPM-Solver++ (Lu et al., 2025) and stays exactly feasible with $4\times$ fewer steps (Appendix I.1.4).

5.2 Scientific Inverse Problems

Setup. We now turn from hard constraints to soft objectives, evaluating ProxDG on InverseBench (Zheng et al., 2025) across compressed-sensing MRI (Lustig et al., 2007), black hole imaging (BHI) (Chael et al., 2019), and full waveform inversion (FWI) (Virieux & Operto, 2010). BHI reconstructs a black hole's sky-plane intensity image from sparse interferometric measurements taken by the Event Horizon Telescope (EHT), and FWI recovers a subsurface velocity field from surface seismic recordings. Compressed-sensing MRI experiments are deferred to Appendix I.2.1. Each inverse problem has mature solvers (Parikh & Boyd, 2014; Neal et al., 2011; Beck & Teboulle, 2009), yet guidance methods rarely exploit them. ProxDG accepts any such solver, and here we design Gauss–Newton (GN) solvers that exploit each problem's least-squares structure (Appendices I.3.2 and I.4.2). We use InverseBench's pre-trained diffusion models, evaluation protocols, and plug-and-play baselines, and additionally run SITCOM (Alkhouri et al., 2025) and DMPlug (Wang et al., 2024). ProxDG (DS) returns the Euler–Maruyama sample, while ProxDG (PP) returns the proximal point from a final correction step (Section 4.3). On BHI, data misfit is measured by two χ^2 statistics, folded as $\tilde{\chi}^2 = \max(\chi^2, 1/\chi^2)$ so that the ideal value is 1.

Results. On BHI (Table 2, full results in Table 15), ProxDG attains the best value of all four metrics at both the 3% and 10% observation-time ratios. On FWI (Figure 5), ProxDG (PP) and ProxDG (DS) rank first and second on every metric, improving PSNR by 2.1 dB and cutting relative ℓ_2 error by 35% over DiffPIR, the strongest baseline. We ablate ProxDG's core components in Appendices I.2.5, I.3.4, and I.4.4, and highlight three findings.

1. *Sample quality scales with a well-matched solver.* Strengthening our Gauss–Newton solver raises PSNR on BHI by up to 4.05dB and lowers $\tilde{\chi}^2_{\text{cp}}$ from 5.77 to 1.23. A poorly matched solver can degrade samples with more iterations. Replacing our solver with gradient descent leaves $\tilde{\chi}^2$ as high as 76.5 on BHI, and additional descent steps often worsen data fit (Appendix I.3.4).

Table 3: Motion planning on three MPD environments. **Best** bold, second best underlined per row.

Env	Metric	ALM	DPS	MPD	PDM	CPS	SafeDiffuser	DiRecT	ProxDG
Planar4Link	FV ↑	0.115	0.462	0.332	0.615	0.544	0.425	0.690	**0.813**
	FP ↑	0.80	0.85	0.69	0.97	0.93	0.81	0.99	**1.00**
	Div ↑	13.5	15.2	31.2	8.7	24.0	15.1	39.5	**64.0**
Spheres3D	FV ↑	0.604	0.506	0.247	0.781	0.833	—	0.724	**0.881**
	FP ↑	**1.00**	0.99	0.94	**1.00**	**1.00**	—	0.99	**1.00**
	Div ↑	46.4	48.1	27.1	72.8	**78.8**	—	68.1	72.4
Warehouse	FV ↑	0.791	0.902	0.717	0.970	0.946	—	**0.996**	0.943
	FP ↑	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	—	**1.00**	**1.00**
	Div ↑	52.3	59.6	68.5	70.7	70.9	—	75.9	**86.4**

2. *Denoiser backpropagation improves sample quality.* We ablate denoiser backpropagation by directly substituting the proximal point for the clean estimate, as in DiffPIR, with the solver fixed. PSNR drops by 0.45 to 0.85 dB at every BHI ratio and by 1.71dB on FWI, where data misfit also rises $2.3\times$. We attribute the drop to the loss of on-manifold guidance (Section 4.2).
3. *ProxDG preserves sample quality with few sampling steps.* On FWI, ProxDG uses $5\times$ fewer diffusion steps than the baselines, yet it leads every metric while sampling fastest (Table 12).

5.3 Robot Motion Planning

Setup. Given a start and goal, we sample 100 joint trajectories that connect them and should be short and collision-free. We score the whole population, since closed-loop deployment relies on diverse fallback modes (Appendix I.5.1). We evaluate on the six environments of Carvalho et al. (2023), which range from a 2-DoF point mass to the 7-DoF Franka Panda. Each environment adds obstacles unseen during training and provides 100 held-out start–goal pairs. All planners sample from the same per-environment diffusion prior (a conditional Temporal U-Net with 15 DDIM steps), so performance differences reflect constraint handling alone. Our baselines span gradient guidance (DPS (Chung et al., 2023), MPD) and proximal guidance (PDM (Christopher et al., 2024), CPS (Narasimhan et al., 2025), SafeDiffuser (Xiao et al., 2025)), plus DiRecT (Giaretta et al., 2026), and every method is capped at 10 cost-gradient evaluations per denoising step. ProxDG solves its proximal subproblems with a batched augmented Lagrangian (ALM) solver (Appendix I.5.3), and the *ALM* baseline runs the same solver directly on samples from the prior. Hard constraints cover object collisions, self-collisions, and joint limits, while soft costs promote smooth trajectories. A goal given in joint space is enforced exactly, while an end-effector goal pose enters as a soft cost. Following Anderson et al. (2018), a trajectory is *valid* if it satisfies all hard constraints and reaches the goal within a small tolerance. We report three metrics. **FV** is the fraction of generated trajectories that are valid. **FP** is the fraction of problems with at least one valid trajectory. **Div** is the Vendi score (Friedman & Dieng, 2023), the effective number of distinct trajectories in the population. Appendix I.5 details the full setup and solver, and reports path length and end-effector error.

Results. ProxDG attains the highest FV in four of six environments, with the largest margin on Planar4Link (0.813 vs. 0.690 for DiRecT) (Tables 3 and 18). It is the only method that solves every problem in every environment, and it has the highest Div in four of six. The ALM baseline, which runs our solver directly on prior samples, degrades as DoF grows (FV 0.115 on Planar4Link), so the gains come from proximal sampling rather than solver strength alone. Appendix I.5.5 shows that ProxDG leads DPS and MPD at every solver budget $N \geq 5$.

6 Limitations and Future Work

ProxDG has three main limitations. First, each step costs a proximal solve and a denoiser backward pass, so ProxDG can sample slower than solver-free guidance (Table 12). Warm starts and skipping the backward pass (He et al., 2024) could reduce both costs. Second, sample quality hinges on a well-matched solver (Section 5.2). Third, ProxDG requires ℓ_2-regularized solvers, a restriction that Bregman–Moreau envelopes (Bauschke et al., 2018; Lau & Liu, 2022) could relax.

AI USE STATEMENT

We used generative AI tools, specifically Claude Code (Anthropic) and Codex (OpenAI), while preparing this work. We used them to implement parts of the batched ALM solver and the evaluation scripts, and to check the proofs of our theoretical results. We also used them to draft and edit parts of the paper text. We did not use generative AI tools to formulate our method, research questions, or theoretical claims. The authors reviewed all AI-assisted work, and we take full responsibility for the final content of this work.

ETHICS STATEMENT

This work develops a training-free guidance method and evaluates it on public benchmarks, namely the PDE setup of Utkarsh et al. (2025), InverseBench (Zheng et al., 2025), and the motion planning environments of Carvalho et al. (2023). It involves no human subjects. The MRI experiments use the de-identified fastMRI knee dataset (Zbontar et al., 2019; Knoll et al., 2020). Although ProxDG enforces constraints and fits measurements, its samples come from a learned prior and can contain plausible but incorrect structure. ProxDG reconstructions should therefore not replace clinical or scientific validation, and planned trajectories still require independent safety checks before deployment on physical robots.

REPRODUCIBILITY STATEMENT

Algorithms 2 and 3 give the full ProxDG samplers. Appendices C through H state every assumption and prove each theoretical result. Appendix I documents each experiment, including the problem setup, solver design, hyperparameter tuning, baselines, and evaluation protocol. All experiments build on public benchmarks and run on a single NVIDIA A100 GPU (80 GB).

REFERENCES

Ismail Alkhouri, Shijun Liang, Cheng-Han Huang, Jimmy Dai, Qing Qu, Saiprasad Ravishankar, and Rongrong Wang. SITCOM: Step-wise triple-consistent diffusion sampling for inverse problems. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 1128–1158. PMLR, 2025. URL https://proceedings.mlr.press/v267/alkhouri25b.html.

Luigi Ambrosio. Transport equation and Cauchy problem for non-smooth vector fields. In Bernard Dacorogna and Paolo Marcellini (eds.), *Calculus of Variations and Nonlinear Partial Differential Equations*, volume 1927 of *Lecture Notes in Mathematics*, pp. 1–41. Springer, Berlin, Heidelberg, 2008. doi: 10.1007/978-3-540-75914-0_1.

Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir R. Zamir. On evaluation of embodied navigation agents, 2018. URL https://arxiv.org/abs/1807.06757.

Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Roni Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=pzpWBbnwiJ.

Heinz H. Bauschke, Minh N. Dao, and Scott B. Lindstrom. Regularizing with bregman–moreau envelopes. *SIAM Journal on Optimization*, 28(4):3208–3228, 2018. doi: 10.1137/17M1130745. URL https://doi.org/10.1137/17M1130745.

Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009. doi: 10.1137/080716542.

Stephen Becker, Jalal Fadili, and Peter Ochs. On quasi-newton forward-backward splitting: Proximal calculus and convergence. *SIAM Journal on Optimization*, 29(4):2445–2481, 2019. doi: 10.1137/18M1167152. URL https://doi.org/10.1137/18M1167152.

Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-Flow: Differentiating through flows for controlled generation. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 3462–3483. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/ben-hamu24a.html.

Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. Nearly d-linear convergence bounds for diffusion models via stochastic localization. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=r5njV3BsuD.

Mario Bertero, Patrizia Boccacci, and Christine De Mol. *Introduction to inverse problems in imaging*. CRC press, 2021.

Matthieu Blanke, Yongquan Qu, Sara Shamekh, and Pierre Gentine. Strictly constrained generative modeling via split augmented langevin sampling, 2025. URL https://arxiv.org/abs/2505.18017.

Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *Transactions on Machine Learning Research*, 2022. URL https://openreview.net/forum?id=MhK5aXo3gB.

Benjamin Boys, Mark Girolami, Jakiw Pidstrigach, Sebastian Reich, Alan Mosca, and Omer Deniz Akyildiz. Tweedie moment projected diffusions for inverse problems. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=4unJi0qrTE. Featured Certification.

Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected langevin monte carlo. *Discrete & Computational Geometry*, 59(4):757–783, 2018.

J. Carvalho, A. T. Le, M. Baierl, D. Koert, and J. Peters. Motion planning diffusion: Learning and planning of robot motions with diffusion models. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023. URL https://arxiv.org/abs/2308.01557.

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

Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 4735–4763. PMLR, 23–29 Jul 2023a. URL https://proceedings.mlr.press/v202/chen23q.html.

Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023b. URL https://openreview.net/forum?id=zyLVMgsZ0U_.

Yuxiao Chen, Ugo Rosolia, Wyatt Ubellacker, Noel Csomay-Shanklin, and Aaron D. Ames. Interactive multi-modal motion planning with branch model predictive control. *IEEE Robotics and Automation Letters*, 7(2):5365–5372, 2022. doi: 10.1109/LRA.2022.3156648.

Chaoran Cheng, Boran Han, Danielle C. Maddix, Abdul Fatir Ansari, Andrew Stuart, Michael W. Mahoney, and Bernie Wang. Gradient-free generation for hard-constrained systems. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=teE4pl9ftK.

Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin CM Burchfiel, and Shuran Song. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.026.

Muthu Chidambaram, Khashayar Gatmiry, Sitan Chen, Holden Lee, and Jianfeng Lu. What does guidance do? a fine-grained analysis in a simple setting. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 84968–85005. Curran Associates, Inc., 2024. doi: 10.52202/079017-2698. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/9a3942c235daa9c5f62a8598ae81a946-Paper-Conference.pdf.

Wonshik Choi, Christopher Fang-Yen, Kamran Badizadegan, Seungeun Oh, Niyom Lue, Ramachandra R Dasari, and Michael S Feld. Tomographic phase microscopy. *Nature methods*, 4(9):717–719, 2007.

Jacob K. Christopher, Stephen Baek, and Ferdinando Fioretto. Constrained synthesis with projected diffusion models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 89307–89333. Curran Associates, Inc., 2024. doi: 10.52202/079017-2834. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/a28af221f2f70be183afc16797a56b91-Paper-Conference.pdf.

Jacob K Christopher, Austin Seamann, Jingyi Cui, Sagar D Khare, and Ferdinando Fioretto. Constrained diffusion for protein design with hard structural constraints. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=kkvqVRu2Zy.

Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical Image Analysis*, pp. 102479, 2022.

Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 25683–25696. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/a48e5877c7bf86a513950ab23b360498-Paper-Conference.pdf.

Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=OnD9zGAGT0k.

Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi S. Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=kKF8_K-mBbS.

Matias G. Delgadino, Sebastien Motsch, Advait Parulekar, William Porteous, and Sanjay Shakkottai. Diffusion-based posterior sampling: A Feynman-Kac analysis of bias and stability, 2026. URL https://arxiv.org/abs/2605.06538.

Alexander Denker, Francisco Vargas, Shreyas Padhy, Kieran Didi, Simon V Mathis, Riccardo Barbano, Vincent Dutordoir, Emile Mathieu, Urszula Julia Komorowska, and Pietro Lio. DEFT: Efficient fine-tuning of diffusion models by learning the generalised h-transform. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=AKBTFQhCjm.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.

Jing Dong, Mustafa Mukadam, Frank Dellaert, and Byron Boots. Motion planning as probabilistic inference using Gaussian processes and factor graphs. In *Proceedings of Robotics: Science and Systems (RSS)*, 2016.

Jing Dong, Mustafa Mukadam, Byron Boots, and Frank Dellaert. Sparse Gaussian processes on matrix Lie groups: A unified framework for optimizing continuous-time trajectories. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6497–6504. IEEE, 2018.

Hongkun Dou, Zeyu Li, Xingyu Jiang, Hongjue Li, Lijun Yang, Wen Yao, and Yue Deng. You Only Look One Step: Accelerating Backpropagation in Diffusion Sampling With Gradient Shortcuts . *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 48(05):5050–5067, May 2026. ISSN 1939-3539. doi: 10.1109/TPAMI.2025.3647857. URL https://doi.ieeecomputersociety.org/10.1109/TPAMI.2025.3647857.

Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 301–308, 2013. doi: 10.1109/HRI.2013.6483603.

Yilun Du, Conor Durkan, Robin Strudel, Joshua B. Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, and Will Sussman Grathwohl. Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and MCMC. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 8489–8510. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/du23a.html.

Alain Durmus, Éric Moulines, and Marcelo Pereyra. Efficient bayesian computation by proximal markov chain monte carlo: When langevin meets moreau. *SIAM Journal on Imaging Sciences*, 11(1):473–506, 2018a. doi: 10.1137/16M1108340. URL https://doi.org/10.1137/16M1108340.

Alain Durmus, Éric Moulines, and Marcelo Pereyra. Efficient bayesian computation by proximal markov chain monte carlo: When langevin meets moreau. *SIAM Journal on Imaging Sciences*, 11(1):473–506, 2018b. doi: 10.1137/16M1108340. URL https://doi.org/10.1137/16M1108340.

Matthias J. Ehrhardt, Lorenz Kuger, and Carola-Bibiane Schönlieb. Proximal langevin sampling with inexact proximal mapping. *SIAM Journal on Imaging Sciences*, 17(3):1729–1760, 2024. doi: 10.1137/23M1593565. URL https://doi.org/10.1137/23M1593565.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Zion English, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Lawrence C. Evans and Ronald F. Gariepy. *Measure Theory and Fine Properties of Functions*. CRC Press, revised edition, 2015.

Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3): 418–491, 1959.

Ruiqi Feng, Chenglei Yu, Wenhao Deng, Peiyan Hu, and Tailin Wu. On the guidance of flow matching. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 16993–17029. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/feng25s.html.

Pete Florence, Corey Lynch, Andy Zeng, Oscar Ramirez, Ayzaan Wahid, Laura Downs, Adrian Wong, Johnny Lee, Igor Mordatch, and Jonathan Tompson. Implicit behavioral cloning. *Conference on Robot Learning (CoRL)*, 2021.

Dan Friedman and Adji Bousso Dieng. The vendi score: A diversity evaluation metric for machine learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=g97OHbQyk1.

Darshan Gadginmath, Ahmed Allibhoy, and Fabio Pasqualetti. Provably safe generative sampling with constricting barrier functions, 2026. URL https://arxiv.org/abs/2602.21429.

Jonathan D Gammell, Siddhartha S Srinivasa, and Timothy D Barfoot. Informed rrt*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *2014 IEEE/RSJ international conference on intelligent robots and systems*, pp. 2997–3004. IEEE, 2014.

Paolo Giaretta, Zeyang Li, and Navid Azizan. DiRecT: Safe diffusion-based planning via receding-horizon denoising, 2026. URL https://arxiv.org/abs/2606.15359.

Alexandros Graikos, Nebojsa Jojic, and Dimitris Samaras. Fast constrained sampling in pre-trained diffusion models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=3kVM0m60Q5.

Yingqing Guo, Hui Yuan, Yukang Yang, Minshuo Chen, and Mengdi Wang. Gradient guidance for diffusion models: An optimization perspective. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 90736–90770. Curran Associates, Inc., 2024. doi: 10.52202/079017-2881. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/a5059a9a389ccc76da85760ea79490d8-Paper-Conference.pdf.

Shivam Gupta, Ajil Jalal, Aditya Parulekar, Eric Price, and Zhiyang Xun. Diffusion posterior sampling is computationally intractable. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 17020–17059. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/gupta24a.html.

Andreas Habring, Alexander Falk, Martin Zach, and Thomas Pock. Diffusion at absolute zero: Langevin sampling using successive moreau envelopes [journal paper], 2025. URL https://arxiv.org/abs/2503.22258.

Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, and Stefano Ermon. Manifold preserving guided diffusion. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=o3BxOLoxm1.

Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 72525–72624. Curran Associates, Inc., 2024. doi: 10.52202/079017-2311. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/84e1b1ec17bb11c57234e96433022a9a-Paper-Conference.pdf.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.

Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 8633–8646. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/39235c56aef13fb05a6adc95eb9d8d66-Paper-Conference.pdf.

William Huang, Yifeng Jiang, Tom Van Wouwe, and C. Karen Liu. Constrained diffusion with trust sampling. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 93849–93873. Curran Associates, Inc., 2024. doi: 10.52202/079017-2976. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/aa8164b8a8e484fe31885657fea14997-Paper-Conference.pdf.

Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust compressed sensing mri with deep generative priors. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 14938–14954. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/7d6044e95a16761171b130dcb476a43e-Paper.pdf.

Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9902–9915. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/janner22a.html.

Abderrahim Jourani, Lionel Thibault, and Dariusz Zagrodny. Differential properties of the moreau envelope. *Journal of Functional Analysis*, 266(3):1185–1237, 2014. ISSN 0022-1236. doi: https://doi.org/10.1016/j.jfa.2013.11.008. URL https://www.sciencedirect.com/science/article/pii/S0022123613004424.

Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=ANvmVS2Yr0.

Olav Kallenberg. *Foundations of Modern Probability*. Probability and Its Applications. Springer, 2nd edition, 2002.

Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 26565–26577. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/a98846e9d9cc01cfb87eb694d946ce6b-Paper-Conference.pdf.

Gavin Kerrigan, Giosue Migliorini, and Padhraic Smyth. Functional flow matching. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3934–3942. PMLR, 02–04 May 2024. URL https://proceedings.mlr.press/v238/kerrigan24a.html.

Andrey Kharitenko, Zebang Shen, Riccardo De Santi, Niao He, and Florian Dorfler. Landing with the score: Riemannian optimization through denoising. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=xZNoeX0z9f.

Hyunjik Kim, George Papamakarios, and Andriy Mnih. The lipschitz constant of self-attention. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5562–5571. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/kim21i.html.

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

Sunwoo Kim, Minkyu Kim, and Dongmin Park. Test-time alignment of diffusion models without reward over-optimization. In *The Thirteenth International Conference on Learning Representations*, 2025b. URL https://openreview.net/forum?id=vi3DjUhFVm.

Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J. Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1):e190007, 2020.

Vince Kurtz and Alexander Davydov. Constrained flow matching via Lagrangian dual flows, 2026. URL https://arxiv.org/abs/2607.04513.

Chieh-Hsin Lai, Yang Song, Dongjun Kim, Yuki Mitsufuji, and Stefano Ermon. The principles of diffusion models, 2025. URL https://arxiv.org/abs/2510.21890.

Tim Tsz-Kit Lau and Han Liu. Bregman proximal Langevin Monte Carlo via Bregman-moreau envelopes. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 12049–12077. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/lau22a.html.

Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 22870–22882. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/8ff87c96935244b63503f542472462b3-Paper-Conference.pdf.

Jason D. Lee, Yuekai Sun, and Michael A. Saunders. Proximal newton-type methods for minimizing composite functions. *SIAM Journal on Optimization*, 24(3):1420–1443, 2014. doi: 10.1137/130921428. URL https://doi.org/10.1137/130921428.

Gunther Leobacher and Alexander Steinicke. Existence, uniqueness and regularity of the projection onto differentiable manifolds. *Annals of Global Analysis and Geometry*, 60:559–587, 2021. doi: 10.1007/s10455-021-09788-z.

Jinhao Liang, Jacob K Christopher, Sven Koenig, and Ferdinando Fioretto. Simultaneous multi-robot motion planning with projected diffusion models. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 37162–37180. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/liang25e.html.

Pierre-Louis Lions and Benjamin Seeger. Transport equations and flows with one-sided Lipschitz velocity fields. *Archive for Rational Mechanics and Analysis*, 248(86), 2024.

Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code, 2024. URL https://arxiv.org/abs/2412.06264.

M. Louboutin, M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, P. Velesko, and G. J. Gorman. Devito (v3.1.0): an embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development*, 12(3):1165–1187, 2019. doi: 10.5194/gmd-12-1165-2019. URL https://gmd.copernicus.org/articles/12/1165/2019/.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-Solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 22:730–751, 2025. doi: 10.1007/s11633-025-1562-4. URL https://doi.org/10.1007/s11633-025-1562-4.

Hao Luan, Yi Xian Goh, See-Kiong Ng, and Chun Kai Ling. Projected coupled diffusion for test-time constrained joint generation. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=1FEm5JLpvg.

Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

Nanye Ma, Mark Goldstein, Michael S. Albergo, Nicholas M. Boffi, Eric Vanden-Eijnden, and Saining Xie. SiT: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, 2024. URL https://arxiv.org/abs/2401.08740.

Yosuke Mizuno. Grmhd simulations and modeling for jet formation and acceleration region in agns. *Universe*, 8(2):85, 2022.

Sreyas Mohan, Zahra Kadkhodaie, Eero P. Simoncelli, and Carlos Fernandez-Granda. Robust and interpretable blind image denoising via bias-free convolutional neural networks. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1906.05478.

Jean-Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965. doi: 10.24033/bsmf.1625. URL http://www.numdam.org/item/BSMF_1965__93__273_0/.

Mustafa Mukadam, Jing Dong, Xinyan Yan, Frank Dellaert, and Byron Boots. Continuous-time Gaussian process motion planning via probabilistic inference. volume 37, pp. 1319–1340, 2018.

Sai Shankar Narasimhan, Shubhankar Agarwal, Litu Rout, Sanjay Shakkottai, and Sandeep P. Chinchali. Constrained posterior sampling: Time series generation with hard constraints. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=iBLAYaEKdL.

Parikh Neal, Chu Eric, Peleato Borja, and Eckstein Jonathan. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.

Gregory Ongie, Ajil Jalal, Christopher A Metzler, Richard G Baraniuk, Alexandros G Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020.

Stanley Osher, Howard Heaton, and Samy Wu Fung. A hamilton–jacobi-based proximal operator. *Proceedings of the National Academy of Sciences*, 120(14):e2220469120, 2023. doi: 10.1073/pnas.2220469120. URL https://www.pnas.org/doi/abs/10.1073/pnas.2220469120.

Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3): 127–239, 2014. doi: 10.1561/2400000003.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.

Xinyu Peng, Ziyang Zheng, Wenrui Dai, Nuoqian Xiao, Chenglin Li, Junni Zou, and Hongkai Xiong. Improving diffusion models for inverse problems using optimal posterior covariance. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 40347–40370. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/peng24h.html.

Marcelo Pereyra. Proximal markov chain monte carlo algorithms. *Statistics and Computing*, 26 (4):745–760, July 2016. ISSN 0960-3174. doi: 10.1007/s11222-015-9567-4. URL https://doi.org/10.1007/s11222-015-9567-4.

Jakiw Pidstrigach. Score-based generative models detect manifolds. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 35852–35865. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/e8fb575e3ede31f9b8c05d53514eb7c6-Paper-Conference.pdf.

Nathan Ratliff, Matt Zucker, J. Andrew Bagnell, and Siddhartha Srinivasa. Chomp: Gradient optimization techniques for efficient motion planning. In *2009 IEEE International Conference on Robotics and Automation*, pp. 489–494, 2009. doi: 10.1109/ROBOT.2009.5152817.

Yinuo Ren, Wenhao Gao, Lexing Ying, Grant M. Rotskoff, and Jiequn Han. Driftlite: Lightweight drift control for inference-time scaling of diffusion models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2509.21655.

Marien Renaud, Valentin De Bortoli, Arthur Leclaire, and Nicolas Papadakis. From stability of langevin diffusion to convergence of proximal MCMC for non-log-concave sampling. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=mcYO3DccbQ.

Severi Rissanen, Markus Heinonen, and Arno Solin. Free hunch: Denoiser covariance estimation for diffusion models without extra costs. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=4JK2XMGUc8.

Omer Rochman-Sharabi and Gilles Louppe. Predict-project-renoise: Sampling diffusion models under hard constraints, 2026. URL https://arxiv.org/abs/2601.21033.

R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*, volume 317 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, 1998. ISBN 978-3-540-62772-2. doi: 10.1007/978-3-642-02431-3. Corrected 3rd printing, 2009.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

Ralf Römer, Alexander von Rohr, and Angela Schoellig. Diffusion predictive control with constraints. In *Proceedings of the 7th Annual Learning for Dynamics & Control Conference*, pp. 791–803. PMLR, 2025.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi (eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4.

Litu Rout, Yujia Chen, Nataniel Ruiz, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu. RB-modulation: Training-free stylization using reference-based modulation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=bnINPG5A32.

Amirmojtaba Sabour, Michael S. Albergo, Carles Domingo-Enrich, Nicholas M. Boffi, Sanja Fidler, Karsten Kreis, and Eric Vanden-Eijnden. Test-time scaling of diffusions with flow maps, 2025. URL https://arxiv.org/abs/2511.22688.

Kartik Sharma, Srijan Kumar, and Rakshit Trivedi. Diffuse, sample, project: Plug-and-play controllable graph generation. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 44545–44564. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/sharma24b.html.

Yifei Shen, Xinyang Jiang, Yifan Yang, Yezhen Wang, Dongqi Han, and Dongsheng Li. Understanding and improving training-free loss-based diffusion guidance. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 108974–109002. Curran Associates, Inc., 2024. doi: 10.52202/079017-3460. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/c4edc5113b4ffd4632718558fb66b9ef-Paper-Conference.pdf.

Changsoo Shin, Seonghyung Jang, and Dong-Joo Min. Improved amplitude preservation for prestack depth migration by inverse scattering theory. *Geophysical Prospecting*, 49(5):592–606, 2001. doi: https://doi.org/10.1046/j.1365-2478.2001.00279.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2478.2001.00279.x.

Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-kac correctors in diffusion: Annealing, guidance, and product of experts. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=Vhc0KrcqWu.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021a. URL https://openreview.net/forum?id=St1giarCHLP.

Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*, 2023a. URL https://openreview.net/forum?id=9_gsMA8MRKQ.

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 Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 32483–32498. PMLR, 23–29 Jul 2023b. URL https://proceedings.mlr.press/v202/song23k.html.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

Jan Pawel Stanczuk, Georgios Batzolis, Teo Deveney, and Carola-Bibiane Schönlieb. Diffusion models encode the intrinsic dimension of data manifolds. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 46412–46440. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/stanczuk24a.html.

He Sun and Katherine L. Bouman. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(3):2628–2637, May 2021. doi: 10.1609/aaai.v35i3.16366. URL https://ojs.aaai.org/index.php/AAAI/article/view/16366.

Simo Särkkä and Arno Solin. *Applied Stochastic Differential Equations*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2019.

Marc Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pp. 1049–1056, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585161. doi: 10.1145/1553374.1553508. URL https://doi.org/10.1145/1553374.1553508.

Utkarsh, Pengfei Cai, Alan Edelman, Rafael Gómez-Bombarelli, and Christopher Rackauckas. Physics-constrained flow matching: Sampling generative models with hard constraints. In *Advances in Neural Information Processing Systems*, 2025.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Enrico Ventura, Beatrice Achilli, Gianluigi Silvestri, Carlo Lucibello, and Luca Ambrogioni. Manifolds, random matrices and spectral gaps: The geometric phases of generative diffusion. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=KlN00vQEY2.

Alexander Yu. Veretennikov. On strong solutions and explicit formulas for solutions of stochastic integral equations. *Mathematics of the USSR-Sbornik*, 39(3):387–403, 1981.

Cédric Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.

Jean Virieux and Stéphane Operto. An overview of full-waveform inversion in exploration geophysics. 2010.

Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/d54e99a6c03704e95e6965532dec148b-Paper.pdf.

Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. https://github.com/huggingface/diffusers, 2022.

Hengkang Wang, Xu Zhang, Taihui Li, Yuxiang Wan, Tiancong Chen, and Ju Sun. DMPlug: A plug-in method for solving inverse problems with diffusion models. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/d5c04aa72b92c53bda5b525b60958295-Abstract-Conference.html.

Grady Williams, Paul Drews, Brian Goldfain, James M. Rehg, and Evangelos A. Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *Trans. Rob.*, 34(6):1603–1622, December 2018. ISSN 1552-3098. doi: 10.1109/TRO.2018.2865891. URL https://doi.org/10.1109/TRO.2018.2865891.

Luhuan Wu, Brian Trippe, Christian Naesseth, David M. Blei, and John P. Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in Neural Information Processing Systems*, volume 36. Curran Associates, Inc., 2023. doi: 10.52202/075280-1363. URL https://papers.neurips.cc/paper_files/paper/2023/hash/63e8bc7bbf1cfea36d1d1b6538aecce5-Abstract-Conference.html.

Wei Xiao, Tsun-Hsuan Wang, Chuang Gan, Ramin Hasani, Mathias Lechner, and Daniela Rus. Safediffuser: Safe planning with diffusion probabilistic models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=ig2wk7kK9J.

Jeongyong Yang, Seunghwan Jang, and SooJean Han. Safeflowmatcher: Safe and fast planning using flow matching with control barrier functions. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=refcXHU1Nh.

Lingxiao Yang, Shutong Ding, Yifan Cai, Jingyi Yu, Jingya Wang, and Ye Shi. Guidance with spherical Gaussian constraint for conditional diffusion. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 56071–56095. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/yang24h.html.

Jiachen Yao, Abbas Mammadov, Julius Berner, Gavin Kerrigan, Jong Chul Ye, Kamyar Azizzadenesheli, and Anima Anandkumar. Guided diffusion sampling on function spaces with applications to PDEs. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=oAgwvZay2U.

Haotian Ye, Haowei Lin, Jiaqi Han, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Zou, and Stefano Ermon. Tfg: Unified training-free guidance for diffusion models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 22370–22417. Curran Associates, Inc., 2024. doi: 10.52202/079017-0704. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/2818054fc6de6dacdda0f142a3475933-Paper-Conference.pdf.

Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. Freedom: Training-free energy-guided conditional diffusion model. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 23117–23127, 2023. doi: 10.1109/ICCV51070.2023.02118.

Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 16010–16021, October 2023.

Stefano Zampini, Jacob K Christopher, Luca Oneto, Davide Anguita, and Ferdinando Fioretto. Training-free constrained generation with stable diffusion models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=TrNB08KuHK.

Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, Nafissa Yakubova, James Pinkerton, Duo Wang, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastmri: An open dataset and benchmarks for accelerated mri, 2019. URL https://arxiv.org/abs/1811.08839.

Zheyuan Zhan, Hongchen Li, Can Wang, Yinfei Ma, Mingzhen Huang, Ruoshi Bai, Jiawei Chen, Siwei Lyu, and Defang Chen. Conditioning matters: Stabilizing inversion and attention in diffusion image editing, 2026. URL https://arxiv.org/abs/2606.14125.

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, June 2025a. doi: 10.1109/cvpr52734.2025.01946. URL http://dx.doi.org/10.1109/CVPR52734.2025.01946.

Jichen Zhang, Liqun Zhao, Antonis Papachristodoulou, and Jack Umenberger. Constrained diffusers for safe planning and control. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b. URL https://openreview.net/forum?id=tahkGZjjWA.

Xiangcheng Zhang, Haowei Lin, Haotian Ye, James Zou, Jianzhu Ma, Yitao Liang, and Yilun Du. Inference-time scaling of diffusion models through classical search, 2025c. URL https://arxiv.org/abs/2505.23614.

Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy Feng, Caifeng Zou, Yu Sun, Nikola Borislavov Kovachki, Zachary E Ross, Katherine Bouman, and Yisong Yue. Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=U3PBITXNG6.

Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 1219–1229, June 2023.

Algorithm 1 Proximal Diffusion Guidance (ProxDG) via Langevin Dynamics

Require: Score estimator $s_\theta(\cdot, t)$; annealing schedule $\{\lambda_t\}_{t=0}^T$; step sizes $\{\tau_t\}_{t=0}^T$; sampling iterations K
Ensure: Clean sample $\mathbf{x}_0$

$\mathbf{x}_T^{(K)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ▷ Initialize from prior
for $t = T-1$ **down to** 0 **do** ▷ Outer loop: diffusion timesteps
  $\mathbf{x}_t^{(0)} \leftarrow \mathbf{x}_{t+1}^{(K)}$
  for $k = 0$ **to** $K-1$ **do** ▷ Inner loop: sampling iterations
    $\hat{\mathbf{x}}_0^{(k)} \leftarrow \hat{\mathbf{x}}_0(\mathbf{x}_t^{(k)}, t)$ ▷ Tweedie's estimate
    $\mathbf{z}_t^{(k)} \leftarrow \mathrm{prox}_{\lambda_t f}(\hat{\mathbf{x}}_0^{(k)})$ ▷ Proximal point
    $\nabla f_t^\lambda(\mathbf{x}_t^{(k)}) \leftarrow \frac{1}{\lambda_t}\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}_t^{(k)}, t)^\top \left[\hat{\mathbf{x}}_0^{(k)} - \mathbf{z}_t^{(k)}\right]$ ▷ Guidance gradient
    $\epsilon_t^{(k+1)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
    $\mathbf{x}_t^{(k+1)} \leftarrow \mathbf{x}_t^{(k)} + \tau_t s_\theta(\mathbf{x}_t^{(k)}, t) - \tau_t \nabla f_t^\lambda(\mathbf{x}_t^{(k)}) + \sqrt{2\tau_t}\epsilon_t^{(k+1)}$
  end for
end for
return $\mathbf{x}_0^{(K)}$

A Algorithms

A.1 The SDE Sampling Algorithm

Algorithm 2 integrates the guided SDE in Eq. 12 using Euler–Maruyama (EM). We initialize $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and integrate backward to $t = 0$ with uniform step size $\Delta t = T/N$. For constrained tasks, we optionally apply the final proximal correction from Section 4.3, replacing $\mathbf{x}_0$ with $\mathrm{prox}_{\lambda_{\Delta t} f}(\mathbf{x}_0)$.

Algorithm 2 Proximal Diffusion Guidance (ProxDG) via SDE Integration

Require: Score estimator $s_\theta(\cdot, t)$; guidance scale $w \geq 0$; annealing schedule $\{\lambda_t\}_{t\in(0,T]}$; step size Δt; proximal correction flag, correct
Ensure: Clean sample $\mathbf{x}_0$

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ▷ Initialize from prior
for $n = N$ **down to** 1 **do**
  $t \leftarrow n\Delta t$
  $\hat{\mathbf{x}}_0 \leftarrow \hat{\mathbf{x}}_0(\mathbf{x}_t, t)$ ▷ Clean-sample estimate
  $\mathbf{z}_t \leftarrow \mathrm{prox}_{\lambda_t f}(\hat{\mathbf{x}}_0)$ ▷ Proximal point
  $\nabla f_t^\lambda(\mathbf{x}_t) \leftarrow \frac{1}{\lambda_t}\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}_t, t)^\top \left[\hat{\mathbf{x}}_0 - \mathbf{z}_t\right]$ ▷ Guidance gradient
  $\overline{\mathbf{w}}_{t-\Delta t} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
  $\mathbf{x}_{t-\Delta t} \leftarrow \mathbf{x}_t - \left[\mu(t)\mathbf{x}_t - g(t)^2\left(s_\theta(\mathbf{x}_t, t) - w\nabla f_t^\lambda(\mathbf{x}_t)\right)\right]\Delta t + g(t)\sqrt{\Delta t}\,\overline{\mathbf{w}}_{t-\Delta t}$
end for
if correct **then**
  $\mathbf{x}_0 \leftarrow \mathrm{prox}_{\lambda_{\Delta t} f}(\mathbf{x}_0)$ ▷ Proximal correction
end if
return $\mathbf{x}_0$

A.2 The ODE Sampling Algorithm

Algorithm 3 integrates the guided ODE using Euler's method. We initialize $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and integrate backward to $t = 0$ with uniform step size $\Delta t = T/N$. For the probability-flow ODE in Eq. 222, we use

$$v_\theta(\mathbf{x}, t) = \mu(t)\mathbf{x} - \tfrac{1}{2}g(t)^2 s_\theta(\mathbf{x}, t), \qquad c(t) = \tfrac{1}{2}g(t)^2 \tag{13}$$

For flow matching, we use the learned velocity field with

$$c(t) = \sigma_t\dot{\sigma}_t - \frac{\dot{\alpha}_t}{\alpha_t}\sigma_t^2 \tag{14}$$

For constrained tasks, we optionally apply the same final proximal correction as in Algorithm 2.

Algorithm 3 Proximal Diffusion Guidance (ProxDG) via ODE Integration

Require: Velocity field $v_\theta(\cdot, t)$; guidance scale $w \geq 0$; annealing schedule $\{\lambda_t\}_{t\in(0,T]}$; step size Δt; proximal correction flag, correct

Ensure: Clean sample $\mathbf{x}_0$

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ▷ Initialize from prior

for $n = N$ **down to** 1 **do**

  $t \leftarrow n\Delta t$

  $\hat{\mathbf{x}}_0 \leftarrow \hat{\mathbf{x}}_0(\mathbf{x}_t, t)$ ▷ Clean-sample estimate

  $\mathbf{z}_t \leftarrow \mathrm{prox}_{\lambda_t f}(\hat{\mathbf{x}}_0)$ ▷ Proximal point

  $\nabla f_t^\lambda(\mathbf{x}_t) \leftarrow \frac{1}{\lambda_t}\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}_t, t)^\top\left[\hat{\mathbf{x}}_0 - \mathbf{z}_t\right]$ ▷ Guidance gradient

  $\mathbf{x}_{t-\Delta t} \leftarrow \mathbf{x}_t - \left[v_\theta(\mathbf{x}_t, t) + wc(t)\nabla f_t^\lambda(\mathbf{x}_t)\right]\Delta t$ ▷ Euler update

end for

if correct **then**

  $\mathbf{x}_0 \leftarrow \mathrm{prox}_{\lambda_{\Delta t} f}(\mathbf{x}_0)$ ▷ Proximal correction

end if

return $\mathbf{x}_0$

B Discussion

B.1 Related Work

The substitution updates come from the proximal sampling literature, where projected (Bubeck et al., 2018) and proximal Langevin algorithms (Ehrhardt et al., 2024; Renaud et al., 2025; Pereyra, 2016; Durmus et al., 2018a) apply the proximal operator to the sample at every step. A recent line of work instead treats the proximal point as a gradient target (Blanke et al., 2025; Durmus et al., 2018b). Most relevant to our setting, DAZ (Habring et al., 2025) performs this gradient-based proximal sampling along an annealed probability path to some target Gibbs distribution. We also propose an annealed proximal sampler, but for diffusion guidance. Unlike prior proximal guidance works, our sampler has gradient form, inheriting TFG advances while still admitting proximal solves in place of plain gradient evaluations.

B.2 Background

The forward SDE in Eq. 1 is linear with additive Gaussian noise, so the transition kernel $p(\mathbf{x}_t|\mathbf{x}_0)$ is Gaussian. The mean and covariance of $p(\mathbf{x}_t|\mathbf{x}_0)$ satisfy linear differential equations, whose solutions give (Särkkä & Solin, 2019)

$$\begin{aligned}&\mathbf{x}_t|\mathbf{x}_0 \sim \mathcal{N}(\alpha_t\mathbf{x}_0, \sigma_t^2\mathbf{I})\\ &\alpha_t = \exp\left(\int_0^t \mu(s)\,ds\right), \qquad \sigma_t^2 = \alpha_t^2\int_0^t \frac{g(s)^2}{\alpha_s^2}\,ds\end{aligned} \tag{15}$$

We can therefore sample $\mathbf{x}_t$ directly as $\mathbf{x}_t = \alpha_t\mathbf{x}_0 + \sigma_t\boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Different choices of μ and g recover the VP and VE noise schedules (Song et al., 2021b), as well as the EDM parameterization (Karras et al., 2022).

B.3 Smoothing Guidance with Moreau Envelopes

B.3.1 Moreau Envelope Properties

The lower envelope of paraboloids: We can construct the Moreau envelope geometrically from the graph of f, defined as the set of points $(\mathbf{z}, f(\mathbf{z}))$ with $f(\mathbf{z}) < \infty$. Fix $\lambda > 0$ and place an upward-opening paraboloid at every point of the graph, with curvature $1/\lambda$.

$$g_{\mathbf{z}}(\mathbf{u}) \coloneqq f(\mathbf{z}) + \frac{1}{2\lambda}\|\mathbf{u}-\mathbf{z}\|^2 \tag{16}$$

For a fixed $\mathbf{z}$, the argument $\mathbf{u}$ ranges over a single paraboloid. If we instead fix $\mathbf{u}$ and vary $\mathbf{z}$, the values $g_{\mathbf{z}}(\mathbf{u})$ describe the heights of all paraboloids at the same location $\mathbf{u}$. The infimum $\inf_{\mathbf{z}} g_{\mathbf{z}}(\mathbf{u})$ gives the lowest of these heights. We define the Moreau envelope as the function that assigns, to each $\mathbf{u}$, the lowest height attained by any paraboloid $g_\cdot(\mathbf{u})$:

$$\Psi_f(\mathbf{u}, \lambda) \coloneqq \inf_{\mathbf{z}} g_{\mathbf{z}}(\mathbf{u}) \tag{17}$$

Its graph therefore forms the lower boundary of the paraboloid family $g_{\mathbf{z}}(\cdot)_{\mathbf{z}}$ (Figure 6).

Several properties of the Moreau envelope follow directly from its paraboloid construction. For example, $\Psi_f(\mathbf{u}, \lambda) \le f(\mathbf{u})$, since the envelope lies below every paraboloid and $g_{\mathbf{u}}(\mathbf{u}) = f(\mathbf{u})$.

To connect the envelope's gradient to the proximal map, consider the paraboloids that touch the envelope at $\mathbf{u}$. A paraboloid $g_{\mathbf{z}^*}$ touches precisely when $g_{\mathbf{z}^*}(\mathbf{u}) = \Psi_f(\mathbf{u}, \lambda)$, meaning that $\mathbf{z}^*$ minimizes the paraboloid height at $\mathbf{u}$. These minimizing indices form the proximal set of f at $\mathbf{u}$:

$$\mathrm{Prox}_{\lambda f}(\mathbf{u}) \coloneqq \arg\min_{\mathbf{z}}\left\{f(\mathbf{z}) + \frac{1}{2\lambda}\|\mathbf{u}-\mathbf{z}\|^2\right\} \tag{18}$$

For any $\mathbf{z}^* \in \mathrm{Prox}_{\lambda f}(\mathbf{u})$, $g_{\mathbf{z}^*}$ lies above the envelope everywhere and touches it at $\mathbf{u}$. The envelope therefore cannot increase from $\mathbf{u}$ faster than $g_{\mathbf{z}^*}$, as the next lemma formalizes.

Lemma B.1 (Touching paraboloid bound)**.** *For any $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$, $h \in \mathbb{R}$, and $\mathbf{z}^* \in \mathrm{Prox}_{\lambda f}(\mathbf{u})$,*

$$\Psi_f(\mathbf{u}+h\mathbf{v}, \lambda) - \Psi_f(\mathbf{u}, \lambda) \le \frac{h}{\lambda}\langle \mathbf{u}-\mathbf{z}^*, \mathbf{v}\rangle + \frac{h^2}{2\lambda}\|\mathbf{v}\|^2 \tag{19}$$

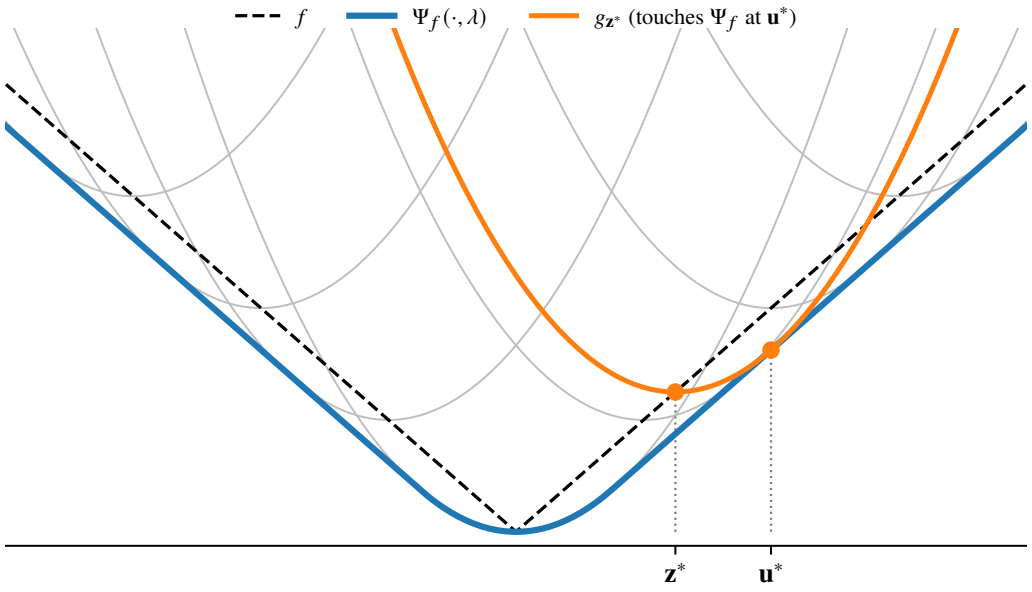

Figure 6: The Moreau envelope $\Psi_f(\cdot,\lambda)$ (blue) of f (dashed) as the lower envelope of the paraboloid family $\{g_{\mathbf{z}}\}_{\mathbf{z}}$ (gray). Each paraboloid has vertex $\mathbf{z}$, height $f(\mathbf{z})$, and curvature $1/\lambda$. At $\mathbf{u}^*$, the unique touching paraboloid $g_{\mathbf{z}^*}$ (orange) is indexed by $\mathbf{z}^* = \mathrm{prox}_{\lambda f}(\mathbf{u}^*)$. Since $g_{\mathbf{z}^*}$ is tangent to Ψ_f at $\mathbf{u}^*$, the two functions share a slope there: $\nabla\Psi_f(\mathbf{u}^*,\lambda) = (\mathbf{u}^* - \mathbf{z}^*)/\lambda$.

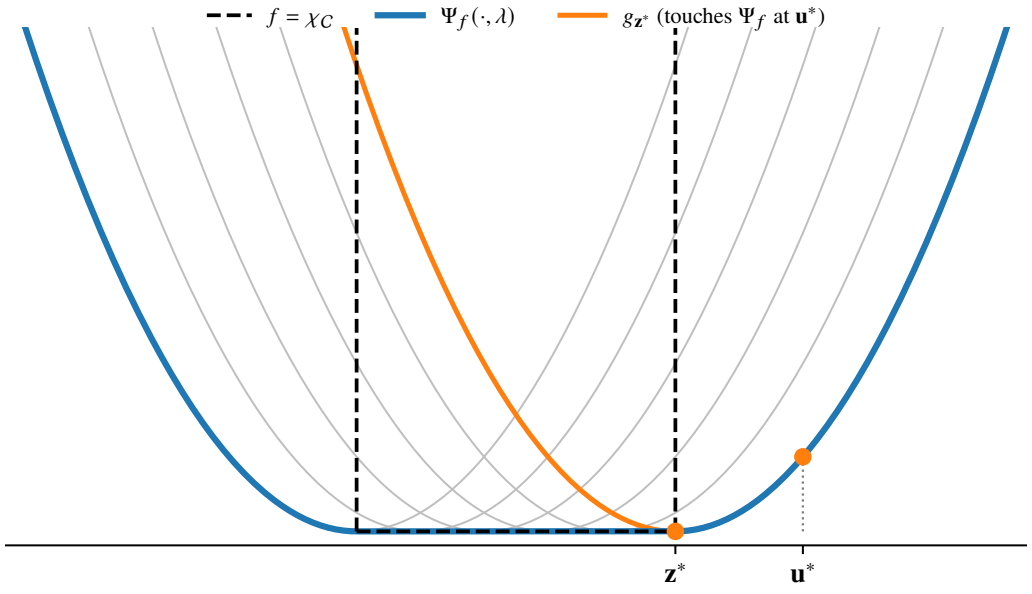

Figure 7: The Moreau envelope $\Psi_f(\cdot,\lambda)$ (blue) of a characteristic function $f = \chi_{\mathcal{C}}$ (dashed), which is zero on the nonempty closed constraint set $\mathcal{C}$ and $+\infty$ elsewhere. Only vertices $\mathbf{z} \in \mathcal{C}$ yield finite paraboloids $g_{\mathbf{z}}$ (gray), so for $\mathbf{u}^*$ outside $\mathcal{C}$ the touching paraboloid $g_{\mathbf{z}^*}$ (orange) is indexed by the projection $\mathbf{z}^* = \mathrm{proj}_{\mathcal{C}}(\mathbf{u}^*)$. $g_{\mathbf{z}^*}$ has height zero, so $\Psi_f(\mathbf{u}^*,\lambda) = 1/(2\lambda)\|\mathbf{u}^* - \mathbf{z}^*\|^2$, the squared distance to $\mathcal{C}$ scaled by $1/(2\lambda)$.

Proof. The paraboloid $g_{\mathbf{z}^*}$ bounds the envelope from above and equals it at $\mathbf{u}$. Hence,

$$\Psi_f(\mathbf{u}+h\mathbf{v},\lambda)-\Psi_f(\mathbf{u},\lambda)\le g_{\mathbf{z}^*}(\mathbf{u}+h\mathbf{v})-g_{\mathbf{z}^*}(\mathbf{u}) \tag{20}$$

The quadratic expension of the right-hand side gives Eq. 19. □

Eq. 19 is one-sided: it bounds Ψ_f from above by a touching paraboloid. This is enough to derive both the unique proximal point and the gradients at points of differentiability.

Corollary B.2 (Unique proximal point and envelope gradient)**.** *Assume C.6, and fix $\lambda > 0$ and $\mathbf{u}^* \in \mathbb{R}^d$. If $\Psi_f(\cdot,\lambda)$ is differentiable at $\mathbf{u}^*$, then $\mathrm{Prox}_{\lambda f}(\mathbf{u}^*)$ contains a unique point $\mathrm{prox}_{\lambda f}(\mathbf{u}^*)$, and*

$$\nabla\Psi_f(\mathbf{u}^*,\lambda)=\frac{1}{\lambda}\left(\mathbf{u}^*-\mathrm{prox}_{\lambda f}(\mathbf{u}^*)\right) \tag{21}$$

Proof. The proximal set is nonempty under Assumption C.6 (Rockafellar & Wets, 1998, Theorem 1.25). Choose any $\mathbf{z}^* \in \mathrm{Prox}_{\lambda f}(\mathbf{u}^*)$. By Lemma B.1, the gap between $g_{\mathbf{z}^*}$ and the envelope is nonnegative everywhere and zero at $\mathbf{u}^*$. Since the gap is minimized at $\mathbf{u}^*$ and differentiable there, its gradient vanishes:

$$\nabla\Psi_f(\mathbf{u}^*,\lambda)=\nabla g_{\mathbf{z}^*}(\mathbf{u}^*)=\frac{\mathbf{u}^*-\mathbf{z}^*}{\lambda} \tag{22}$$

Thus, every proximal point equals $\mathbf{u}^*-\lambda\nabla\Psi_f(\mathbf{u}^*,\lambda)$, so the proximal point is unique. We can therefore write $\mathbf{z}^*=\mathrm{prox}_{\lambda f}(\mathbf{u}^*)$, which yields Eq. 21. □

Critical points. The gradient formula in Eq. 21 also explains Property 3 from Section 4.1: a critical point of the envelope is a critical point of f, provided both functions are differentiable there.

Proposition B.3 (Critical points of the envelope)**.** *Suppose f and $\Psi_f(\cdot,\lambda)$ are differentiable at $\mathbf{u}^*$. If $\nabla\Psi_f(\mathbf{u}^*,\lambda)=\mathbf{0}$, then $\nabla f(\mathbf{u}^*)=\mathbf{0}$.*

Proof. By Corollary B.2, $\nabla\Psi_f(\mathbf{u}^*,\lambda)=\mathbf{0}$ implies $\mathrm{prox}_{\lambda f}(\mathbf{u}^*)=\mathbf{u}^*$. Thus, $\mathbf{z}=\mathbf{u}^*$ minimizes

$$g_{\mathbf{z}}(\mathbf{u}^*)=f(\mathbf{z})+\frac{1}{2\lambda}\|\mathbf{u}^*-\mathbf{z}\|^2 \tag{23}$$

At $\mathbf{z}=\mathbf{u}^*$, the quadratic penalty has zero gradient, so first-order optimality reduces to

$$\mathbf{0}=\nabla_{\mathbf{z}}g_{\mathbf{z}}(\mathbf{u}^*)|_{\mathbf{z}=\mathbf{u}^*}=\nabla f(\mathbf{u}^*) \tag{24}$$

□

B.3.2 On-Manifold Guidance

Below, we express the forward process $\mathbf{x}_t=\alpha_t\mathbf{x}_0+\sigma_t\epsilon$ in normalized coordinates $\mathbf{u}_t=\mathbf{x}_t/\alpha_t=\mathbf{x}_0+\rho_t\epsilon$, where $\rho_t=\sigma_t/\alpha_t$ is the effective noise level.

Jacobian as tangent projection (locally linear manifold). Manifold Constrained Gradient (MCG) (Chung et al., 2022) examines how the denoised estimate $\hat{\mathbf{x}}_0$ changes when the noisy sample $\mathbf{x}_t$ is perturbed. Under a locally linear manifold $\mathcal{M}$ and uniform data density on $\mathcal{M}$, $\hat{\mathbf{x}}_0$ is the nearest point on $\mathcal{M}$ to $\mathbf{u}_t=\mathbf{x}_t/\alpha_t$. Consider a perturbation $\mathbf{h}=\mathbf{h}_{\|}+\mathbf{h}_{\perp}$, decomposed into components tangent and normal to $\mathcal{M}$ at $\hat{\mathbf{x}}_0$. For sufficiently small ε,

$$\hat{\mathbf{x}}_0(\mathbf{x}_t+\varepsilon\mathbf{h},t)=\hat{\mathbf{x}}_0(\mathbf{x}_t,t)+\frac{\varepsilon}{\alpha_t}\mathbf{h}_{\|} \tag{25}$$

Normal perturbations leave $\hat{\mathbf{x}}_0$ unchanged, so only $\mathbf{h}_{\|}$ affects the denoised estimate. The denoiser Jacobian measures how $\hat{\mathbf{x}}_0$ changes under small perturbations, so we differentiate Eq. 25 at $\varepsilon=0$ to obtain

$$\mathbf{J}_{\hat{\mathbf{x}}_0}\mathbf{h}=\frac{1}{\alpha_t}\mathbf{h}_{\|} \tag{26}$$

Since $\mathbf{h}_{\|}=\Pi_{\mathcal{T}_{\hat{\mathbf{x}}_0}\mathcal{M}}\mathbf{h}$ and $\mathbf{h}$ is arbitrary, the denoiser Jacobian equals the scaled tangent projector:

$$\mathbf{J}_{\hat{\mathbf{x}}_0}=\frac{1}{\alpha_t}\Pi_{\mathcal{T}_{\hat{\mathbf{x}}_0}\mathcal{M}} \tag{27}$$

For MCG's measurement loss $\ell(\mathbf{x}_0) = \|\mathbf{W}(\mathbf{y} - \mathbf{H}\mathbf{x}_0)\|^2$, the chain rule gives

$$\nabla_{\mathbf{x}_t}\ell(\hat{\mathbf{x}}_0) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \nabla_{\mathbf{x}_0}\ell(\hat{\mathbf{x}}_0) = \frac{1}{\alpha_t}\Pi_{\mathcal{T}_{\hat{\mathbf{x}}_0}\mathcal{M}}\nabla_{\mathbf{x}_0}\ell(\hat{\mathbf{x}}_0) \tag{28}$$

Thus, the denoiser Jacobian projects the measurement gradient onto the tangent space (up to the scalar $1/\alpha_t$).

Jacobian as tangent projection (curved manifold). On a curved manifold with nonuniform density, $\hat{\mathbf{x}}_0$ generally differs from the nearest point $\pi(\mathbf{u}_t)$ at finite noise. Thus, the nearest-point projection no longer determines $\mathbf{J}_{\hat{\mathbf{x}}_0}$. We instead characterize $\mathbf{J}_{\hat{\mathbf{x}}_0}$ through the posterior covariance Σ_t; for an exact score,

$$\alpha_t \mathbf{J}_{\hat{\mathbf{x}}_0} = \frac{\Sigma_t}{\rho_t^2}, \qquad \Sigma_t \coloneqq \mathrm{Cov}[\mathbf{x}_0 \mid \mathbf{x}_t] \tag{29}$$

The scaled denoiser Jacobian measures posterior uncertainty, normalized by the noies variance ρ_t^2.

Let $\mathcal{M}$ be a compact, embedded C^3 manifold without boundary, with a strictly positive C^3 data density. For a normalized noisy sample $\mathbf{u}_t$ near $\mathcal{M}$, the posterior $q_0(\mathbf{x}_0|\mathbf{x}_t)$ concentrates around $\pi(\mathbf{u}_t)$ in the small-noise limit. A local Laplace expansion gives

$$\alpha_t \mathbf{J}_{\hat{\mathbf{x}}_0} = \frac{\Sigma_t}{\rho_t^2} = D\pi(\mathbf{u}_t) + \mathcal{O}\big(\rho_t \|\log \rho_t\|^3\big) \tag{30}$$

where $D\pi$ is the derivative of the nearest-point projection. The bound is uniform in any fixed tube around $\mathcal{M}$ with radius below the reach of $\mathcal{M}$. (Kharitenko et al., 2026, Theorem 1).

On $\mathcal{M}$, $D\pi(\mathbf{u}_t)$ equals the orthogonal tangent projector. Away from $\mathcal{M}$, $D\pi(\mathbf{u}_t)$ still removes normal components at $\pi(\mathbf{u}_t)$, but curvature can rescale the tangential components. Thus, the difference between $D\pi(\mathbf{u}_t)$ and the tangent projector is $\mathcal{O}(\mathrm{dist}(\mathbf{u}_t, \mathcal{M}))$ (Kharitenko et al., 2026, Corollary 2 and Lemma 11). The noise and curvature errors both vanish in the limit, so $\alpha_t \mathbf{J}_{\hat{\mathbf{x}}_0}$ approaches the tangent projector.

B.4 Practical Design Choices

B.4.1 Stabilizing Proximal Guidance

The annealing schedule λ_t introduces numerical instabilities to ProxDG. Inside the guidance vector $\Delta_t = (1/\lambda_t)(\hat{\mathbf{x}}_0 - \mathrm{prox}_{\lambda_t f}(\hat{\mathbf{x}}_0))$, the scale $1/\lambda_t$ diverges as $\lambda_t \to 0$. Inside the proximal objective, the regularization coefficient $1/(2\lambda_t)$ also diverges, a problem encountered in other proximal guidance methods (Zhu et al., 2023; Rout et al., 2025; Narasimhan et al., 2025; Christopher et al., 2026). We stabilize each instability in turn.

We address the first instability by defining λ_t to cancel the $1/\lambda_t$ scaling in Δ_t. Standard guidance methods scale the guidance drift by σ_t^2 (Ho et al., 2020; 2022; Peng et al., 2024), so the full ProxDG guidance term scales as σ_t^2/λ_t. We make this ratio constant in t by setting

$$\lambda_t \coloneqq \beta\sigma_t^2 \tag{31}$$

with $\beta > 0$ a hyperparameter. This reparameterization keeps the guidance from vanishing artificially as $t \to 0$, which improves optimality with respect to f (Yu et al., 2023; Ye et al., 2024; Zhang et al., 2025c). From a theoretical perspective, the ideal λ_t matches the posterior covariance $\mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t]$ of the prior given the noisy sample (see Appendix B.4.2), which can be expensive to approximate (Peng et al., 2024; Rissanen et al., 2025). The choice $\lambda_t \propto \sigma_t^2$ is a tractable alternative, validated empirically by DiffPIR (Zhu et al., 2023).

We address the second instability by rescaling the proximal objective:

$$\mathrm{prox}_{\lambda_t f}(\hat{\mathbf{x}}_0) \coloneqq \arg\min_{\mathbf{z}_t \in \mathbb{R}^d} \left\{ \lambda_t f(\mathbf{z}_t) + \frac{1}{2}\|\hat{\mathbf{x}}_0 - \mathbf{z}_t\|^2 \right\} \tag{32}$$

This keeps the proximal point invariant, since positive rescaling does not change the optimum. Under $\lambda_t = \beta\sigma_t^2$, the scaling λ_t on f shrinks to zero at the final denoising steps, so the proximal point collapses to $\hat{\mathbf{x}}_0$ and the guidance vanishes. This is correct behavior for posterior sampling, but sacrifices optimality with respect to f. For tasks concerned with raw optimality, we floor λ_t at a positive $\lambda_{\min}$ past a threshold timestep $t_{\min}$, keeping guidance active through the end of sampling.

B.4.2 ANNEALING SCHEDULE

To derive the optimal annealing schedule, we first lift the scalar λ_t to an anisotropic kernel Λ_t. We then derive Λ_t that minimizes the approximation error against ideal diffusion guidance. The optimal Λ_t turns out to be the posterior covariance $\mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t]$, and the scalar λ_t is its isotropic restriction.

Anisotropic ProxDG. We replace the isotropic kernel $\lambda_t \mathbf{I}$ with a positive definite matrix Λ_t, allowing the proximal penalty to weight displacements differently across directions. The proximal point then solves

$$\mathbf{z}_t^\Lambda \coloneqq \arg\min_{\mathbf{z}\in\mathbb{R}^d} \left\{ f(\mathbf{z}) + \frac{1}{2}(\mathbf{z}-\hat{\mathbf{x}}_0)^\top \Lambda^{-1}(\mathbf{z}-\hat{\mathbf{x}}_0) \right\} \tag{33}$$

The minimum value of this objective, viewed as a function of $\hat{\mathbf{x}}_0$, defines a Bregman–Moreau envelope of f (Bauschke et al., 2018; Lau & Liu, 2022).

$$\Psi_f(\mathbf{u}, \Lambda_t) \coloneqq \inf_{\mathbf{z}\in\mathbb{R}^d} \left\{ f(\mathbf{z}) + \frac{1}{2}(\mathbf{z}-\mathbf{u})^\top \Lambda_t^{-1}(\mathbf{z}-\mathbf{u}) \right\} \tag{34}$$

Where differentiable, the envelope has gradient $\Lambda_t^{-1}(\hat{\mathbf{x}}_0 - \mathbf{z}_t^\Lambda)$ with respect to $\hat{\mathbf{x}}_0$ (Lau & Liu, 2022, Proposition 2). We define our smoothed guidance objective by $f_t^\Lambda(\mathbf{x}) = \Psi_f(\hat{\mathbf{x}}_0(\mathbf{x}, t), \Lambda_t)$. The chain rule then gives

$$\nabla f_t^\Lambda(\mathbf{x}) = \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)^\top \Lambda^{-1} \left(\hat{\mathbf{x}}_0 - \mathbf{z}_t^\Lambda \right) \tag{35}$$

We employ ∇f_t^Λ in the Langevin updates of Algorithm 1, recovering ProxDG when $\Lambda_t = \lambda_t \mathbf{I}$. The remaining task is to choose Λ_t.

Ideal guidance and its Gaussian surrogate. By Peng et al. (2024, Proposition 3.1), the ideal guidance is

$$\sigma_t^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t) = \alpha_t \left(\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}] - \hat{\mathbf{x}}_0 \right) \tag{36}$$

The conditional mean $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}]$ is generally intractable. To approximate this mean, we first replace $p(\mathbf{x}_0|\mathbf{x}_t)$ with a Gaussian that matches its mean and covariance (Boys et al., 2024, Section 3.1):

$$\hat{q}_0(\mathbf{x}_0|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_0; \hat{\mathbf{x}}_0, \Sigma_t), \qquad \Sigma_t \coloneqq \mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t] \tag{37}$$

The second-order Tweedie identity expresses the covariance in terms of the exact denoiser Jacobian:

$$\Sigma_t = \frac{\sigma_t^2}{\alpha_t} \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}_t, t) \tag{38}$$

We then condition the Gaussian approximation on $\mathbf{y}$, yielding the surrogate posterior

$$\hat{q}_0^\Sigma(\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}) \propto \exp(-f(\mathbf{x}_0)) \mathcal{N}(\mathbf{x}_0; \hat{\mathbf{x}}_0, \Sigma_t) \tag{39}$$

The surrogate mean $\overline{\mathbf{x}}_0^\Sigma$ approximates $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}]$. The surrogate mode $\mathbf{z}_t^\Sigma$ minimizes the negative log-density, which is exactly the proximal objective in Eq. 33 with $\Lambda_t = \Sigma_t$.

Matching the two drifts. We use the covariance identity $\mathbf{J}_{\hat{\mathbf{x}}_0} = (\alpha_t/\sigma_t^2)\Sigma_t$ to express the anisotropic guidance in Eq. 35 as

$$-\sigma_t^2 \nabla f_t^\Lambda(\mathbf{x}_t) = \alpha_t \Sigma_t \Lambda_t^{-1} (\mathbf{z}_t^\Lambda - \hat{\mathbf{x}}_0) \tag{40}$$

We approximate the ideal guidance by replacing its conditional mean $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}]$ with the surrogate mean $\overline{\mathbf{x}}_0^\Sigma$:

$$\sigma_t^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y} \mid \mathbf{x}_t) \approx \alpha_t (\overline{\mathbf{x}}_0^\Sigma - \hat{\mathbf{x}}_0) \tag{41}$$

The anisotropic and surrogate guidance share the prefactor α_t. We compare the remaining guidance terms through the *approximation gap*

$$\begin{aligned} \epsilon(\Lambda_t) &\coloneqq \Sigma_t \Lambda_t^{-1} (\mathbf{z}_t^\Lambda - \hat{\mathbf{x}}_0) - (\overline{\mathbf{x}}_0^\Sigma - \hat{\mathbf{x}}_0) \\ &= \underbrace{\Sigma_t \Lambda_t^{-1} (\mathbf{z}_t^\Lambda - \hat{\mathbf{x}}_0) - (\mathbf{z}_t^\Sigma - \hat{\mathbf{x}}_0)}_{\text{covariance mismatch}} + \underbrace{\mathbf{z}_t^\Sigma - \overline{\mathbf{x}}_0^\Sigma}_{\text{mode-mean mismatch}} \end{aligned} \tag{42}$$

The gap is zero precisely when the anisotropic guidance matches the surrogate guidance. With $\Lambda_t = \Sigma_t$, we have $\Sigma_t \Lambda_t^{-1} = \mathbf{I}$ and $\mathbf{z}_t^\Lambda = \mathbf{z}_t^\Sigma$, so the covariance mismatch vanishes. The gap then reduces to the difference between the surrogate mode and mean. For convex quadratic f, the surrogate posterior $\hat{q}_0^\Sigma$ is Gaussian and has equal mode and mean, so $\epsilon(\Sigma_t) = \mathbf{0}$.

B.5 Connections to Prior Work

B.5.1 DPS

ProxDG recovers DPS when the proximal solver is a single gradient step initialized at $\hat{\mathbf{x}}_0 = \hat{\mathbf{x}}_0(\mathbf{x}_t, t)$. In the rescaled proximal form,

$$\mathbf{z}_t^* \in \arg\min_{\mathbf{z}\in\mathbb{R}^d} \left\{ \lambda_t f(\mathbf{z}) + \frac{1}{2}\|\hat{\mathbf{x}}_0 - \mathbf{z}\|^2 \right\} \tag{43}$$

the ProxDG guidance vector is

$$\nabla f_t^\lambda(\mathbf{x}) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \Delta_t, \qquad \Delta_t \coloneqq \frac{1}{\lambda_t}(\hat{\mathbf{x}}_0 - \mathbf{z}_t^*) \tag{44}$$

We approximate $\mathbf{z}_t^*$ by a gradient step on the proximal objective, starting from $\mathbf{z}_t^{(0)} = \hat{\mathbf{x}}_0$. Since

$$\nabla \left(\lambda_t f(\mathbf{z}) + \frac{1}{2}\|\hat{\mathbf{x}}_0 - \mathbf{z}\|^2 \right) = \lambda_t \nabla f(\mathbf{z}) + (\mathbf{z} - \hat{\mathbf{x}}_0) \tag{45}$$

and the quadratic term has zero gradient at $\mathbf{z} = \hat{\mathbf{x}}_0$, one gradient step yields

$$\mathbf{z}_t^{(1)} = \hat{\mathbf{x}}_0 - \lambda_t \nabla f(\hat{\mathbf{x}}_0) \tag{46}$$

Substituting $\mathbf{z}_t^{(1)}$ into the proximal displacement,

$$\Delta_t^{(1)} = \frac{1}{\lambda_t}\left(\hat{\mathbf{x}}_0 - \mathbf{z}_t^{(1)}\right) = \nabla f(\hat{\mathbf{x}}_0) \tag{47}$$

Therefore the ProxDG guidance vector becomes

$$\mathbf{J}_{\hat{\mathbf{x}}_0}^\top \nabla f(\hat{\mathbf{x}}_0) \tag{48}$$

which is exactly the DPS guidance vector.

B.5.2 LGD & TFG

Recall that ProxDG uses the guidance function

$$f_t^\lambda(\mathbf{x}_t) = \inf_{\mathbf{z}_t\in\mathbb{R}^d} \left\{ f(\mathbf{z}_t) + \frac{1}{2\lambda_t}\|\mathbf{z}_t - \hat{\mathbf{x}}_0\|_2^2 \right\} \tag{49}$$

and the guidance vector

$$\nabla f_t^\lambda(\mathbf{x}_t) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \frac{1}{\lambda_t}\left(\hat{\mathbf{x}}_0 - \mathrm{prox}_{\lambda_t f}\left(\hat{\mathbf{x}}_0\right)\right) \tag{50}$$

HJ-Prox solver and the finite-temperature potential. HJ-Prox (Osher et al., 2023) approximates the proximal point using a Gaussian soft average. For temperature $\upsilon > 0$ and $\mathbf{u} = \hat{\mathbf{x}}_0(\mathbf{x}_t, t)$, define the soft average

$$\overline{\mathbf{z}}_{\upsilon,t}(\mathbf{u}) \coloneqq \frac{\int_{\mathbb{R}^d} \mathbf{z} \exp\left(-f(\mathbf{z})/\upsilon\right) \mathcal{N}(\mathbf{z}; \mathbf{u}, \upsilon\lambda_t \mathbf{I})\, d\mathbf{z}}{\int_{\mathbb{R}^d} \exp(-f(\mathbf{z})/\upsilon) \mathcal{N}(\mathbf{z}; \mathbf{u}, \upsilon\lambda_t \mathbf{I})\, d\mathbf{z}} \tag{51}$$

together with the temperature-scale potential

$$\mathcal{F}_{\upsilon,t}(\mathbf{x}_t) \coloneqq -\upsilon \log \int_{\mathbb{R}^d} \exp(-f(\mathbf{z})/\upsilon) \mathcal{N}\left(\mathbf{z}; \mathbf{u}, \upsilon\lambda_t \mathbf{I}\right)\, d\mathbf{z} \tag{52}$$

We differentiate the potential through the denoiser to obtain

$$\nabla \mathcal{F}_{\upsilon,t}(\mathbf{x}_t) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \frac{1}{\lambda_t}\left(\mathbf{u} - \overline{\mathbf{z}}_{\upsilon,t}(\mathbf{u})\right) \tag{53}$$

Eq. 53 has the form of ProxDG's guidance, with $\overline{\mathbf{z}}_{\upsilon,t}$ approximating the proximal point.

Connection to LGD and TFG. At $\upsilon = 1$, Eq. 52 becomes

$$\mathcal{F}_{1,t}(\mathbf{x}_t) = -\log \int_{\mathbb{R}^d} \exp(-f(\mathbf{z})) \mathcal{N}(\mathbf{z}; \mathbf{u}, \lambda_t \mathbf{I})\, d\mathbf{z} \tag{54}$$

$\mathcal{F}_{1,t}$ approximates the ideal guidance function by replacing $q_0(\mathbf{x}_0|\mathbf{x}_t)$ with the Gaussian $\mathcal{N}(\mathbf{u}, \lambda_t \mathbf{I})$, centered at $\mathbf{u} = \hat{\mathbf{x}}_0$. LGD (Song et al., 2023b), and TFG (Ye et al., 2024; Zhang et al., 2025c) use the same Gaussian approximation but choose different variances. At $\upsilon = 1$, ProxDG with HJ-Prox therefore recovers the LGD and TFG objectives when the variances match.

Note: At $\upsilon = 1$, ProxDG with HJ-Prox, LGD, and TFG share the guidance function $\mathcal{F}_{1,t}$ but compute its gradient differently. For $\mathbf{z} \sim \mathcal{N}(\mathbf{u}, \lambda_t \mathbf{I})$, define the normalized weight

$$w_t(\mathbf{z}; \mathbf{u}) \coloneqq \frac{\exp(-f(\mathbf{z}))}{\mathbb{E}_{\mathbf{z}' \sim \mathcal{N}(\mathbf{u}, \lambda_t \mathbf{I})}[\exp(-f(\mathbf{z}'))]} \tag{55}$$

LGD and TFG compute the weighted gradient

$$\mathbf{g}_{\mathrm{LGD/TFG}}(\mathbf{u}) = \mathbb{E}\left[w_t(\mathbf{z}; \mathbf{u}) \nabla f(\mathbf{z})\right] \tag{56}$$

HJ-Prox instead computes $\overline{\mathbf{z}}_{1,t}(\mathbf{u})$ to form a proximal displacement:

$$\overline{\mathbf{z}}_{1,t}(\mathbf{u}) = \mathbb{E}\left[w_t(\mathbf{z}; \mathbf{u}) \mathbf{z}\right], \qquad \mathbf{g}_{\mathrm{HJ}}(\mathbf{u}) = \frac{\mathbf{u} - \overline{\mathbf{z}}_{1,t}(\mathbf{u})}{\lambda_t} \tag{57}$$

For sufficiently regular f, Stein's lemma relates these expressions:

$$\mathbb{E}\left[w_t(\mathbf{z}; \mathbf{u})(\mathbf{z} - \mathbf{u})\right] = -\lambda_t \mathbb{E}\left[w_t(\mathbf{z}; \mathbf{u}) \nabla f(\mathbf{z})\right] \tag{58}$$

Thus, $\mathbf{g}_{\mathrm{HJ}} = \mathbf{g}_{\mathrm{LGD/TFG}}$ at the population level, and the chain rule gives

$$\nabla_{\mathbf{x}_t} \mathcal{F}_{1,t}(\mathbf{x}_t) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \mathbf{g}_{\mathrm{HJ}}(\mathbf{u}) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \mathbf{g}_{\mathrm{LGD/TFG}}(\mathbf{u}), \qquad \mathbf{u} = \hat{\mathbf{x}}_0(\mathbf{x}_t, t) \tag{59}$$

The finite-sample estimators may not agree in expectation.

Zero-temperature connection. As $\upsilon \downarrow 0$, $\mathcal{F}_{\upsilon,t}$ converges to the Moreau envelope (Habring et al., 2025),

$$\lim_{\upsilon \downarrow 0} \mathcal{F}_{\upsilon,t}(\mathbf{x}_t) = f_t^{\lambda}(\mathbf{x}_t) \tag{60}$$

Thus, ProxDG is the zero-temperature limit of potentials $\mathcal{F}_{\upsilon,t}$, which recover LGD and TFG at $\upsilon = 1$.

B.5.3 ΠGDM

In the linear-Gaussian setting, ProxDG admits a closed-form proximal solve and reduces to ΠGDM (Song et al., 2023a). Consider the measurement model

$$\mathbf{y} = \mathbf{H}\mathbf{x}_0 + \epsilon, \qquad \epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) \tag{61}$$

where $\mathbf{H} \in \mathbb{R}^{m \times d}$ and the known measurement noise is $\sigma^2 > 0$. The data-fit objective becomes the convex quadratic

$$f(\mathbf{x}) = \frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 \tag{62}$$

Because f is quadratic, our proximal point has a closed form solution. Anchored at $\hat{\mathbf{x}}_0$,

$$\mathbf{z}_t^* = \mathrm{prox}_{\lambda_t f}(\hat{\mathbf{x}}_0) = \arg\min_{\mathbf{z}_t \in \mathbb{R}^d} \left\{ \frac{\lambda_t}{2\sigma^2} \|\mathbf{y} - \mathbf{H}\mathbf{z}_t\|^2 + \frac{1}{2} \|\hat{\mathbf{x}}_0 - \mathbf{z}_t\|^2 \right\} \tag{63}$$

First-order optimality yields the unique minimizer

$$\mathbf{z}_t^* = \left(\mathbf{I} + \frac{\lambda_t}{\sigma^2} \mathbf{H}^\top \mathbf{H} \right)^{-1} \left(\hat{\mathbf{x}}_0 + \frac{\lambda_t}{\sigma^2} \mathbf{H}^\top \mathbf{y} \right) \tag{64}$$

so the ProxDG displacement $\Delta_t \coloneqq (1/\lambda_t)(\hat{\mathbf{x}}_0 - \mathbf{z}_t^*)$ simplifies to

$$\Delta_t = \mathbf{H}^\top (\lambda_t \mathbf{H}\mathbf{H}^\top + \sigma^2 \mathbf{I})^{-1} (\mathbf{H}\hat{\mathbf{x}}_0 - \mathbf{y}) \tag{65}$$

Applying the chain rule through Tweedie's estimator yields

$$\nabla f_t^{\lambda}(\mathbf{x}_t) = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \mathbf{H}^\top \left(\lambda_t \mathbf{H}\mathbf{H}^\top + \sigma^2 \mathbf{I} \right)^{-1} (\mathbf{H}\hat{\mathbf{x}}_0 - \mathbf{y}) \tag{66}$$

Thus, ProxDG recovers ΠGDM's pseudoinverse correction up to sign (with appropriate choice of λ_t).

C f_t^λ PROPERTIES

C.1 ASSUMPTIONS

Assumption C.1. $\lambda : (0, T] \to (0, \infty)$ is a continuously differentiable function defined by $t \mapsto \lambda_t$. Furthermore, $\lim_{t\to 0} \lambda_t = 0$ and $\lim_{t\to 0} \frac{\sigma_t^4}{\lambda_t} = 0$.

Assumption C.2. $\alpha, \sigma : [0, T] \to \mathbb{R}$ are continuous functions with $\alpha_t > 0$ for all $t \in [0, T]$, $\alpha_0 = 1$, and $\sigma_0 = 0$.

Assumption C.3. $|1 - \alpha_t|$ decays at least as fast as σ_t^2, i.e. $|1 - \alpha_t| \leq K\sigma_t^2$ for some constant K.

Assumption C.4. The score estimator s_θ is continuous on $\mathbb{R}^d \times [0, T]$.

Assumption C.5. For each $t \in [0, T]$, the score estimator s_θ is locally Lipschitz on $\mathbb{R}^d$. That is, for each compact $K \subset \mathbb{R}^d$, there exists $L_{K,t} < \infty$ such that

$$\|s_\theta(\mathbf{x}, t) - s_\theta(\mathbf{x}', t)\| \leq L_{K,t}\|\mathbf{x} - \mathbf{x}'\|, \qquad \mathbf{x}, \mathbf{x}' \in K \tag{67}$$

Assumption C.6. $f : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ is proper, lower semi-continuous, and bounded from below. That is, there exists a constant $C_f \in \mathbb{R}$ such that $f(\mathbf{x}) \geq C_f$ for all $\mathbf{x} \in \mathbb{R}^d$

In Section 4.2, we define λ_t to satisfy Assumption C.1. Assumption C.2 ensures that the Tweedie's estimator $\hat{\mathbf{x}}_0(\mathbf{x}_t, t)$ is well-defined at all times t. In practice, standard implementations (von Platen et al., 2022; Dhariwal & Nichol, 2021; Nichol & Dhariwal, 2021) directly enforce $\alpha_t > 0$ as $t \to T$ by clipping α_t to a positive lower bound. Assumption C.3 states that $|1 - \alpha_t|$ decays at least as fast as σ_t^2. Remark C.7 shows that popular noise schedules satisfy this decay property. Assumption C.4 holds when s_θ is directly parameterized by a network with continuous components. Assumption C.5 is satisfied by standard score-network architectures, including U-Nets (Ronneberger et al., 2015) and transformers (Vaswani et al., 2017), whose components are locally Lipschitz on compact sets (Virmaux & Scaman, 2018; Kim et al., 2021). Finally, Assumption C.6 is very permissive: it allows nonconvex and discontinuous objectives, provided f is bounded below. For $f = \chi_{\mathcal{C}}$, Assumption C.6 holds exactly when $\mathcal{C}$ is nonempty and closed.

Remark C.7. Under VP-, VE- (Song et al., 2021b), and EDM (Karras et al., 2022) noise schedules, we have that $|1 - \alpha_t| \leq K\sigma_t^2$ for some constant $K \geq 0$.

Proof. **VP Schedules:** In the VP formulation,

$$\alpha_t^2 + \sigma_t^2 = 1 \tag{68}$$

Since $\alpha_t > 0$, Eq. 68 gives $\alpha_t \leq 1$ and $\alpha_t = \sqrt{1 - \sigma_t^2}$. Therefore,

$$|1 - \alpha_t| = 1 - \alpha_t \tag{69}$$

$$= 1 - \sqrt{1 - \sigma_t^2} \tag{70}$$

We multiply the right-hand side by $\frac{1+\sqrt{1-\sigma_t^2}}{1+\sqrt{1-\sigma_t^2}}$ to obtain

$$|1 - \alpha_t| = \frac{\left(1 - \sqrt{1 - \sigma_t^2}\right)\left(1 + \sqrt{1 - \sigma_t^2}\right)}{1 + \sqrt{1 - \sigma_t^2}} \tag{71}$$

$$= \frac{1 - (1 - \sigma_t^2)}{1 + \sqrt{1 - \sigma_t^2}} \tag{72}$$

$$= \frac{\sigma_t^2}{1 + \sqrt{1 - \sigma_t^2}} \tag{73}$$

Since the denominator is at least 1,

$$\frac{\sigma_t^2}{1 + \sqrt{1 - \sigma_t^2}} \leq \sigma_t^2 \implies |1 - \alpha_t| \leq \sigma_t^2 \tag{74}$$

Thus, $|1 - \alpha_t| \leq K\sigma_t^2$ for $K = 1$.

VE and EDM (Karras et al., 2022) Schedules: For VE and EDM schedules, $\alpha_t = 1$ for all $t \in [0, T]$, so $|1 - \alpha_t| = 0$. Therefore, $|1 - \alpha_t| \leq K\sigma_t^2$ holds with $K = 0$. □

C.2 Finiteness and Continuity of f_t^λ

Lemma C.8. *Assume C.2 and C.4. Then the map $(\mathbf{x}, t) \mapsto \hat{\mathbf{x}}_0(\mathbf{x}, t)$ is continuous on $\mathbb{R}^d \times [0, T]$, with $\hat{\mathbf{x}}_0(\mathbf{x}, 0) = \mathbf{x}$.*

Proof. Tweedie's formula gives

$$\hat{\mathbf{x}}_0(\mathbf{x}, t) = \frac{1}{\alpha_t}\mathbf{x} + \frac{\sigma_t^2}{\alpha_t} s_\theta(\mathbf{x}, t) \tag{75}$$

By Assumption C.2, α_t and σ_t are continuous on $[0, T]$ with $\alpha_t > 0$, so $1/\alpha_t$ and σ_t^2/α_t are continuous on $[0, T]$. By Assumption C.4, s_θ is also continuous on $\mathbb{R}^d \times [0, T]$. Therefore, $\hat{\mathbf{x}}_0$ is continuous on $\mathbb{R}^d \times [0, T]$. At $t = 0$, Assumption C.2 gives $\alpha_0 = 1$ and $\sigma_0 = 0$, so $\hat{\mathbf{x}}_0(\mathbf{x}, 0) = \mathbf{x}$. □

Proposition C.9. *Assume C.1, C.2, C.4, and C.6. Then the map $(\mathbf{x}, t) \mapsto f_t^\lambda(\mathbf{x})$ is finite and continuous on $\mathbb{R}^d \times (0, T]$.*

Proof. **Finiteness.** By Assumption C.6, $f \geq C_f$ and f is proper, so there exists $\mathbf{z}_0$ with $f(\mathbf{z}_0) < +\infty$. For every $(\mathbf{x}, t) \in \mathbb{R}^d \times (0, T]$,

$$C_f \leq f_t^\lambda(\mathbf{x}) \leq f(\mathbf{z}_0) + \frac{1}{2\lambda_t}\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{z}_0\|^2 < +\infty \tag{76}$$

The lower bound follows from the nonnegative quadratic penalty; the upper bound follows from evaluating the proximal objective at $\mathbf{z}_0$. Hence $f_t^\lambda(\mathbf{x})$ is finite on $\mathbb{R}^d \times (0, T]$.

Continuity. Since f is proper, lower semicontinuous, and bounded below, the Moreau envelope Ψ_f is jointly continuous on $\mathbb{R}^d \times (0, \infty)$ (Rockafellar & Wets, 1998, Exercise 1.24 and Theorem 1.25). The denoiser is jointly continuous by Lemma C.8, and the schedule $t \mapsto \lambda_t$ is continuous and positive by Assumption C.1. Therefore,

$$(\mathbf{x}, t) \mapsto \Psi_f\left(\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t\right) = f_t^\lambda(\mathbf{x}) \tag{77}$$

is continuous on $\mathbb{R}^d \times (0, T]$ as a composition of continuous maps. □

C.3 Pointwise Convergence of f_t^λ to f

Lemma C.10. *Assume C.1, C.2, C.3, and C.4. Then for every $\mathbf{x} \in \mathbb{R}^d$,*

$$\frac{1}{2\lambda_t}\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x}\|^2 \to 0 \tag{78}$$

as $t \to 0$.

Proof. Tweedie's formula gives

$$\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x} = \frac{1 - \alpha_t}{\alpha_t}\mathbf{x} + \frac{\sigma_t^2}{\alpha_t} s_\theta(\mathbf{x}, t) \tag{79}$$

By the triangle inequality and Assumptions C.2, C.3,

$$\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x}\| \leq \left|\frac{1 - \alpha_t}{\alpha_t}\right| \|\mathbf{x}\| + \frac{\sigma_t^2}{\alpha_t}\|s_\theta(\mathbf{x}, t)\| \tag{80}$$

$$\leq \frac{\sigma_t^2}{\alpha_t}\left(K\|\mathbf{x}\| + \|s_\theta(\mathbf{x}, t)\|\right) \tag{81}$$

$\|s_\theta(\mathbf{x}, t)\|$ is bounded near $t = 0$ by Assumption C.4, and $\alpha_t \to 1$ by Assumption C.2. Therefore, the right-hand side is $\mathcal{O}(\sigma_t^2)$. As a result,

$$\frac{1}{\lambda_t}\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x}\|^2 = \mathcal{O}\left(\frac{\sigma_t^4}{\lambda_t}\right) \tag{82}$$

which vanishes by Assumption C.1. □

Proposition C.11. *Assume C.1, C.2, C.3, C.4, and C.6. Then $f_t^\lambda(\mathbf{x}) \to f(\mathbf{x})$ for every $\mathbf{x} \in \mathbb{R}^d$ as $t \to 0$.*

Proof. Write $\mathbf{u}_t \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}, t)$, so $f_t^\lambda(\mathbf{x}) = \Psi_f(\mathbf{u}_t, \lambda_t)$. Lemma C.8 states $\mathbf{u}_t \to \mathbf{x}$, and Assumption C.1 states $\lambda_t \to 0$.

Upper bound. Since $\Psi_f(\mathbf{u}_t, \lambda_t)$ is an infimum, its objective evaluated at $\mathbf{x}$ yields

$$f_t^\lambda(\mathbf{x}) \le f(\mathbf{x}) + \frac{1}{2\lambda_t}\|\mathbf{u}_t - \mathbf{x}\|^2 \tag{83}$$

Lemma C.10 proves the quadratic penalty decays to zero, so $\limsup_{t\to 0} f_t^\lambda(\mathbf{x}) \le f(\mathbf{x})$.

Lower bound. As $\lambda \to 0$, the envelopes $\Psi_f(\cdot, \lambda)$ epi-converge to f (Rockafellar & Wets, 1998, Section 7B). In particular, for every $(\mathbf{u}_t, \lambda_t) \to (\mathbf{x}, 0)$ (Rockafellar & Wets, 1998, Proposition 7.2),

$$\liminf_{t\to 0} \Psi_f(\mathbf{u}_t, \lambda_t) \ge f(\mathbf{x}) \tag{84}$$

Since $\mathbf{u}_t \to \mathbf{x}$ by Lemma C.8, and $\lambda_t \to 0$ by Assumption C.1, $\liminf_{t\to 0} f_t^\lambda(\mathbf{x}) \ge f(\mathbf{x})$. □

C.4 Critical Points of f_t^λ

Proposition C.12. *Fix $t \in (0, T]$ and $\mathbf{x}^* \in \mathbb{R}^d$, and let $\mathbf{u}^* \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}^*, t)$. Assume:*

1. *f satisfies Assumption C.6 and is differentiable at $\mathbf{u}^*$.*
2. *$\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable at $\mathbf{x}^*$, with nonsingular Jacobian $\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}^*, t)$.*
3. *$\Psi_f(\cdot, \lambda_t)$ is differentiable at $\mathbf{u}^*$.*

If $\nabla f_t^\lambda(\mathbf{x}^*) = \mathbf{0}$, then $\nabla f(\mathbf{u}^*) = \mathbf{0}$.

Proof. The chain rule gives

$$\nabla f_t^\lambda(\mathbf{x}^*) = \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}^*, t)^\top \nabla_{\mathbf{u}} \Psi_f(\mathbf{u}^*, \lambda_t) \tag{85}$$

Since $\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}^*, t)$ is nonsingular, $\nabla f_t^\lambda(\mathbf{x}^*) = \mathbf{0}$ forces $\nabla_{\mathbf{u}} \Psi_f(\mathbf{u}^*, \lambda_t) = \mathbf{0}$.

By Proposition B.3, $\nabla_{\mathbf{u}} \Psi_f(\mathbf{u}^*, \lambda_t) = \mathbf{0}$ implies $\nabla f(\mathbf{u}^*) = \mathbf{0}$. □

C.5 Gradients of f_t^λ

Lemma C.13. *Assume C.1, C.2, C.5, and C.6, and fix $t \in (0, T]$. Then f_t^λ is locally Lipschitz on $\mathbb{R}^d$, and therefore differentiable almost everywhere.*

Proof. By Assumption C.5, $s_\theta(\cdot, t)$ is locally Lipschitz. Since $\alpha_t > 0$, the denoiser

$$\hat{\mathbf{x}}_0(\cdot, t) = \frac{1}{\alpha_t}\mathrm{Id} + \frac{\sigma_t^2}{\alpha_t} s_\theta(\cdot, t) \tag{86}$$

is also locally Lipschitz. By Assumption C.6 and Jourani et al. (2014, Proposition 3.3(b)), $\Psi_f(\cdot, \lambda_t)$ is locally Lipschitz. Therefore,

$$f_t^\lambda = \Psi_f(\cdot, \lambda_t) \circ \hat{\mathbf{x}}_0(\cdot, t) \tag{87}$$

is locally Lipschitz as a composition of locally Lipschitz maps, and hence differentiable almost everywhere by Rademacher's theorem. □

Proposition C.14. *Assume C.1, C.2, C.5, and C.6, and fix $t \in (0, T]$. At every $\mathbf{x} \in \mathbb{R}^d$ where both $\hat{\mathbf{x}}_0(\cdot, t)$ and f_t^λ are differentiable,*

$$\nabla f_t^\lambda(\mathbf{x}) = \frac{1}{\lambda_t} \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)^\top \left(\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{z}\right) \tag{88}$$

for every $\mathbf{z} \in \mathrm{Prox}_{\lambda_t f}(\hat{\mathbf{x}}_0(\mathbf{x}, t))$. In particular, Eq. 88 holds at almost every $\mathbf{x} \in \mathbb{R}^d$.

Proof. Fix $\mathbf{x}$ where $\hat{\mathbf{x}}_0(\cdot, t)$ and f_t^λ are differentiable. Assumption C.6 guarantees that $\mathrm{Prox}_{\lambda_t f}(\hat{\mathbf{x}}_0(\mathbf{x}, t))$ is nonempty (Rockafellar & Wets, 1998, Theorem 1.25), so choose any $\mathbf{z} \in \mathrm{Prox}_{\lambda_t f}(\hat{\mathbf{x}}_0(\mathbf{x}, t))$. By definition of f_t^λ,

$$f_t^\lambda(\mathbf{y}) \leq f(\mathbf{z}) + \frac{1}{2\lambda_t}\|\hat{\mathbf{x}}_0(\mathbf{y}, t) - \mathbf{z}\|^2, \qquad \mathbf{y} \in \mathbb{R}^d \tag{89}$$

Eq. 89 holds with equality at $\mathbf{y} = \mathbf{x}$ because $\mathbf{z}$ is a proximal minimizer. Therefore, $\mathbf{y} = \mathbf{x}$ minimizes

$$f(\mathbf{z}) + \frac{1}{2\lambda_t}\|\hat{\mathbf{x}}_0(\mathbf{y}, t) - \mathbf{z}\|^2 - f_t^\lambda(\mathbf{y}) \tag{90}$$

Since $\hat{\mathbf{x}}_0(\cdot, t)$ and f_t^λ are differentiable at $\mathbf{x}$, first-order optimality gives

$$\nabla f_t^\lambda(\mathbf{x}) = \frac{1}{\lambda_t}\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)^\top \left(\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{z}\right) \tag{91}$$

Since $\mathbf{z}$ was arbitrary, Eq. 88 holds at $\mathbf{x}$.

By Assumption C.5, the denoiser $\hat{\mathbf{x}}_0(\cdot, t)$ is locally Lipschitz. By Lemma C.13, f_t^λ is locally Lipschitz. Thus, by Rademacher's theorem, both functions are differentiable almost everywhere, and Eq. 88 holds at almost every $\mathbf{x}$. □

D p_t^λ CONTINUITY AND CONVERGENCE

We prove that the guided probability path from Section 4.1,

$$p_t^\lambda(\mathbf{x}) \propto q_t(\mathbf{x}) \exp(-f_t^\lambda(\mathbf{x})) \tag{92}$$

extends continuously to $[0, T]$ in W_p for every $p \in [1, \infty)$, with the target posterior as its endpoint at $t = 0$. Our assumptions allow nonsmooth, nonconvex, and extended-valued objectives, including hard constraints. We state our assumptions on the unconditional path $(q_t)_{t\in[0,T]}$ in terms of its densities and potentials

$$V_t(\mathbf{x}) := -\log q_t(\mathbf{x}), \qquad t \in [0, T] \tag{93}$$

We consider two regimes according to the Lebesgue measure of $\mathrm{dom}(f) = \{\mathbf{x} : f(\mathbf{x}) < \infty\}$. When $\mathrm{dom}(f)$ has positive measure, we are in the *volume regime*, which covers soft objectives and inequality constraints. The target posterior then has density p_{tar}, and we prove that the path converges to p_{tar} in total variation (TV). When $\mathrm{dom}(f)$ has measure zero, we are in the *surface regime*, which covers equality constraints. The target is then the surface posterior π_{tar} (Eq. 136), and we prove that the path converges to π_{tar} in W_p for every $p \in [1, \infty)$.

Let g be a finite soft objective and $\mathcal{C}$ a closed constraint set. Then $f = g + \chi_{\mathcal{C}}$ has $\mathrm{dom}(f) = \mathcal{C}$. In this appendix, we write $f = g + \chi_{\mathcal{S}}$ instead. The set $\mathcal{S}$ carries only the measure-zero constraints (e.g., equalities). The function g carries the soft objective and every other constraint (e.g., inequalities), and $g = +\infty$ where a constraint is violated. We assume that $\mathcal{S} \subseteq \mathbb{R}^d$ is nonempty and closed, and that $g : \mathcal{S} \to (-\infty, +\infty]$ is proper, lower semicontinuous on $\mathcal{S}$, and bounded below. The volume regime takes $\mathcal{S} = \mathbb{R}^d$. The surface regime requires $\mathcal{S}$ to be a C^2 surface satisfying the geometric assumptions below.

In this appendix, c and C denote positive constants independent of t and the integration variables; their values may change from line to line. We write $a \lesssim b$ for $a \le Cb$. When proving limits as $t \to 0$, we assert bounds only for sufficiently small t.

D.1 ASSUMPTIONS

We group the assumptions into three sets. The first set concerns the prior path $(q_t)_{t\in[0,T]}$ and applies to both the volume and surface regimes. The second and third sets are specific to the volume and surface regimes respectively.

D.1.1 SHARED ASSUMPTIONS

Assumption D.1. There exists a family of probability densities $(q_t)_{t\in[0,T]}$ on $\mathbb{R}^d$ such that:

1. $q_t(\mathbf{x}) > 0$ for all $(\mathbf{x}, t) \in \mathbb{R}^d \times [0, T]$.
2. The map $(\mathbf{x}, t) \mapsto q_t(\mathbf{x})$ is jointly continuous on $\mathbb{R}^d \times [0, T]$.

Assumption D.2. There exist $M > 0$, $a_V > 0$, and $b_V \in \mathbb{R}$ such that

$$V_t(\mathbf{x}) \ge a_V \|\mathbf{x}\| + b_V, \qquad t \in [0, T],\ \|\mathbf{x}\| > M \tag{94}$$

Assumption D.1 ensures that the prior potentials V_t are well-defined and converge as $t \to 0$. Assumption D.2 prevents probability mass from escaping to infinity as $t \to 0$.

D.1.2 VOLUME REGIME

Assumption D.3. The set $\{\mathbf{x} \in \mathbb{R}^d : f(\mathbf{x}) < \infty\}$ has positive Lebesgue measure.

D.1.3 SURFACE REGIME

Under equality constraints, Assumption D.3 fails, and we instead impose geometric assumptions on $\mathcal{S}$. We write $\mathcal{H}^n$ for the n-dimensional Hausdorff measure on $\mathbb{R}^d$ and $B_R := \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \le R\}$ for the closed ball of radius R centered at the origin. The reach of a closed set $\mathcal{S}$ is the largest radius $r \in [0, \infty]$ such that every $\mathbf{x}$ with $\mathrm{dist}(\mathbf{x}, \mathcal{S}) < r$ has a unique nearest point in $\mathcal{S}$ (Federer, 1959).

Assumption D.4. $\mathcal{S} \subseteq \mathbb{R}^d$ is a nonempty, closed, embedded C^2 submanifold without boundary, of dimension $n \in \{0, \dots, d-1\}$ and codimension $m := d - n$. There exists $\rho \in (0, \infty)$ such that $\mathrm{reach}(\mathcal{S}) \geq \rho$.

Assumption D.5. $g : \mathcal{S} \to (-\infty, +\infty]$ is lower semicontinuous on $\mathcal{S}$ and bounded below by a constant $C_g \in \mathbb{R}$.

Assumption D.6. There exist $t_d \in (0, T]$, $0 \leq C_d < \infty$, and $k \geq 1$ such that

$$\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x}\| \leq C_d \sigma_t^2 (1 + \|\mathbf{x}\|)^k, \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in (0, t_d] \tag{95}$$

Assumption D.7. $\mathcal{H}^n(\{\mathbf{y} \in \mathcal{S} : g(\mathbf{y}) < \infty\}) > 0$.

Assumption D.4 states that points sufficiently close to $\mathcal{S}$ has a unique nearest-point projection onto $\mathcal{S}$. Assumption D.5 states that g is lower semicontinuous and bounded below, mirroring Assumption C.6. Assumption D.6 states that the denoiser moves each point $\mathbf{x}$ by at most $\mathcal{O}(\sigma_t^2)$, up to a polynomial factor in $\|\mathbf{x}\|$. This assumption holds under a uniform Lipschitz condition on the score estimator (Lee et al., 2022, Assumption 2 and Theorem 3.1). Assumption D.7 states the feasible set must have positive surface measure.

D.2 TV-Continuity of $(p_t^\lambda)_t$ on $(0, T]$

We now prove that the path $t \mapsto p_t^\lambda$ is continuous in TV on the interior $(0, T]$. Interior continuity does not require the volume or surface assumptions.

Lemma D.8. *Assume C.1, C.2, C.4, C.6, and D.1. Then for every $t^* \in (0, T]$,*

$$\lim_{t \to t^*} \|p_t^\lambda - p_{t^*}^\lambda\|_{\mathrm{TV}} = 0 \tag{96}$$

Proof. Fix $t^* \in (0, T]$ and a sequence $t_n \to t^*$ in $(0, T]$, and write $\tilde{p}_t^\lambda := q_t \exp(-f_t^\lambda)$.

Convergence of normalization constants. By Assumption D.1 and Proposition C.9, $q_{t_n} \to q_{t^*}$ and $f_{t_n}^\lambda \to f_{t^*}^\lambda$ pointwise. Scheffé's lemma gives $\|q_{t_n} - q_{t^*}\|_{L^1} \to 0$, and $f \geq C_f$ implies $f_t^\lambda \geq C_f$.

We write

$$\tilde{p}_{t_n}^\lambda - \tilde{p}_{t^*}^\lambda = \underbrace{(q_{t_n} - q_{t^*}) \exp(-f_{t_n}^\lambda)}_{\text{Term 1}} + \underbrace{q_{t^*} \left(\exp(-f_{t_n}^\lambda) - \exp(-f_{t^*}^\lambda)\right)}_{\text{Term 2}} \tag{97}$$

Term 1 tends to zero in L^1 because its norm is bounded by $\exp(-C_f)\|q_{t_n} - q_{t^*}\|_{L^1} \to 0$. Term 2 tends to zero in L^1 by dominated convergence: it vanishes pointwise and is bounded by $\exp(-C_f) q_{t^*} \in L^1$. Thus,

$$|Z_{t_n}^{\mathrm{prox}} - Z_{t^*}^{\mathrm{prox}}| \leq \|\tilde{p}_{t_n}^\lambda - \tilde{p}_{t^*}^\lambda\|_{L^1} \to 0 \tag{98}$$

Convergence in total variation. Since $q_{t^*} > 0$ and $f_{t^*}^\lambda < \infty$ (Proposition C.9), $\tilde{p}_{t^*}^\lambda > 0$ and $Z_{t^*}^{\mathrm{prox}} > 0$. Therefore, $p_{t_n}^\lambda(\mathbf{x}) \to p_{t^*}^\lambda(\mathbf{x})$ for every $\mathbf{x} \in \mathbb{R}^d$. Then, Scheffé's lemma yields

$$\|p_{t_n}^\lambda - p_{t^*}^\lambda\|_{\mathrm{TV}} \to 0 \tag{99}$$

□

D.3 Convergence in the Volume Regime

In the volume regime, the target has density $p_{\mathrm{tar}} \propto \exp(-f - V_0)$, and we prove that $p_t^\lambda \to p_{\mathrm{tar}}$ in TV.

Proposition D.9. *Assume C.1, C.2, C.3, C.4, C.6, D.1, and D.3. Then*

$$\lim_{t \to 0} \|p_t^\lambda - p_{\mathrm{tar}}\|_{\mathrm{TV}} = 0 \tag{100}$$

Proof. Define

$$\tilde{p}_t^{\lambda}(\mathbf{x}) \coloneqq q_t(\mathbf{x}) \exp(-f_t^{\lambda}(\mathbf{x})), \qquad \tilde{p}_{\mathrm{tar}}(\mathbf{x}) \coloneqq q_0(\mathbf{x}) \exp(-f(\mathbf{x})) \tag{101}$$

with the convention $\exp(-\infty) = 0$.

Positive and finite normalization constants. By Assumption C.6, f and f_t^{λ} are bounded below by C_f. Consequently,

$$0 \le \tilde{p}_t^{\lambda} \le \exp(-C_f) q_t, \qquad 0 \le \tilde{p}_{\mathrm{tar}} \le \exp(-C_f) q_0 \tag{102}$$

so $Z_t^{\mathrm{prox}}, Z_{\mathrm{tar}} \le \exp(-C_f) < \infty$. The proof of Lemma D.8 already establishes $Z_t^{\mathrm{prox}} > 0$ for every $t \in (0, T]$. Furthermore, $q_0 > 0$ by Assumption D.1, so $\tilde{p}_{\mathrm{tar}}$ is strictly positive on $\mathrm{dom}(f)$. The set $\mathrm{dom}(f)$ has positive Lebesgue measure by Assumption D.3; hence $Z_{\mathrm{tar}} > 0$.

Convergence of unnormalized densities. Proposition C.11 gives

$$\exp(-f_t^{\lambda}(\mathbf{x})) \to \exp(-f(\mathbf{x})), \qquad \mathbf{x} \in \mathbb{R}^d \tag{103}$$

including points where $f(\mathbf{x}) = \infty$. By Assumption D.1, $q_t \to q_0$ pointwise, and Scheffé's lemma gives $\|q_t - q_0\|_{L^1} \to 0$. We write

$$\tilde{p}_t^{\lambda} - \tilde{p}_{\mathrm{tar}} = (q_t - q_0) \exp(-f_t^{\lambda}) + q_0 \left(\exp(-f_t^{\lambda}) - \exp(-f) \right) \tag{104}$$

The triangle inequality gives

$$\|\tilde{p}_t^{\lambda} - \tilde{p}_{\mathrm{tar}}\|_{L^1} \le \underbrace{\exp(-C_f) \|q_t - q_0\|_{L^1}}_{\text{Term 1}} + \underbrace{\int_{\mathbb{R}^d} q_0(\mathbf{x}) \left| \exp(-f_t^{\lambda}(\mathbf{x})) - \exp(-f(\mathbf{x})) \right| \, d\mathbf{x}}_{\text{Term 2}} \tag{105}$$

Term 1 tends to zero by the L^1 convergence of q_t. Term 2 tends to zero by dominated convergence: its integrand vanishes pointwise and is bounded by $\exp(-C_f) q_0 \in L^1$. Thus,

$$|Z_t^{\mathrm{prox}} - Z_{\mathrm{tar}}| \le \|\tilde{p}_t^{\lambda} - \tilde{p}_{\mathrm{tar}}\|_{L^1} \to 0 \tag{106}$$

Convergence in total variation. Since $\tilde{p}_t^{\lambda} \to \tilde{p}_{\mathrm{tar}}$ pointwise and $Z_t^{\mathrm{prox}} \to Z_{\mathrm{tar}} > 0$, we have $p_t^{\lambda} \to p_{\mathrm{tar}}$ pointwise. Scheffé's lemma then yields

$$\|p_t^{\lambda} - p_{\mathrm{tar}}\|_{\mathrm{TV}} \to 0 \qquad \text{as } t \to 0 \tag{107}$$

□

Finally, the volume regime's main result.

Theorem D.10. *Assume the hypotheses of Lemma D.8 and Proposition D.9. Then the map $t \mapsto p_t^{\lambda}$ is continuous on $[0, T]$ in total variation, with the convention $p_0^{\lambda} \coloneqq p_{\mathrm{tar}}$.*

D.4 Convergence in the Surface Regime

We prove that p_t^{λ} converges to the surface posterior π_{tar} in W_p for every $p \in [1, \infty)$ as $t \to 0$.

Lemma D.11. *Assume D.1 and D.2. Then*

1. *There exists $0 < C < \infty$ such that*

$$q_t(\mathbf{x}) \le C \exp(-a_V \|\mathbf{x}\|), \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [0, T] \tag{108}$$

Consequently, for every $p \in [1, \infty)$,

$$\sup_{t \in [0,T]} \int_{\mathbb{R}^d} (1 + \|\mathbf{x}\|^p) q_t(\mathbf{x}) \, d\mathbf{x} < \infty \tag{109}$$

2. *Assume additionally C.1 and D.6, and define*

$$R_t \coloneqq \lambda_t^{-1/(2k)}, \qquad \varepsilon_t \coloneqq \frac{\sigma_t^2}{\sqrt{\lambda_t}} \tag{110}$$

Then $R_t \to \infty$ as $t \to 0$ and, for all sufficiently small $t > 0$,

$$\sup_{\|\mathbf{x}\| \le R_t} \|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x}\| \le 2^k C_d \varepsilon_t \to 0 \tag{111}$$

Proof. **Part 1.** By joint continuity on the compact set $B_M \times [0,T]$,

$$Q \coloneqq \max_{\|\mathbf{x}\| \le M,\ t\in[0,T]} q_t(\mathbf{x}) < \infty \tag{112}$$

For $\|\mathbf{x}\| > M$, Assumption D.2 implies $q_t(\mathbf{x}) \le \exp(-b_V)\exp(-a_V\|\mathbf{x}\|)$. Thus,

$$q_t(\mathbf{x}) \le C\exp(-a_V\|\mathbf{x}\|), \qquad C \coloneqq \max\{Q\exp(a_V M), \exp(-b_V)\} \tag{113}$$

throughout $\mathbb{R}^d \times [0,T]$. For every $p \in [1,\infty)$, the function $(1+\|\mathbf{x}\|^p)\exp(-a_V\|\mathbf{x}\|)$ is integrable over $\mathbb{R}^d$, which proves Eq. 109.

Part 2. By Assumption C.1, $\lambda_t \to 0$ and $\varepsilon_t \to 0$, so $R_t \to \infty$. For sufficiently small $t > 0$, we have $t \le t_d$ and $R_t \ge 1$. Assumption D.6 therefore yields

$$\sup_{\|\mathbf{x}\|\le R_t} \|\hat{\mathbf{x}}_0(\mathbf{x},t) - \mathbf{x}\| \le C_d\sigma_t^2(1+R_t)^k \le 2^k C_d\sigma_t^2 R_t^k = 2^k C_d\varepsilon_t \to 0 \tag{114}$$

which proves Eq. 111. □

D.4.1 Tube Preliminaries

Define the distance to $\mathcal{S}$ and the open tube of radius $r > 0$ by

$$d_{\mathcal{S}}(\mathbf{x}) \coloneqq \inf_{\mathbf{z}\in\mathcal{S}} \|\mathbf{x}-\mathbf{z}\|, \qquad U_r \coloneqq \{\mathbf{x} : d_{\mathcal{S}}(\mathbf{x}) < r\} \tag{115}$$

For $\mathbf{y} \in \mathcal{S}$, let $T_{\mathbf{y}}\mathcal{S}$ and $N_{\mathbf{y}}\mathcal{S}$ denote the tangent and normal spaces at $\mathbf{y}$. Define the normal bundle by

$$\mathcal{N} \coloneqq \{(\mathbf{y},\mathbf{v}) : \mathbf{y} \in \mathcal{S},\ \mathbf{v} \in N_{\mathbf{y}}\mathcal{S}\} \tag{116}$$

We equip $\mathcal{N}$ with the measure μ defined by

$$\int_{\mathcal{N}} H\, d\mu \coloneqq \int_{\mathcal{S}}\int_{N_{\mathbf{y}}\mathcal{S}} H(\mathbf{y},\mathbf{v})\,\mathcal{L}^m(d\mathbf{v})\,\mathcal{H}^n(d\mathbf{y}) \tag{117}$$

for every nonnegative Borel function H, where $\mathcal{L}^m$ denotes Euclidean volume on each normal space.

Let II denote the second fundamental form. For $\mathbf{v} \in N_{\mathbf{y}}\mathcal{S}$, we define the shape operator $A_{\mathbf{y}}(\mathbf{v})$ by

$$\langle A_{\mathbf{y}}(\mathbf{v})\mathbf{a},\mathbf{b}\rangle = \langle \mathrm{II}_{\mathbf{y}}(\mathbf{a},\mathbf{b}),\mathbf{v}\rangle, \qquad \mathbf{a},\mathbf{b} \in T_{\mathbf{y}}\mathcal{S} \tag{118}$$

and set $J(\mathbf{y},\mathbf{v}) \coloneqq \det(I_n - A_{\mathbf{y}}(\mathbf{v}))$.

Fix $r_0 \coloneqq \rho/2$, and recall $R_t = \lambda_t^{-1/(2k)}$ and $\varepsilon_t = \sigma_t^2/\sqrt{\lambda_t}$ from Lemma D.11.

Lemma D.12. *Assume D.4. Then for every $(\mathbf{y},\mathbf{v}) \in \mathcal{N}$ with $\|\mathbf{v}\| < \rho$,*

$$d_{\mathcal{S}}(\mathbf{y}+\mathbf{v}) = \|\mathbf{v}\|, \qquad 0 < \left(1-\frac{\|\mathbf{v}\|}{\rho}\right)^n \le J(\mathbf{y},\mathbf{v}) \le \left(1+\frac{\|\mathbf{v}\|}{\rho}\right)^n \tag{119}$$

For every $0 < r < \rho$, the map $E(\mathbf{y},\mathbf{v}) = \mathbf{y}+\mathbf{v}$ is a bijection from $\{(\mathbf{y},\mathbf{v}) \in \mathcal{N} : \|\mathbf{v}\| < r\}$ onto U_r. Consequently, for every nonnegative Borel function $h : \mathbb{R}^d \to [0,\infty]$,

$$\int_{U_r} h(\mathbf{x})\,d\mathbf{x} = \int_{\mathcal{S}}\int_{\substack{\mathbf{v}\in N_{\mathbf{y}}\mathcal{S}\\ \|\mathbf{v}\|<r}} h(\mathbf{y}+\mathbf{v})J(\mathbf{y},\mathbf{v})\,\mathcal{L}^m(d\mathbf{v})\,\mathcal{H}^n(d\mathbf{y}) \tag{120}$$

Proof. **Normal coordinates.** Let $P_{\mathcal{S}}$ denote the nearest-point projection on U_ρ. Since $\mathrm{reach}(\mathcal{S}) \ge \rho$, every $\mathbf{x} \in U_\rho$ decomposes uniquely as

$$\mathbf{x} = \mathbf{y}+\mathbf{v}, \qquad (\mathbf{y},\mathbf{v}) \in \mathcal{N}, \quad \|\mathbf{v}\| < \rho \tag{121}$$

with $P_{\mathcal{S}}(\mathbf{x}) = \mathbf{y}$ (Federer, 1959, Remark 4.6 and Theorem 4.8(13)). Hence $d_{\mathcal{S}}(\mathbf{y}+\mathbf{v}) = \|\mathbf{v}\|$, E is a bijection from $\{\|\mathbf{v}\| < r\}$ onto U_r, and for every $\mathbf{y} \in \mathcal{S}$,

$$P_{\mathcal{S}}^{-1}(\{\mathbf{y}\}) \cap U_r = \{\mathbf{y}+\mathbf{v} : \mathbf{v} \in N_{\mathbf{y}}\mathcal{S},\ \|\mathbf{v}\| < r\} \tag{122}$$

Jacobian bounds. For $n = 0$, $J = 1$, so the bounds hold directly. For $n \geq 1$, Leobacher & Steinicke (2021, Proposition 4, Lemma 5, Proposition 6) imply

$$\|A_{\mathbf{y}}(\boldsymbol{\eta})\|_{\mathrm{op}} \leq \rho^{-1}, \qquad \boldsymbol{\eta} \in N_{\mathbf{y}}\mathcal{S}, \quad \|\boldsymbol{\eta}\| = 1 \tag{123}$$

Since $A_{\mathbf{y}}(\mathbf{v})$ is linear in $\mathbf{v}$,

$$\|A_{\mathbf{y}}(\mathbf{v})\|_{\mathrm{op}} \leq \frac{\|\mathbf{v}\|}{\rho} \tag{124}$$

Thus, the eigenvalues of $I_n - A_{\mathbf{y}}(\mathbf{v})$ lie between $1 - \|\mathbf{v}\|/\rho$ and $1 + \|\mathbf{v}\|/\rho$. Since $J(\mathbf{y}, \mathbf{v}) = \det\left(I_n - A_{\mathbf{y}}(\mathbf{v})\right)$ is the product of the n eigenvalues,

$$\left(1 - \frac{\|\mathbf{v}\|}{\rho}\right)^n \leq J(\mathbf{y}, \mathbf{v}) \leq \left(1 + \frac{\|\mathbf{v}\|}{\rho}\right)^n \tag{125}$$

Tube formula. Fix $0 < r < \rho$. By Federer (1959, Theorem 4.8(8)), $P_{\mathcal{S}}$ is Lipschitz on U_r. By Leobacher & Steinicke (2021, Theorem C),

$$DP_{\mathcal{S}}(\mathbf{y} + \mathbf{v}) = \left(I_n - A_{\mathbf{y}}(\mathbf{v})\right)^{-1} P_{T_{\mathbf{y}}\mathcal{S}} \tag{126}$$

where $P_{T_{\mathbf{y}}\mathcal{S}}$ denotes orthogonal projection onto the tangent space. Since $A_{\mathbf{y}}(\mathbf{v})$ is self-adjoint and $J(\mathbf{y}, \mathbf{v}) > 0$, the n-dimensional Jacobian of $P_{\mathcal{S}}$ satisfies

$$\mathrm{Jac}_n\, P_{\mathcal{S}}(\mathbf{y} + \mathbf{v}) = \frac{1}{J(\mathbf{y}, \mathbf{v})} \tag{127}$$

The coarea formula (Federer, 1959, Theorem 3.1) with weight $h/\,\mathrm{Jac}_n\, P_{\mathcal{S}}$ gives

$$\int_{U_r} h(\mathbf{x})\, d\mathbf{x} = \int_{\mathcal{S}} \int_{P_{\mathcal{S}}^{-1}(\{\mathbf{y}\}) \cap U_r} h(\mathbf{x}) J(\mathbf{y}, \mathbf{x} - \mathbf{y})\, \mathcal{H}^m(d\mathbf{x})\, \mathcal{H}^n(d\mathbf{y}) \tag{128}$$

By Eq. 122, $\mathbf{x} = \mathbf{y} + \mathbf{v}$ with $\mathbf{v} \in N_{\mathbf{y}}\mathcal{S}$ and $\|\mathbf{v}\| < r$, and $\mathcal{H}^m(d\mathbf{x}) = \mathcal{L}^m(d\mathbf{v})$. Substituting $\mathbf{x} = \mathbf{y} + \mathbf{v}$ yields Eq. 120. □

Corollary D.13. *Assume D.4. There exists $C < \infty$ such that*

$$\mathcal{H}^n(\mathcal{S} \cap B_R) \leq C(R + r_0)^d, \qquad R > 0 \tag{129}$$

Consequently, for every $a > 0$,

$$\int_{\mathcal{S}} \exp(-a\|\mathbf{y}\|)\, \mathcal{H}^n(d\mathbf{y}) < \infty \tag{130}$$

Proof. **Surface-area bound.** Fix $R > 0$. Substitute $h = \mathbf{1}_{B_{R+r_0}}$ into the tube formula (Eq. 120):

$$\underbrace{\int_{U_{r_0}} \mathbf{1}_{B_{R+r_0}}(\mathbf{x})\, d\mathbf{x}}_{\text{Term 1}} = \underbrace{\int_{\mathcal{S}} \int_{\substack{\mathbf{v} \in N_{\mathbf{y}}\mathcal{S} \\ \|\mathbf{v}\| < r_0}} \mathbf{1}_{B_{R+r_0}}(\mathbf{y} + \mathbf{v}) J(\mathbf{y}, \mathbf{v})\, \mathcal{L}^m(d\mathbf{v})\, \mathcal{H}^n(d\mathbf{y})}_{\text{Term 2}} \tag{131}$$

Let ω_j denote the volume of the unit ball in $\mathbb{R}^j$. Since the integrand of Term 1 is nonnegative, enlarge its domain from U_{r_0} to $\mathbb{R}^d$ to obtain Term 1 $\leq \omega_d(R + r_0)^d$. Since the integrand of Term 2 is nonnegative, restrict its outer integral to $\mathcal{S} \cap B_R$, where $\mathbf{y} + \mathbf{v} \in B_{R+r_0}$ and $h(\mathbf{y} + \mathbf{v}) = 1$. Apply $J \geq 2^{-n}$ from Lemma D.12 to obtain

$$\text{Term 2} \geq 2^{-n} \int_{\mathcal{S} \cap B_R} \int_{\substack{\mathbf{v} \in N_{\mathbf{y}}\mathcal{S} \\ \|\mathbf{v}\| < r_0}} \mathcal{L}^m(d\mathbf{v})\, \mathcal{H}^n(d\mathbf{y}) = 2^{-n} \omega_m r_0^m\, \mathcal{H}^n(\mathcal{S} \cap B_R) \tag{132}$$

Hence $\mathcal{H}^n(\mathcal{S} \cap B_R) \leq C(R + r_0)^d$.

Exponential integrability. For every $a > 0$,

$$\exp(-a\|\mathbf{y}\|) = a \int_{\|\mathbf{y}\|}^{\infty} \exp(-aR)\, dR \tag{133}$$

By Tonelli's theorem and the surface-area bound,

$$\int_{\mathcal{S}} \exp(-a\|\mathbf{y}\|)\, \mathcal{H}^n(d\mathbf{y}) = a \int_0^{\infty} \exp(-aR) \mathcal{H}^n(\mathcal{S} \cap B_R)\, dR \tag{134}$$

$$\leq Ca \int_0^{\infty} \exp(-aR)(R + r_0)^d\, dR < \infty \tag{135}$$

□

D.4.2 THE SURFACE POSTERIOR

Our target is the *surface posterior*

$$\pi_{\mathrm{tar}}(A) \coloneqq \frac{1}{Z_{\mathrm{surf}}} \int_{A\cap\mathcal{S}} \exp(-g) q_0 \, d\mathcal{H}^n, \qquad Z_{\mathrm{surf}} \coloneqq \int_{\mathcal{S}} \exp(-g) q_0 \, d\mathcal{H}^n \tag{136}$$

for Borel sets $A \subseteq \mathbb{R}^d$. The posterior is supported on $\mathcal{S}$ and has density proportional to $\exp(-g)q_0$.
Remark D.14. Under Assumptions D.1, D.2, D.4, D.5, and D.7, $Z_{\mathrm{surf}} \in (0,\infty)$. Thus π_{tar} is a well-defined probability measure with finite moments of every order $p \in [1,\infty)$.

Proof. **Positivity.** By Assumption D.1, $q_0 > 0$, so $\exp(-g)q_0$ is strictly positive on $\{g < \infty\}$. By Assumption D.7, the domain $\{g < \infty\}$ has positive $\mathcal{H}^n$ measure. Hence $Z_{\mathrm{surf}} > 0$.

Finiteness and moments. Fix $p \in [1,\infty)$. The prior bound in Lemma D.11 and $g \geq C_g$ imply

$$\int_{\mathcal{S}} (1 + \|\mathbf{y}\|^p) \exp(-g(\mathbf{y})) q_0(\mathbf{y}) \, \mathcal{H}^n(d\mathbf{y}) \leq C_p \int_{\mathcal{S}} \exp(-(a_V/2)\|\mathbf{y}\|) \, \mathcal{H}^n(d\mathbf{y}) \tag{137}$$

By Eq. 130, the integral on the right-hand side is finite. Therefore, $Z_{\mathrm{surf}} < \infty$ and

$$\int_{\mathbb{R}^d} \|\mathbf{x}\|^p \, d\pi_{\mathrm{tar}} = \frac{1}{Z_{\mathrm{surf}}} \int_{\mathcal{S}} \|\mathbf{y}\|^p \exp(-g(\mathbf{y})) q_0(\mathbf{y}) \, \mathcal{H}^n(d\mathbf{y}) < \infty \tag{138}$$

□

D.4.3 THE TV OBSTRUCTION

For probability measures, we use the normalization

$$\|\mu - \nu\|_{\mathrm{TV}} \coloneqq \sup_{A\subseteq\mathbb{R}^d,\ A \text{ Borel}} |\mu(A) - \nu(A)| \in [0,1] \tag{139}$$

Proposition D.15. *Assume D.4 and the hypotheses of Lemma D.8. For every $t \in (0,T]$ and every probability measure π with $\pi(\mathcal{S}) = 1$,*

$$\|p_t^\lambda - \pi\|_{\mathrm{TV}} = 1 \tag{140}$$

Proof. Since $\mathcal{S}$ is a closed submanifold of dimension $n < d$, $\mathcal{S}$ is Borel and has zero Lebesgue measure. Since p_t^λ has a Lebesgue density, $p_t^\lambda(\mathcal{S}) = 0$. Therefore,

$$1 = |p_t^\lambda(\mathcal{S}) - \pi(\mathcal{S})| \leq \|p_t^\lambda - \pi\|_{\mathrm{TV}} \leq 1 \tag{141}$$

□

Under the assumptions of Remark D.14, π_{tar} is a probability measure with $\pi_{\mathrm{tar}}(\mathcal{S}) = 1$. Hence $\|p_t^\lambda - \pi_{\mathrm{tar}}\|_{\mathrm{TV}} = 1$ for every $t \in (0,T]$, and p_t^λ cannot converge to π_{tar} in TV.

D.4.4 THE LAPLACE-TYPE LIMIT

Lemma D.16. *Assume C.1, D.4, D.5, and D.6, with $f = g + \chi_{\mathcal{S}}$. Fix $\mathbf{y} \in \mathcal{S}$ and $\mathbf{w} \in N_{\mathbf{y}}\mathcal{S}$, and set*

$$\mathbf{x}_t \coloneqq \mathbf{y} + \sqrt{\lambda_t}\mathbf{w}, \qquad \mathbf{u}_t \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}_t, t) \tag{142}$$

As $t \to 0$,

$$\exp(-\Psi_f(\mathbf{u}_t, \lambda_t)) \to \exp(-g(\mathbf{y}) - \|\mathbf{w}\|^2/2) \tag{143}$$

Proof. **Step 1: The denoised point.** For sufficiently small t, we have $t \le t_d$ and $\lambda_t \le 1$, so $\|\mathbf{x}_t\| \le \|\mathbf{y}\| + \|\mathbf{w}\|$. Assumption D.6 then gives

$$\|\mathbf{u}_t - \mathbf{x}_t\| \le c_1 \sigma_t^2 = c_1 \varepsilon_t \sqrt{\lambda_t}, \qquad c_1 := C_d(1 + \|\mathbf{y}\| + \|\mathbf{w}\|)^k \tag{144}$$

By Assumption C.1, $\varepsilon_t \to 0$, hence $\mathbf{u}_t = \mathbf{y} + \sqrt{\lambda_t}\mathbf{w} + o(\sqrt{\lambda_t})$. Once $\sqrt{\lambda_t}\|\mathbf{w}\| < \rho$, $d_{\mathcal{S}}(\mathbf{x}_t) = \sqrt{\lambda_t}\|\mathbf{w}\|$ by Lemma D.12. The distance $d_{\mathcal{S}}$ is 1-Lipschitz, so

$$\frac{d_{\mathcal{S}}(\mathbf{u}_t)}{\sqrt{\lambda_t}} \to \|\mathbf{w}\| \tag{145}$$

Step 2: Lower bound on the envelope. Let $E_t := \Psi_f(\mathbf{u}_t, \lambda_t)$. Since $f = +\infty$ off $\mathcal{S}$,

$$E_t = \inf_{\mathbf{z} \in \mathcal{S}} \left\{ g(\mathbf{z}) + \frac{\|\mathbf{u}_t - \mathbf{z}\|^2}{2\lambda_t} \right\} \tag{146}$$

Fix any real $b < g(\mathbf{y})$. Since g is lower semicontinuous, there exists $\eta > 0$ such that $g(\mathbf{z}) \ge b$ whenever $\mathbf{z} \in \mathcal{S}$ and $\|\mathbf{z} - \mathbf{y}\| < \eta$. Since $\mathbf{u}_t \to \mathbf{y}$, we have $\|\mathbf{u}_t - \mathbf{y}\| < \eta/2$ for sufficiently small t. At every $\mathbf{z} \in \mathcal{S}$ with $\|\mathbf{z} - \mathbf{y}\| < \eta$, the objective in Eq. 146 is at least $b + d_{\mathcal{S}}(\mathbf{u}_t)^2/(2\lambda_t)$, because $g(\mathbf{z}) \ge b$ and $\|\mathbf{u}_t - \mathbf{z}\| \ge d_{\mathcal{S}}(\mathbf{u}_t)$. At every remaining $\mathbf{z} \in \mathcal{S}$, the triangle inequality yields $\|\mathbf{u}_t - \mathbf{z}\| \ge \eta/2$, so the objective is at least $C_g + \eta^2/(8\lambda_t)$. Therefore

$$E_t \ge \min\Bigg\{ \underbrace{b + \frac{d_{\mathcal{S}}(\mathbf{u}_t)^2}{2\lambda_t}}_{\text{Term 1}}, \underbrace{C_g + \frac{\eta^2}{8\lambda_t}}_{\text{Term 2}} \Bigg\} \tag{147}$$

Term 1 converges to $b + \|\mathbf{w}\|^2/2$ by Eq. 145, and Term 2 diverges as $t \to 0$. Thus $\liminf_{t\to 0} E_t \ge b + \|\mathbf{w}\|^2/2$. Since $b < g(\mathbf{y})$ was arbitrary,

$$\liminf_{t \to 0} E_t \ge g(\mathbf{y}) + \frac{\|\mathbf{w}\|^2}{2} \tag{148}$$

including when $g(\mathbf{y}) = \infty$. In that case $E_t \to \infty$, which proves Eq. 143 under the convention $\exp(-\infty) = 0$.

Step 3: Upper bound on the envelope. If $g(\mathbf{y}) < \infty$, set $\mathbf{z} = \mathbf{y}$ in Eq. 146 to obtain

$$E_t \le g(\mathbf{y}) + \frac{\|\mathbf{u}_t - \mathbf{y}\|^2}{2\lambda_t} \to g(\mathbf{y}) + \frac{\|\mathbf{w}\|^2}{2} \tag{149}$$

The bounds in Eq. 148 and Eq. 149 prove Eq. 143. □

Proposition D.17. *Assume D.2, D.4, D.5, D.6, and the hypotheses of Lemma D.8, with $f = g + \chi_{\mathcal{S}}$. Let $\varphi : \mathbb{R}^d \to \mathbb{R}$ be continuous with $|\varphi(\mathbf{x})| \le C_\varphi(1 + \|\mathbf{x}\|^p)$ for some $C_\varphi < \infty$ and $p \in [1, \infty)$. As $t \to 0$,*

$$(2\pi\lambda_t)^{-m/2} \int_{\mathbb{R}^d} \varphi(\mathbf{x}) q_t(\mathbf{x}) \exp(-f_t^\lambda(\mathbf{x}))\, d\mathbf{x} \to \int_{\mathcal{S}} \varphi(\mathbf{y}) \exp(-g(\mathbf{y})) q_0(\mathbf{y})\, \mathcal{H}^n(d\mathbf{y}) \tag{150}$$

Proof. Fix C_φ and p as in the statement. Implicit constants may depend on C_φ and p. Since $f = +\infty$ off $\mathcal{S}$ and $g \ge C_g$ on $\mathcal{S}$,

$$f_t^\lambda(\mathbf{x}) \ge C_g + \frac{d_{\mathcal{S}}(\hat{\mathbf{x}}_0(\mathbf{x}, t))^2}{2\lambda_t}, \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in (0, T] \tag{151}$$

In particular, $\exp(-f_t^\lambda) \le \exp(-C_g)$. The growth bound on φ and the prior moment bound in Eq. 109 therefore ensure that $\varphi q_t \exp(-f_t^\lambda)$ is absolutely integrable for every $t \in (0, T]$.

Step 1: The tube complement. Recall $r_0 = \rho/2$ and $R_t = \lambda_t^{-1/(2k)}$. Call $\{d_{\mathcal{S}} \ge r_0,\ \|\mathbf{x}\| \le R_t\}$ the inner region and $\{\|\mathbf{x}\| > R_t\}$ the outer region. The inner and outer regions cover $\mathbb{R}^d \setminus U_{r_0}$.

On the inner region, the 1-Lipschitz property of $d_{\mathcal{S}}$ and Eq. 111 imply

$$d_{\mathcal{S}}(\hat{\mathbf{x}}_0(\mathbf{x}, t)) \ge r_0 - 2^k C_d \varepsilon_t \ge r_0/2 \tag{152}$$

once $2^k C_d \varepsilon_t \le r_0/2$. Substitute Eq. 152 into Eq. 151 to obtain

$$\exp(-f_t^\lambda(\mathbf{x})) \le \exp(-C_g)\exp(-r_0^2/(8\lambda_t)), \qquad d_{\mathcal{S}}(\mathbf{x}) \ge r_0, \quad \|\mathbf{x}\| \le R_t \tag{153}$$

On the outer region, Part 1 of Lemma D.11 and $\exp(-a_V\|\mathbf{x}\|/2) \le \exp(-a_V R_t/2)$ yield

$$\int_{\{\|\mathbf{x}\|>R_t\}} (1+\|\mathbf{x}\|^p) q_t(\mathbf{x})\, d\mathbf{x} \lesssim \exp(-a_V R_t/2) \int_{\mathbb{R}^d} (1+\|\mathbf{x}\|^p)\exp(-a_V\|\mathbf{x}\|/2)\, d\mathbf{x} \tag{154}$$

$$\lesssim \exp(-a_V R_t/2) \tag{155}$$

Bound $|\varphi|$ by $C_\varphi(1+\|\mathbf{x}\|^p)$ throughout $\mathbb{R}^d \setminus U_{r_0}$. On the inner region, apply Eq. 153 and the moment bound in Eq. 109. On the outer region, apply $\exp(-f_t^\lambda) \le \exp(-C_g)$ and Eq. 155. Sum over the inner and outer regions to obtain

$$(2\pi\lambda_t)^{-m/2} \int_{\mathbb{R}^d\setminus U_{r_0}} |\varphi(\mathbf{x})| q_t(\mathbf{x})\exp(-f_t^\lambda(\mathbf{x}))\, d\mathbf{x}$$

$$\lesssim \lambda_t^{-m/2}\Big(\underbrace{\exp(-r_0^2/(8\lambda_t))}_{\text{inner region}} + \underbrace{\exp(-a_V R_t/2)}_{\text{outer region}}\Big) \to 0 \tag{156}$$

since both exponentials decay faster than every power of λ_t.

Step 2: Normal coordinates and the pointwise limit. For $(\mathbf{y},\mathbf{w}) \in \mathcal{N}$, set $\mathbf{x}_t := \mathbf{y} + \sqrt{\lambda_t}\mathbf{w}$ and define the rescaled tube integrand

$$F_t(\mathbf{y},\mathbf{w}) := \mathbf{1}\{\sqrt{\lambda_t}\|\mathbf{w}\| < r_0\}\varphi(\mathbf{x}_t) q_t(\mathbf{x}_t)\exp(-f_t^\lambda(\mathbf{x}_t)) J(\mathbf{y}, \sqrt{\lambda_t}\mathbf{w}) \tag{157}$$

The tube formula (Eq. 120) requires a nonnegative integrand, so write $\varphi = \varphi_+ - \varphi_-$ with $\varphi_\pm := \max\{\pm\varphi, 0\}$. Apply Eq. 120 with $r = r_0$ to $\varphi_\pm q_t \exp(-f_t^\lambda)$ and subtract. Substituting $\mathbf{v} = \sqrt{\lambda_t}\mathbf{w}$ in each normal fibre contributes the factor $\lambda_t^{m/2}$, so

$$(2\pi\lambda_t)^{-m/2} \int_{U_{r_0}} \varphi q_t \exp(-f_t^\lambda)\, d\mathbf{x} = (2\pi)^{-m/2} \int_{\mathcal{N}} F_t\, d\mu \tag{158}$$

Fix $(\mathbf{y},\mathbf{w}) \in \mathcal{N}$ and take the limit of each factor in Eq. 157. The indicator equals one for sufficiently small t. Since $\mathbf{x}_t \to \mathbf{y}$, continuity of φ and joint continuity of $(\mathbf{x},t) \mapsto q_t(\mathbf{x})$ yield $\varphi(\mathbf{x}_t)q_t(\mathbf{x}_t) \to \varphi(\mathbf{y})q_0(\mathbf{y})$. The Jacobian bounds in Lemma D.12 imply $J(\mathbf{y}, \sqrt{\lambda_t}\mathbf{w}) \to 1$. Lemma D.16 supplies the limit of $\exp(-f_t^\lambda(\mathbf{x}_t))$, so

$$F_t(\mathbf{y},\mathbf{w}) \to \varphi(\mathbf{y}) q_0(\mathbf{y})\exp(-g(\mathbf{y}))\exp(-\|\mathbf{w}\|^2/2) \tag{159}$$

Step 3: An integrable dominating function. We bound each factor of Eq. 157 on $\{\sqrt{\lambda_t}\|\mathbf{w}\| < r_0\}$, where $\|\mathbf{x}_t - \mathbf{y}\| < r_0$. The growth bound on φ, Part 1 of Lemma D.11, and the Jacobian bounds in Lemma D.12 give

$$|\varphi(\mathbf{x}_t)| \le C_\varphi\,(1 + (\|\mathbf{y}\| + r_0)^p) \lesssim (1+\|\mathbf{y}\|)^p \tag{160}$$

$$q_t(\mathbf{x}_t) \lesssim \exp(-a_V\|\mathbf{y}\|) \tag{161}$$

$$J(\mathbf{y}, \sqrt{\lambda_t}\mathbf{w}) \le (3/2)^n \tag{162}$$

By Eq. 151, the remaining factor $\exp(-f_t^\lambda(\mathbf{x}_t))$ requires a lower bound on $d_{\mathcal{S}}(\hat{\mathbf{x}}_0(\mathbf{x}_t,t))$. Since $\|\mathbf{x}_t\| \le \|\mathbf{y}\| + r_0$, Assumption D.6 implies

$$\|\hat{\mathbf{x}}_0(\mathbf{x}_t,t) - \mathbf{x}_t\| \le \sigma_t^2 s_{\mathbf{y}}, \qquad s_{\mathbf{y}} := C_d(1+r_0)^k(1+\|\mathbf{y}\|)^k \tag{163}$$

Lemma D.12 gives $d_{\mathcal{S}}(\mathbf{x}_t) = \sqrt{\lambda_t}\|\mathbf{w}\|$. The distance $d_{\mathcal{S}}$ is nonnegative and 1-Lipschitz, so

$$d_{\mathcal{S}}(\hat{\mathbf{x}}_0(\mathbf{x}_t,t)) \ge \sqrt{\lambda_t}(\|\mathbf{w}\| - \varepsilon_t s_{\mathbf{y}})_+ \ge \sqrt{\lambda_t}(\|\mathbf{w}\| - s_{\mathbf{y}})_+ \tag{164}$$

once $\varepsilon_t \le 1$. Here $a_+ := \max\{a, 0\}$. Substitute Eq. 164 into Eq. 151 to obtain

$$\exp(-f_t^\lambda(\mathbf{x}_t)) \le \exp(-C_g)\exp\left(-\frac{1}{2}(\|\mathbf{w}\| - s_{\mathbf{y}})_+^2\right) \tag{165}$$

Multiply the factor bounds to control $|F_t|$ on $\{\sqrt{\lambda_t}\|\mathbf{w}\| < r_0\}$. F_t vanishes elsewhere, so on all of $\mathcal{N}$,

$$|F_t(\mathbf{y},\mathbf{w})| \lesssim K_p(\mathbf{y},\mathbf{w}) \coloneqq (1+\|\mathbf{y}\|)^p \exp(-a_V\|\mathbf{y}\|) \exp\left(-\frac{1}{2}(\|\mathbf{w}\| - s_{\mathbf{y}})_+^2\right) \tag{166}$$

Next, we prove that K_p is integrable over $\mathcal{N}$. By definition of μ,

$$\int_{\mathcal{N}} K_p \, d\mu = \int_{\mathcal{S}} (1+\|\mathbf{y}\|)^p \exp(-a_V\|\mathbf{y}\|) \int_{N_{\mathbf{y}}\mathcal{S}} \exp(-(\|\mathbf{w}\| - s_{\mathbf{y}})_+^2/2) \, \mathcal{L}^m(d\mathbf{w}) \, \mathcal{H}^n(d\mathbf{y}) \tag{167}$$

Bound the inner integral first. Since $r - s \geq r/2$ whenever $r > 2s$,

$$\exp(-(r-s)_+^2/2) \leq \mathbf{1}\{r \leq 2s\} + \exp(-r^2/8), \qquad r, s \geq 0 \tag{168}$$

Substitute $r = \|\mathbf{w}\|$ and $s = s_{\mathbf{y}}$ into Eq. 168 and integrate over $N_{\mathbf{y}}\mathcal{S}$. The term $\mathbf{1}\{\|\mathbf{w}\| \leq 2s_{\mathbf{y}}\}$ integrates to $\omega_m(2s_{\mathbf{y}})^m$, where ω_m is the volume of the unit ball in $\mathbb{R}^m$. The term $\exp(-\|\mathbf{w}\|^2/8)$ integrates to $(8\pi)^{m/2}$. Hence, for each $\mathbf{y} \in \mathcal{S}$,

$$\begin{aligned}\int_{N_{\mathbf{y}}\mathcal{S}} \exp(-(\|\mathbf{w}\| - s_{\mathbf{y}})_+^2/2) \, \mathcal{L}^m(d\mathbf{w}) &\leq \omega_m(2s_{\mathbf{y}})^m + (8\pi)^{m/2} \\ &\lesssim (1+s_{\mathbf{y}})^m \end{aligned} \tag{169}$$

Substitute Eq. 169 into Eq. 167 and apply $s_{\mathbf{y}} \lesssim (1+\|\mathbf{y}\|)^k$. Then apply $(1+\|\mathbf{y}\|)^{km+p}\exp(-(a_V/2)\|\mathbf{y}\|) \lesssim 1$ to obtain

$$\begin{aligned}\int_{\mathcal{N}} K_p \, d\mu &\lesssim \int_{\mathcal{S}} (1+\|\mathbf{y}\|)^{km+p} \exp(-a_V\|\mathbf{y}\|) \, \mathcal{H}^n(d\mathbf{y}) \\ &\lesssim \int_{\mathcal{S}} \exp(-(a_V/2)\|\mathbf{y}\|) \, \mathcal{H}^n(d\mathbf{y}) < \infty \end{aligned} \tag{170}$$

where the final integral is finite by Eq. 130.

Then, we apply dominated convergence to F_t on $\mathcal{N}$, with the pointwise limit in Eq. 159 and the dominating function K_p in Eq. 166. The factor $\exp(-\|\mathbf{w}\|^2/2)$ of the pointwise limit integrates to $(2\pi)^{m/2}$ over $N_{\mathbf{y}}\mathcal{S}$, so

$$(2\pi)^{-m/2} \int_{\mathcal{N}} F_t \, d\mu \to \int_{\mathcal{S}} \varphi(\mathbf{y}) \exp(-g(\mathbf{y})) q_0(\mathbf{y}) \, \mathcal{H}^n(d\mathbf{y}) \tag{171}$$

Finally, we split the integral in Eq. 150 over U_{r_0} and $\mathbb{R}^d \setminus U_{r_0}$ to write

$$\begin{aligned}(2\pi\lambda_t)^{-m/2} \int_{\mathbb{R}^d} \varphi q_t \exp(-f_t^\lambda) \, d\mathbf{x} = &\underbrace{(2\pi\lambda_t)^{-m/2} \int_{U_{r_0}} \varphi q_t \exp(-f_t^\lambda) \, d\mathbf{x}}_{\text{Term 1}} \\ &+ \underbrace{(2\pi\lambda_t)^{-m/2} \int_{\mathbb{R}^d \setminus U_{r_0}} \varphi q_t \exp(-f_t^\lambda) \, d\mathbf{x}}_{\text{Term 2}} \end{aligned} \tag{172}$$

By Eq. 158 and Eq. 171, Term 1 converges to the right-hand side of Eq. 150. By Eq. 156, Term 2 converges to zero. The two limits prove Eq. 150. □

D.4.5 Convergence to the Surface Posterior

Theorem D.18. *Assume D.7 and the hypotheses of Proposition D.17. As $t \to 0$,*

1. $(2\pi\lambda_t)^{-m/2} Z_t^{\mathrm{prox}} \to Z_{\mathrm{surf}} \in (0,\infty)$.
2. $\int \varphi \, dp_t^\lambda \to \int \varphi \, d\pi_{\mathrm{tar}}$ *for every continuous* $\varphi : \mathbb{R}^d \to \mathbb{R}$ *satisfying* $|\varphi(\mathbf{x})| \leq C_\varphi(1+\|\mathbf{x}\|^p)$ *for some* $C_\varphi < \infty$ *and* $p \in [1,\infty)$.

3. $W_p(p_t^\lambda, \pi_{\mathrm{tar}}) \to 0$ *for every* $p \in [1, \infty)$.

Proof. **Part 1.** Proposition D.17 with $\varphi \equiv 1$ yields $(2\pi\lambda_t)^{-m/2} Z_t^{\mathrm{prox}} \to Z_{\mathrm{surf}}$. By Remark D.14, $Z_{\mathrm{surf}} \in (0, \infty)$ and π_{tar} has finite moments of every order $p \in [1, \infty)$.

Part 2. Fix φ and p as in the statement. The function $\varphi q_t \exp(-f_t^\lambda)$ is integrable by the proof of Proposition D.17, so φ is integrable under p_t^λ. By Part 1, π_{tar} has finite p-th moment, so φ is integrable under π_{tar}. By the definition of p_t^λ,

$$\int_{\mathbb{R}^d} \varphi \, dp_t^\lambda = \frac{(2\pi\lambda_t)^{-m/2} \int_{\mathbb{R}^d} \varphi(\mathbf{x}) q_t(\mathbf{x}) \exp(-f_t^\lambda(\mathbf{x})) \, d\mathbf{x}}{(2\pi\lambda_t)^{-m/2} Z_t^{\mathrm{prox}}} \tag{173}$$

Proposition D.17 gives the limit of the numerator, and Part 1 gives the limit of the denominator. Hence

$$\int_{\mathbb{R}^d} \varphi \, dp_t^\lambda \to \frac{\int_{\mathcal{S}} \varphi(\mathbf{y}) \exp(-g(\mathbf{y})) q_0(\mathbf{y}) \, \mathcal{H}^n(d\mathbf{y})}{Z_{\mathrm{surf}}} = \int_{\mathbb{R}^d} \varphi \, d\pi_{\mathrm{tar}} \tag{174}$$

Part 3. Fix $p \in [1, \infty)$. Every bounded continuous function satisfies the growth bound in Part 2, so p_t^λ converges weakly to π_{tar}. Part 2 with $\varphi(\mathbf{x}) = \|\mathbf{x}\|^p$ yields

$$\int_{\mathbb{R}^d} \|\mathbf{x}\|^p \, dp_t^\lambda \to \int_{\mathbb{R}^d} \|\mathbf{x}\|^p \, d\pi_{\mathrm{tar}} < \infty \tag{175}$$

By Villani (2009, Definition 6.8 and Theorem 6.9), weak convergence and Eq. 175 imply $W_p(p_t^\lambda, \pi_{\mathrm{tar}}) \to 0$. □

D.5 W_p-Continuity of p_t^λ on $[0, T]$

We now establish continuity on $[0, T]$ in every Wasserstein distance of finite order.

Theorem D.19. *Assume C.1, C.2, C.4, C.6, D.1, and D.2. In the volume regime, assume C.3 and D.3, and set* $p_0^\lambda := p_{\mathrm{tar}}$. *In the surface regime, assume D.4 to D.7 and set* $p_0^\lambda := \pi_{\mathrm{tar}}$. *Then* $t \mapsto p_t^\lambda$ *is continuous on* $[0, T]$ *in* W_p *for every* $p \in [1, \infty)$.

Proof. Fix $p \in [1, \infty)$. In the surface regime, continuity at $t = 0$ follows from Part 3 of Theorem D.18. It remains to prove continuity at positive times in both regimes and at $t = 0$ in the volume regime.

For $t > 0$, define

$$h_t(\mathbf{x}) := q_t(\mathbf{x}) \exp(-f_t^\lambda(\mathbf{x})) \tag{176}$$

In the volume regime, define

$$h_0(\mathbf{x}) := q_0(\mathbf{x}) \exp(-f(\mathbf{x})) \tag{177}$$

The normalizing constants $Z_t := \int_{\mathbb{R}^d} h_t \, d\mathbf{x}$ lie in $(0, \infty)$ by Lemma D.8 and Proposition D.9.

Since $f_t^\lambda \geq C_f$ and $f \geq C_f$, Part 1 of Lemma D.11 implies

$$0 \leq h_t(\mathbf{x}) \lesssim \exp(-a_V \|\mathbf{x}\|), \qquad \mathbf{x} \in \mathbb{R}^d \tag{178}$$

with a constant independent of t. In particular, each p_t^λ with density h_t / Z_t has finite p-th moment.

Fix $t^* \in (0, T]$, or $t^* = 0$ in the volume regime. Assumption D.1 gives $q_t \to q_{t^*}$ pointwise. Proposition C.9 covers $t^* > 0$, and Proposition C.11 covers $t^* = 0$. Hence, as $t \to t^*$ through positive times,

$$h_t(\mathbf{x}) \to h_{t^*}(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^d \tag{179}$$

Let $\varphi : \mathbb{R}^d \to \mathbb{R}$ be continuous with $|\varphi(\mathbf{x})| \leq C_\varphi(1 + \|\mathbf{x}\|^p)$. The bound

$$|\varphi(\mathbf{x}) h_t(\mathbf{x})| \leq C C_\varphi (1 + \|\mathbf{x}\|^p) \exp(-a_V \|\mathbf{x}\|) \tag{180}$$

has an integrable right-hand side. Dominated convergence therefore yields

$$\int_{\mathbb{R}^d} \varphi h_t \, d\mathbf{x} \to \int_{\mathbb{R}^d} \varphi h_{t^*} \, d\mathbf{x} \tag{181}$$

The choice $\varphi \equiv 1$ gives $Z_t \to Z_{t^*} > 0$. Consequently,

$$\int_{\mathbb{R}^d} \varphi \, dp_t^\lambda = \frac{\int_{\mathbb{R}^d} \varphi h_t \, d\mathbf{x}}{Z_t} \to \frac{\int_{\mathbb{R}^d} \varphi h_{t^*} \, d\mathbf{x}}{Z_{t^*}} = \int_{\mathbb{R}^d} \varphi \, dp_{t^*}^\lambda \tag{182}$$

Every bounded continuous φ satisfies the growth bound, so Eq. 182 gives weak convergence of p_t^λ to $p_{t^*}^\lambda$. Eq. 182 with $\varphi(\mathbf{x}) = \|\mathbf{x}\|^p$ gives convergence of p-th moments. By Villani (2009, Definition 6.8 and Theorem 6.9), $W_p(p_t^\lambda, p_{t^*}^\lambda) \to 0$. □

Theorem D.19, together with Theorem D.10, establishes Theorem 4.2.

E GUIDED SDE PROPERTIES

E.1 SETUP

Fix $0 < \epsilon < T$ and a guidance weight $w \in \mathbb{R}$. Recall that the guided SDE in Eq. 12 has the form

$$d\mathbf{x}_t = b_t^\lambda(\mathbf{x}_t)\,dt + g(t)\,d\overline{\mathbf{w}}_t, \qquad \mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{183}$$

where

$$b_t^\lambda(\mathbf{x}) \coloneqq \mu(t)\mathbf{x} - g(t)^2\,(\xi_t(\mathbf{x}) - wG_t(\mathbf{x})) \tag{184}$$

Here ξ_t is the score component of the drift. The guidance component G_t is a jointly Borel representative of ∇f_t^λ, constructed in Lemma E.4. Both components are induced by the same denoiser:

$$\xi_t(\mathbf{x}) \coloneqq \frac{\alpha_t \hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{x}}{\sigma_t^2}, \qquad F_t(\mathbf{x}) \coloneqq f_t^\lambda(\mathbf{x}) = \Psi_f(\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t) \tag{185}$$

When $\hat{\mathbf{x}}_0$ is the exact posterior mean $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$, Tweedie's identity gives $\xi_t = \nabla \log q_t$. When $\hat{\mathbf{x}}_0$ is a learned denoiser, $\xi_t = s_\theta$.

E.2 ASSUMPTIONS

Assumption E.1. The functions $\mu, g, \alpha, \sigma, \lambda$ are continuous on $[\epsilon, T]$, with

$$g(t) > 0, \qquad \alpha_t > 0, \qquad \sigma_t > 0, \qquad \lambda_t > 0 \tag{186}$$

for every $t \in [\epsilon, T]$.

Assumption E.2. The map $(\mathbf{x}, t) \mapsto \hat{\mathbf{x}}_0(\mathbf{x}, t)$ is jointly continuous on $\mathbb{R}^d \times [\epsilon, T]$. There exists $C_J < \infty$ such that

$$\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \hat{\mathbf{x}}_0(\mathbf{y}, t)\| \le C_J \|\mathbf{x} - \mathbf{y}\|, \qquad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{187}$$

Assumption E.1 holds for standard VP, VE, and EDM schedules on $[\epsilon, T]$ (Song et al., 2021b; Karras et al., 2022). Assumption E.2 holds for the exact denoiser when the data have bounded support, since then $\nabla \log q_t$ is Lipschitz uniformly over $[\epsilon, T]$ (Bortoli, 2022, Lemma C.2). For learned denoisers, joint continuity follows from Lemma C.8, and empirical studies of trained networks support the Lipschitz bound (Mohan et al., 2020; Kadkhodaie et al., 2024; Zhan et al., 2026).

By continuity and compactness, we define

$$\begin{aligned} \lambda_{\min} &\coloneqq \min_{t\in[\epsilon,T]} \lambda_t > 0, & \lambda_{\max} &\coloneqq \max_{t\in[\epsilon,T]} \lambda_t < \infty \\ g_{\min} &\coloneqq \min_{t\in[\epsilon,T]} g(t) > 0, & g_{\max} &\coloneqq \max_{t\in[\epsilon,T]} g(t) < \infty \\ \sigma_* &\coloneqq \min_{t\in[\epsilon,T]} \sigma_t > 0, & \alpha_* &\coloneqq \max_{t\in[\epsilon,T]} \alpha_t < \infty \end{aligned} \tag{188}$$

Joint continuity of the denoiser also gives

$$H_0 \coloneqq \sup_{t\in[\epsilon,T]} \|\hat{\mathbf{x}}_0(\mathbf{0}, t)\| < \infty \tag{189}$$

The Lipschitz bound consequently implies uniform linear growth:

$$\|\hat{\mathbf{x}}_0(\mathbf{x}, t)\| \le H_0 + C_J\|\mathbf{x}\|, \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{190}$$

E.3 REGULARITY OF THE GUIDED DRIFT

We show that the guided drift b_t^λ is jointly Borel with linear growth.

Lemma E.3. *Assume C.1, C.2, C.5, C.6, E.1, and E.2. Fix $\mathbf{z}^\dagger \in \mathrm{dom}(f)$, and define*

$$A \coloneqq \sqrt{2\lambda_{\max}\,(f(\mathbf{z}^\dagger) - C_f)} + \|\mathbf{z}^\dagger\|, \qquad K_G(\mathbf{x}) \coloneqq \frac{C_J}{\lambda_{\min}}\,(A + H_0 + C_J\|\mathbf{x}\|) \tag{191}$$

Then for every $t \in [\epsilon, T]$ and almost every $\mathbf{x} \in \mathbb{R}^d$,

$$\|\nabla F_t(\mathbf{x})\| \le K_G(\mathbf{x}) \tag{192}$$

Proof. **Proximal displacement.** Fix $t \in [\epsilon, T]$, $\mathbf{u} \in \mathbb{R}^d$, and $\mathbf{z} \in \mathrm{Prox}_{\lambda_t f}(\mathbf{u})$. Since $\mathbf{z}$ minimizes the proximal objective at $\mathbf{u}$, evaluating the objective at $\mathbf{z}^\dagger$ yields

$$f(\mathbf{z}) + \frac{\|\mathbf{u}-\mathbf{z}\|^2}{2\lambda_t} \le f(\mathbf{z}^\dagger) + \frac{\|\mathbf{u}-\mathbf{z}^\dagger\|^2}{2\lambda_t} \tag{193}$$

Since $f(\mathbf{z}) \ge C_f$ and $\lambda_t \le \lambda_{\max}$, the inequality $\sqrt{a+b} \le \sqrt{a} + \sqrt{b}$ and the triangle inequality yield

$$\begin{aligned}\|\mathbf{u}-\mathbf{z}\| &\le \sqrt{2\lambda_{\max}\left(f(\mathbf{z}^\dagger) - C_f\right)} + \|\mathbf{u}-\mathbf{z}^\dagger\| \\ &\le A + \|\mathbf{u}\|\end{aligned} \tag{194}$$

Gradient bound. Fix $t \in [\epsilon, T]$. Proposition C.14 gives, at almost every $\mathbf{x}$ and for any $\mathbf{z} \in \mathrm{Prox}_{\lambda_t f}(\hat{\mathbf{x}}_0(\mathbf{x}, t))$,

$$\nabla F_t(\mathbf{x}) = \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)^\top \frac{\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{z}}{\lambda_t} \tag{195}$$

Assumption E.2 gives $\|\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)\|_{\mathrm{op}} \le C_J$. Set $\mathbf{u} = \hat{\mathbf{x}}_0(\mathbf{x}, t)$ in Eq. 194 to obtain

$$\|\hat{\mathbf{x}}_0(\mathbf{x}, t) - \mathbf{z}\| \le A + \|\hat{\mathbf{x}}_0(\mathbf{x}, t)\| \le A + H_0 + C_J\|\mathbf{x}\| \tag{196}$$

Since $\lambda_t \ge \lambda_{\min}$, the norm of $\nabla F_t(\mathbf{x})$ is at most $C_J(A + H_0 + C_J\|\mathbf{x}\|)/\lambda_{\min} = K_G(\mathbf{x})$, which proves Eq. 192. □

Lemma E.4. *Assume the hypotheses of Lemma E.3. There is a jointly Borel map $G : \mathbb{R}^d \times [\epsilon, T] \to \mathbb{R}^d$ with slices $G_t := G(\cdot, t)$ such that*

$$\|G_t(\mathbf{x})\| \le K_G(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{197}$$

For every $t \in [\epsilon, T]$, the slice G_t equals ∇F_t at almost every $\mathbf{x} \in \mathbb{R}^d$.

Proof. **Construction.** Fix positive numbers $q_n \downarrow 0$. For each coordinate $i \in \{1, \dots, d\}$, define the difference quotients and their upper limit

$$\delta_{i,n}(\mathbf{x}, t) := \frac{F_t(\mathbf{x} + q_n \mathbf{e}_i) - F_t(\mathbf{x})}{q_n}, \qquad \widetilde{G}_i := \limsup_{n\to\infty} \delta_{i,n} \tag{198}$$

The map $(\mathbf{x}, t) \mapsto F_t(\mathbf{x})$ is jointly continuous as a composition of Ψ_f (Rockafellar & Wets, 1998, Exercise 1.24 and Theorem 1.25), $\hat{\mathbf{x}}_0$ (Assumption E.2), and λ (Assumption E.1). Hence each $\delta_{i,n}$ is jointly continuous. By Lemma C.13, each F_t is locally Lipschitz, so each $\widetilde{G}_i$ is finite. Thus $\widetilde{G} := (\widetilde{G}_1, \dots, \widetilde{G}_d)$ is jointly Borel, and $\widetilde{G}(\mathbf{x}, t) = \nabla F_t(\mathbf{x})$ wherever F_t is differentiable. The candidate $\widetilde{G}$ may not satisfy Eq. 197 at every point, so define

$$G(\mathbf{x}, t) := \begin{cases} \widetilde{G}(\mathbf{x}, t) & \text{if } \|\widetilde{G}(\mathbf{x}, t)\| \le K_G(\mathbf{x}) \\ \mathbf{0} & \text{otherwise} \end{cases} \tag{199}$$

The set $\{(\mathbf{x}, t) : \|\widetilde{G}(\mathbf{x}, t)\| \le K_G(\mathbf{x})\}$ is Borel because $\widetilde{G}$ is jointly Borel and K_G is continuous. Hence G is jointly Borel and satisfies Eq. 197.

Agreement almost everywhere. Fix $t \in [\epsilon, T]$. By Lemma C.13, F_t is differentiable off a Lebesgue-null set N_t. By Lemma E.3, the bound $\|\nabla F_t(\mathbf{x})\| \le K_G(\mathbf{x})$ holds off a Lebesgue-null set N_t'. For $\mathbf{x} \notin N_t \cup N_t'$, the candidate satisfies $\|\widetilde{G}(\mathbf{x}, t)\| = \|\nabla F_t(\mathbf{x})\| \le K_G(\mathbf{x})$, so

$$G(\mathbf{x}, t) = \widetilde{G}(\mathbf{x}, t) = \nabla F_t(\mathbf{x}) \tag{200}$$

□

The next lemma applies to any jointly Borel guidance field with linear growth, not only to the field constructed in Lemma E.4.

Lemma E.5. *Assume E.1 and E.2. Let $G : \mathbb{R}^d \times [\epsilon, T] \to \mathbb{R}^d$ be jointly Borel and satisfy, for some $C_G < \infty$,*

$$\|G_t(\mathbf{x})\| \le C_G(1 + \|\mathbf{x}\|), \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{201}$$

Then the guided drift b_t^λ in Eq. 184 is jointly Borel, and there exists $C_b < \infty$ such that

$$\|b_t^\lambda(\mathbf{x})\| \le C_b(1 + \|\mathbf{x}\|), \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{202}$$

Proof. The score component ξ_t in Eq. 185 is jointly continuous by Assumptions E.1 and E.2. The growth bound $\|\hat{\mathbf{x}}_0(\mathbf{x}, t)\| \le H_0 + C_J\|\mathbf{x}\|$ gives

$$\|\xi_t(\mathbf{x})\| \le \frac{\alpha_* H_0 + (\alpha_* C_J + 1)\|\mathbf{x}\|}{\sigma_*^2} \tag{203}$$

Since μ and g are continuous and G is jointly Borel, the drift

$$b_t^\lambda(\mathbf{x}) = \mu(t)\mathbf{x} - g(t)^2\xi_t(\mathbf{x}) + wg(t)^2 G_t(\mathbf{x}) \tag{204}$$

is jointly Borel. Each of its terms has linear growth uniformly in t, because μ and g are bounded on $[\epsilon, T]$. The triangle inequality then proves Eq. 202. □

E.4 Strong Well-Posedness

Let $\nu \in \mathcal{P}(\mathbb{R}^d)$ be an initial distribution. Fix a probability space supporting a standard d-dimensional Brownian motion $(\mathbf{w}_s)_{s\ge 0}$ and an independent initial condition $\mathbf{y}_0 \sim \nu$. Let $(\mathcal{F}_s)_{s\ge 0}$ be the complete, right-continuous augmentation of

$$\mathcal{F}_s^0 := \sigma(\mathbf{y}_0, \mathbf{w}_r : 0 \le r \le s) \tag{205}$$

The guided SDE in Eq. 206 runs backward in t, so we solve it in the sampling time $s := T - t$. A *strong solution* is a continuous, $(\mathcal{F}_s)$-adapted process $(\mathbf{y}_s)_{s\in[0,T-\epsilon]}$ that satisfies the integral form of Eq. 207 almost surely. The generated path is $\mathbf{x}_t := \mathbf{y}_{T-t}$. *Pathwise uniqueness* holds if any two strong solutions driven by the same $(\mathbf{y}_0, \mathbf{w})$ are indistinguishable.

Theorem E.6. *Assume the hypotheses of Lemma E.5, and let b_t^λ be the guided drift in Eq. 184. Set $S := T - \epsilon$ and consider the guided SDE*

$$d\mathbf{x}_t = b_t^\lambda(\mathbf{x}_t)\, dt + g(t)\, d\overline{\mathbf{w}}_t, \qquad t \in [\epsilon, T], \qquad \mathbf{x}_T \sim \nu \tag{206}$$

together with its reparameterization in the sampling time $s := T - t$,

$$\begin{aligned} d\mathbf{y}_s &= \beta(\mathbf{y}_s, s)\, ds + g(T-s)\, d\mathbf{w}_s, \qquad s \in [0, S], \qquad \mathbf{y}_0 \sim \nu \\ \beta(\mathbf{y}, s) &:= -b_{T-s}^\lambda(\mathbf{y}) \end{aligned} \tag{207}$$

Then, for every initial distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$, the guided SDE in Eq. 206 admits a strong solution $\mathbf{y}$ and pathwise uniqueness holds. Furthermore, the distribution of $\mathbf{y}_s$ has a Lebesgue density for every $s \in (0, S]$.

Proof. **Step 1: Truncated equations.** Lemma E.5 gives a jointly Borel drift β with

$$\|\beta(\mathbf{y}, s)\| \le C_b(1 + \|\mathbf{y}\|), \qquad (\mathbf{y}, s) \in \mathbb{R}^d \times [0, S] \tag{208}$$

For each $n \ge 1$, truncate the drift outside the closed ball $\overline{B}_n$ by setting

$$\beta_n(\mathbf{y}, s) := \mathbf{1}\{\|\mathbf{y}\| \le n\}\beta(\mathbf{y}, s) \tag{209}$$

The drift β_n is bounded and jointly Borel. The diffusion matrix $g(T-s)\mathbf{I}$ is continuous, bounded, uniformly elliptic, and independent of $\mathbf{y}$. Extend β_n and $g(T-s)\mathbf{I}$ to $s > S$ by their values at $s = S$. By Veretennikov (1981, Theorem 1 and Section 2, Remark 3), there is a unique strong solution $\mathbf{y}^n$ of

$$d\mathbf{y}_s^n = \beta_n(\mathbf{y}_s^n, s)\, ds + g(T-s)\, d\mathbf{w}_s, \qquad \mathbf{y}_0^n = \mathbf{y}_0 \tag{210}$$

All truncated equations use the same initial condition $\mathbf{y}_0$ and Brownian motion $\mathbf{w}$.

Define the exit times

$$\tau_n \coloneqq \inf\{s \in [0, S] : \|\mathbf{y}_s^n\| \geq n\}, \qquad \inf \varnothing \coloneqq \infty \tag{211}$$

For $m \geq n$, the drifts β_m and β_n agree on $\overline{B}_n$, so $\mathbf{y}^n$ solves the m-th truncated equation up to τ_n. The process $\mathbf{y}^m$ solves the same equation on $[0, S]$, and pathwise uniqueness gives, almost surely,

$$\mathbf{y}^m_{s\wedge\tau_n} = \mathbf{y}^n_{s\wedge\tau_n}, \qquad s \in [0, S] \tag{212}$$

Step 2: Nonexplosion and pathwise uniqueness. Define the continuous martingale

$$\mathbf{M}_s \coloneqq \int_0^s g(T - r)\, d\mathbf{w}_r \tag{213}$$

Since $\|\beta_n\| \leq \|\beta\|$, every β_n satisfies the linear-growth bound in Eq. 208. The integral form of the n-th truncated equation then gives

$$\sup_{0\leq r\leq s} \|\mathbf{y}_r^n\| \leq \|\mathbf{y}_0\| + \sup_{0\leq r\leq S} \|\mathbf{M}_r\| + C_b s + C_b \int_0^s \sup_{0\leq u\leq r} \|\mathbf{y}_u^n\|\, dr \tag{214}$$

Gronwall's inequality yields

$$\sup_{0\leq s\leq S} \|\mathbf{y}_s^n\| \leq \left(\|\mathbf{y}_0\| + C_b S + \sup_{0\leq s\leq S} \|\mathbf{M}_s\|\right) \exp(C_b S) =: R \tag{215}$$

The random variable R is finite almost surely and does not depend on n. Thus $\tau_n = \infty$ whenever $n > R$. By Eq. 212, the sequence $(\mathbf{y}^n)_n$ is eventually constant on $[0, S]$, almost surely. Let $\mathbf{y}$ denote the limit of $(\mathbf{y}^n)_n$. The process $\mathbf{y}$ is continuous and $(\mathcal{F}_s)$-adapted. For $n > R$, the path of $\mathbf{y}^n$ stays in $\overline{B}_n$, where $\beta_n = \beta$, so $\mathbf{y}$ satisfies the integral form of Eq. 207 on $[0, S]$.

Let $\mathbf{y}'$ be any strong solution driven by the same $(\mathbf{y}_0, \mathbf{w})$, with exit times $\rho_n \coloneqq \inf\{s \in [0, S] : \|\mathbf{y}'_s\| \geq n\}$. Since $\beta = \beta_n$ on $\overline{B}_n$, the process $\mathbf{y}'$ solves the n-th truncated equation up to ρ_n. The process $\mathbf{y}^n$ solves the same equation on $[0, S]$, and pathwise uniqueness gives $\mathbf{y}' = \mathbf{y}^n$ on $[0, \rho_n]$ almost surely. For all large n, the bounded path of $\mathbf{y}'$ gives $\rho_n = \infty$, and $n > R$ gives $\mathbf{y}^n = \mathbf{y}$. Thus $\mathbf{y}' = \mathbf{y}$ on $[0, S]$ almost surely. Any two strong solutions are therefore indistinguishable.

Step 3: Positive-time densities. For each n, the process

$$\boldsymbol{\theta}_s^n \coloneqq \frac{\beta_n(\mathbf{y}_s^n, s)}{g(T - s)} \tag{216}$$

is bounded because β_n is bounded and $g \geq g_{\min} > 0$, so Novikov's condition holds. By Girsanov's theorem (Kallenberg, 2002), there is a measure $\mathbb{Q}_n$ equivalent to $\mathbb{P}$ on $\mathcal{F}_S$ under which

$$\mathbf{y}_s^n = \mathbf{y}_0 + \int_0^s g(T - r)\, d\mathbf{w}_r^{\mathbb{Q}_n} \tag{217}$$

where $\mathbf{w}^{\mathbb{Q}_n}$ is an $(\mathcal{F}_s)$-Brownian motion under $\mathbb{Q}_n$, hence independent of the $\mathcal{F}_0$-measurable $\mathbf{y}_0$. Conditional on $\mathbf{y}_0$, the variable $\mathbf{y}_s^n$ is Gaussian under $\mathbb{Q}_n$ with covariance

$$\left(\int_0^s g(T - r)^2\, dr\right) \mathbf{I} \tag{218}$$

Since $g > 0$, the covariance is positive definite for every $s > 0$. Thus $\mathbf{y}_s^n$ has a Lebesgue density under $\mathbb{Q}_n$, and also under $\mathbb{P}$ because $\mathbb{Q}_n$ and $\mathbb{P}$ are equivalent.

Let $N \subset \mathbb{R}^d$ be a Lebesgue-null Borel set and fix $s \in (0, S]$. By Eq. 212, $\mathbf{y}_s = \mathbf{y}_s^n$ on $\{\tau_n > s\}$. Since $\mathbf{y}_s^n$ has a Lebesgue density under $\mathbb{P}$,

$$\mathbb{P}(\mathbf{y}_s \in N,\ \tau_n > s) = \mathbb{P}(\mathbf{y}_s^n \in N,\ \tau_n > s) = 0 \tag{219}$$

Step 2 gives $\mathbb{P}(\tau_n \leq s) \to 0$ as $n \to \infty$, so $\mathbb{P}(\mathbf{y}_s \in N) = 0$. Thus the distribution of $\mathbf{y}_s$ has a Lebesgue density. □

We now show that Theorem E.6 does not depend on the choice of Borel representative.

Corollary E.7. *Let G and G' be two guidance fields satisfying the hypotheses of Lemma E.5, with $G_t = G'_t$ at a.e. $\mathbf{x}$ for every $t \in [\epsilon, T]$. Let $\mathbf{y}$ and $\mathbf{y}'$ be the strong solutions of Eq. 207 with guidance G and G', driven by the same $(\mathbf{y}_0, \mathbf{w})$. Then $\mathbf{y}$ and $\mathbf{y}'$ are indistinguishable. In particular, every jointly Borel representative of ∇F_t with uniform linear growth yields the same solution.*

Proof. Let D be the set on which G and G' differ,

$$D \coloneqq \{(\mathbf{x}, t) \in \mathbb{R}^d \times [\epsilon, T] : G_t(\mathbf{x}) \neq G'_t(\mathbf{x})\} \tag{220}$$

The set D is Borel, and each slice $D_t \coloneqq \{\mathbf{x} : (\mathbf{x}, t) \in D\}$ is Lebesgue-null by assumption. Tonelli's theorem gives

$$\mathbb{E}\left[\int_0^S \mathbf{1}_D(\mathbf{y}_s, T - s)\, ds\right] = \int_0^S \mathbb{P}(\mathbf{y}_s \in D_{T-s})\, ds = 0 \tag{221}$$

The last equality holds because $\mathbf{y}_s$ has a Lebesgue density for every $s > 0$ (Theorem E.6) and each slice D_{T-s} is Lebesgue-null. The time that $(\mathbf{y}_s, T-s)$ spends in D is nonnegative with zero mean, so it vanishes almost surely. Thus the drifts the guidance G and G' agree along $\mathbf{y}$ for a.e. $s \in [0, S]$, and $\mathbf{y}$ also solves Eq. 207 with guidance G'. By pathwise uniqueness $\mathbf{y}$ and $\mathbf{y}'$ are indistinguishable. □

Two caveats qualify the scope of Theorem E.6.

Note: SDE marginals may not equal $(p_t^\lambda)_t$. Theorem E.6 proves that the guided SDE is a well-defined sampler, but its marginals may not equal $(p_t^\lambda)_t$. In general, training-free guidance is biased, and its marginals do not match the target path (Lai et al., 2025; Skreta et al., 2025; Ren et al., 2026). The bias cannot always be removed efficiently because posterior sampling is computationally hard (Gupta et al., 2024).

Note: Well-posedness only on positive-time windows. Theorem E.6 holds on every sampling interval $[\epsilon, T]$ with $\epsilon > 0$, but not at $t = 0$. As $\epsilon \downarrow 0$, both σ_ϵ and λ_ϵ vanish, so the score and the guidance can blow up. Prior works handle the score blow-up by early stopping at $t = \epsilon$ (Bortoli, 2022; Chen et al., 2023b;a; Benton et al., 2024). In our setting, the cost of early stopping is small, since p_ϵ^λ converges to the target as $\epsilon \to 0$ (Appendices D and G.2).

F PROBABILITY-FLOW ODE PROPERTIES

Our PDE experiments guide a flow matching prior (Kerrigan et al., 2024). To analyze the continuous-time sampler, fix $0 < \epsilon < T' \leq T$ and consider the guided probability-flow ODE

$$d\mathbf{x}_t = u_t(\mathbf{x}_t)\, dt, \qquad t \in [\epsilon, T'], \qquad \mathbf{x}_{T'} \sim \nu \tag{222}$$

which carries a seed $\mathbf{x}_{T'}$ to a sample $\mathbf{x}_\epsilon$ along the guided velocity field

$$u_t(\mathbf{x}) \coloneqq \mu(t)\mathbf{x} - \frac{1}{2} g(t)^2 \left(\xi_t(\mathbf{x}) - w G_t(\mathbf{x})\right) \tag{223}$$

Here ξ_t is the score component and G_t is a Borel representative of $\nabla_{\mathbf{x}} f_t^\lambda$. For flow matching, $\mu(t) = \dot{\alpha}_t/\alpha_t$ and $g(t)^2/2$ is the coefficient $c(t)$ in Appendix A.2. We may take $T' = T$ when the assumptions below hold through T. For the linear flow matching path, $\alpha_T = 0$, so we instead take $T' < T$. The distribution ν specifies the initialization at T'.

For Eq. 222 to be a well-defined sampler, the velocity field u_t must have two properties (Lipman et al., 2024, Theorems 1 and 2). The first is a **flow**, so that each seed determines a unique sample and each sample determines its seed. The second is **density transport**, so that the distribution of the sample solves the continuity equation. Both properties follow from the semiconcavity of the Moreau envelope, and hold in the following sense:

1. Every sample has a unique inversion trajectory, and almost every seed has a unique sampling trajectory.
2. The inversion trajectory recovers the original seed from its sample.
3. If the seed distribution has a density, so does the sample distribution throughout sampling, and its evolution is uniquely determined by the continuity equation.

F.1 ASSUMPTIONS

We impose the hypotheses of Lemma E.3 with T' in place of T, and define the constants of Appendix E over $[\epsilon, T']$. Assumption C.3 proves pointwise convergence as $t \to 0$; it is not needed for the positive-time results established here.

Assumption F.1. For each $t \in [\epsilon, T']$, the map $\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable on $\mathbb{R}^d$. Furthermore, there exists a constant $L_J < \infty$ such that

$$\|\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t) - \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{y}, t)\|_{\mathrm{op}} \leq L_J \|\mathbf{x} - \mathbf{y}\|, \qquad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d, \quad t \in [\epsilon, T'] \tag{224}$$

Assumption F.1 extends the Lipschitz bound of Assumption E.2 to a Lipschitz Jacobian. For data supported in $\overline{B}_{R_0}$, the exact denoiser has a Lipschitz Jacobian with constant $L_J = 8\alpha_*^2 R_0^3/\sigma_*^4$. The learned denoiser with fixed conditioning also has a Lipschitz Jacobian, when its blocks use SiLU/GELU activations and stabilized pre-normalization. Examples include DiT (Peebles & Xie, 2023), SiT (Ma et al., 2024), and Stable Diffusion (Rombach et al., 2022; Esser et al., 2024). Under Assumptions E.2 and F.1, $\mathbf{J}_{\hat{\mathbf{x}}_0}$ is jointly Borel with $\|\mathbf{J}_{\hat{\mathbf{x}}_0}\|_{\mathrm{op}} \leq C_J$ everywhere.

Constants. For each radius $R > 0$, define

$$M_{\zeta,R} \coloneqq \frac{A + H_0 + C_J R}{\lambda_{\min}}, \qquad c_{t,R} \coloneqq \frac{C_J^2}{\lambda_t} + M_{\zeta,R} L_J, \qquad t \in [\epsilon, T'] \tag{225}$$

The map $t \mapsto c_{t,R}$ is continuous on $[\epsilon, T']$ because λ is continuous and positive there.

F.2 THE GUIDANCE FIELD

We construct the guidance field G_t and show that it is one-sided Lipschitz on bounded sets, which is the regularity Theorem F.3 requires.

Lemma F.2. *Assume F.1 and the hypotheses of Lemma E.3. There is a jointly Borel map $G : \mathbb{R}^d \times [\epsilon, T'] \to \mathbb{R}^d$, with $G_t \coloneqq G(\cdot, t)$, that satisfies the following for every $t \in [\epsilon, T']$.*

1. *$\|G_t(\mathbf{x})\| \le K_G(\mathbf{x})$ for every $\mathbf{x} \in \mathbb{R}^d$, and $G_t = \nabla F_t$ wherever F_t is differentiable, hence at a.e. $\mathbf{x}$.*

2. *For every $R > 0$*

$$F_t(\mathbf{y}) \le F_t(\mathbf{x}) + \langle G_t(\mathbf{x}), \mathbf{y} - \mathbf{x}\rangle + \frac{c_{t,R}}{2}\|\mathbf{y} - \mathbf{x}\|^2, \qquad \mathbf{x}, \mathbf{y} \in \overline{B}_R \tag{226}$$

3. *For every $R > 0$,*

$$\langle G_t(\mathbf{x}) - G_t(\mathbf{y}), \mathbf{x} - \mathbf{y}\rangle \le c_{t,R}\|\mathbf{x} - \mathbf{y}\|^2, \qquad \mathbf{x}, \mathbf{y} \in \overline{B}_R \tag{227}$$

Proof. **Measurable selection.** The function f is proper, lower semicontinuous, and bounded below (Assumption C.6), and λ is continuous (Assumption E.1). Hence the mapping

$$(\mathbf{u}, t) \mapsto \mathrm{Prox}_{\lambda_t f}(\mathbf{u}) \tag{228}$$

is outer semicontinuous with nonempty compact values (Rockafellar & Wets, 1998, Theorem 1.25), and so Borel measurable (Rockafellar & Wets, 1998, Exercise 14.9). It admits a jointly Borel selection $\mathbf{z}(\mathbf{u}, t) \in \mathrm{Prox}_{\lambda_t f}(\mathbf{u})$ (Rockafellar & Wets, 1998, Corollary 14.6). Define the jointly Borel map

$$\boldsymbol{\zeta}(\mathbf{u}, t) \coloneqq \frac{\mathbf{u} - \mathbf{z}(\mathbf{u}, t)}{\lambda_t} \tag{229}$$

Fix $t \in [\epsilon, T']$ and $\mathbf{u}, \mathbf{u}' \in \mathbb{R}^d$. The point $\mathbf{z}(\mathbf{u}, t)$ attains the infimum $\Psi_f(\mathbf{u}, \lambda_t)$, so the infimum $\Psi_f(\mathbf{u}', \lambda_t)$ satisfies

$$\begin{aligned}\Psi_f(\mathbf{u}', \lambda_t) &\le f(\mathbf{z}(\mathbf{u}, t)) + \frac{\|\mathbf{u}' - \mathbf{z}(\mathbf{u}, t)\|^2}{2\lambda_t} \\ &= \Psi_f(\mathbf{u}, \lambda_t) + \langle \boldsymbol{\zeta}(\mathbf{u}, t), \mathbf{u}' - \mathbf{u}\rangle + \frac{\|\mathbf{u}' - \mathbf{u}\|^2}{2\lambda_t}\end{aligned} \tag{230}$$

Since $\mathbf{u}'$ is arbitrary, $\boldsymbol{\zeta}(\mathbf{u}, t)$ is a supergradient of $\Psi_f(\cdot, \lambda_t)$ at $\mathbf{u}$, and $\Psi_f(\cdot, \lambda_t)$ is semiconcave with constant $1/\lambda_t$.

The proximal displacement bound in Eq. 194 and the denoiser growth bound in Eq. 190 give

$$\|\boldsymbol{\zeta}(\hat{\mathbf{x}}_0(\mathbf{x}, t), t)\| \le \frac{A + H_0 + C_J\|\mathbf{x}\|}{\lambda_t}, \qquad \mathbf{x} \in \mathbb{R}^d \tag{231}$$

Since $\lambda_t \ge \lambda_{\min}$, the right-hand side is at most $M_{\zeta,R}$ on $\overline{B}_R$.

Construction of G. Define

$$G(\mathbf{x}, t) \coloneqq \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)^\top \boldsymbol{\zeta}(\hat{\mathbf{x}}_0(\mathbf{x}, t), t) \tag{232}$$

G is jointly Borel because $\mathbf{J}_{\hat{\mathbf{x}}_0}$ and $\boldsymbol{\zeta}$ are jointly Borel and $\hat{\mathbf{x}}_0$ is jointly continuous. Since $\|\mathbf{J}_{\hat{\mathbf{x}}_0}\|_{\mathrm{op}} \le C_J$, the growth bound in Eq. 231 yields

$$\|G(\mathbf{x}, t)\| \le \frac{C_J}{\lambda_{\min}}\left(A + H_0 + C_J\|\mathbf{x}\|\right) = K_G(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T'] \tag{233}$$

which proves the norm bound in Part 1.

Fix $t \in [\epsilon, T']$. Let $\mathbf{x}$ be a point at which F_t is differentiable. By Assumption F.1, the denoiser $\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable at $\mathbf{x}$. Set $\mathbf{z} = \mathbf{z}(\hat{\mathbf{x}}_0(\mathbf{x}, t), t)$ in Proposition C.14 to obtain

$$\nabla F_t(\mathbf{x}) = \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t)^\top \boldsymbol{\zeta}(\hat{\mathbf{x}}_0(\mathbf{x}, t), t) = G(\mathbf{x}, t) \tag{234}$$

By Lemma C.13, F_t is differentiable almost everywhere. Hence $G_t = \nabla F_t$ almost everywhere, which proves Part 1.

Semiconcavity. For Part 2, fix $R > 0$, $t \in [\epsilon, T']$, and $\mathbf{x}, \mathbf{y} \in \overline{B}_R$. Write

$$\mathbf{u} \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}, t), \qquad \mathbf{u}' \coloneqq \hat{\mathbf{x}}_0(\mathbf{y}, t), \qquad \mathbf{J} \coloneqq \mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}, t) \tag{235}$$

so that

$$F_t(\mathbf{x}) = \Psi_f(\mathbf{u}, \lambda_t), \qquad F_t(\mathbf{y}) = \Psi_f(\mathbf{u}', \lambda_t), \qquad G(\mathbf{x}, t) = \mathbf{J}^\top \boldsymbol{\zeta}(\mathbf{u}, t) \tag{236}$$

Since $\mathbf{J}_{\hat{\mathbf{x}}_0}(\cdot, t)$ is L_J-Lipschitz (Assumption F.1), Taylor's theorem gives

$$\mathbf{u}' - \mathbf{u} = \mathbf{J}(\mathbf{y} - \mathbf{x}) + \boldsymbol{\rho}, \qquad \|\boldsymbol{\rho}\| \le \frac{L_J}{2}\|\mathbf{y} - \mathbf{x}\|^2 \tag{237}$$

Apply Eq. 230 at $(\mathbf{u}, \mathbf{u}')$ and substitute Eq. 237 to obtain

$$F_t(\mathbf{y}) \le F_t(\mathbf{x}) + \langle G(\mathbf{x}, t), \mathbf{y} - \mathbf{x}\rangle + \underbrace{\langle \boldsymbol{\zeta}(\mathbf{u}, t), \boldsymbol{\rho}\rangle}_{\text{Term 1}} + \underbrace{\frac{\|\mathbf{u}' - \mathbf{u}\|^2}{2\lambda_t}}_{\text{Term 2}} \tag{238}$$

Since $\mathbf{x} \in \overline{B}_R$, Eq. 231 and $\lambda_t \ge \lambda_{\min}$ give $\|\boldsymbol{\zeta}(\mathbf{u}, t)\| \le M_{\zeta,R}$. The Cauchy–Schwarz inequality and Eq. 237 then bound Term 1 by

$$\langle \boldsymbol{\zeta}(\mathbf{u}, t), \boldsymbol{\rho}\rangle \le M_{\zeta,R}\|\boldsymbol{\rho}\| \le \frac{M_{\zeta,R} L_J}{2}\|\mathbf{y} - \mathbf{x}\|^2 \tag{239}$$

Assumption E.2 gives $\|\mathbf{u}' - \mathbf{u}\| \le C_J\|\mathbf{y} - \mathbf{x}\|$, which bounds Term 2 by

$$\frac{\|\mathbf{u}' - \mathbf{u}\|^2}{2\lambda_t} \le \frac{C_J^2}{2\lambda_t}\|\mathbf{y} - \mathbf{x}\|^2 \tag{240}$$

By Eq. 225, Term 1 and Term 2 sum to at most $\frac{c_{t,R}}{2}\|\mathbf{y} - \mathbf{x}\|^2$. Hence Eq. 238 gives

$$F_t(\mathbf{y}) \le F_t(\mathbf{x}) + \langle G(\mathbf{x}, t), \mathbf{y} - \mathbf{x}\rangle + \frac{c_{t,R}}{2}\|\mathbf{y} - \mathbf{x}\|^2 \tag{241}$$

which proves Part 2.

Part 2 at $(\mathbf{x}, \mathbf{y})$ and at $(\mathbf{y}, \mathbf{x})$ reads

$$\begin{aligned} F_t(\mathbf{y}) &\le F_t(\mathbf{x}) + \langle G(\mathbf{x}, t), \mathbf{y} - \mathbf{x}\rangle + \frac{c_{t,R}}{2}\|\mathbf{x} - \mathbf{y}\|^2 \\ F_t(\mathbf{x}) &\le F_t(\mathbf{y}) + \langle G(\mathbf{y}, t), \mathbf{x} - \mathbf{y}\rangle + \frac{c_{t,R}}{2}\|\mathbf{x} - \mathbf{y}\|^2 \end{aligned} \tag{242}$$

The sum of the two inequalities is Part 3. □

F.3 Flow and Density Transport

The sampling ODE in Eq. 222 runs backward in t, so we solve it in the sampling time $s \coloneqq T' - t \in [0, S]$, where $S \coloneqq T' - \epsilon$. Let G be the guidance field in Eq. 232, and define the sampling field

$$v(\mathbf{y}, s) \coloneqq -u_t(\mathbf{y}) = -\mu(t)\mathbf{y} + \frac{g(t)^2}{2}\left(\xi_t(\mathbf{y}) - wG(\mathbf{y}, t)\right), \qquad t = T' - s \tag{243}$$

The path $\mathbf{y}_s \coloneqq \mathbf{x}_{T'-s}$ then runs from the seed $\mathbf{y}_0 = \mathbf{x}_{T'} \sim \nu$ to the sample $\mathbf{y}_S = \mathbf{x}_\epsilon$ and solves

$$\dot{\mathbf{y}}_s = v(\mathbf{y}_s, s), \qquad s \in [0, S] \tag{244}$$

The guidance field $G(\cdot, t)$ can be discontinuous, so Eq. 244 may not have classical solutions. Following Lions & Seeger (2024, Section 2), we relax Eq. 244 in Filippov's sense. The Filippov set of v at $(\mathbf{x}, s)$ is the convex hull of the nearby velocities,

$$\mathcal{V}(\mathbf{x}, s) \coloneqq \operatorname{conv}\{\text{limit points of } v(\mathbf{x}', s) \text{ as } \mathbf{x}' \to \mathbf{x}\} \tag{245}$$

and a solution, or Filippov trajectory of v, is an absolutely continuous path $\mathbf{y}$ that satisfies

$$\dot{\mathbf{y}}_s \in \mathcal{V}(\mathbf{y}_s, s), \qquad \text{a.e. } s \in [0, S] \tag{246}$$

If $v(\cdot, s)$ is continuous at $\mathbf{y}_s$, then $\mathcal{V}(\mathbf{y}_s, s) = \{v(\mathbf{y}_s, s)\}$, and Eq. 246 agrees with Eq. 244. At a discontinuity, $\dot{\mathbf{y}}_s$ may be any convex combination of the limit points of v.

Theorem F.3. *Assume the hypotheses of Lemma F.2, and fix a guidance weight $w \ge 0$. There is a constant $C_0 > 0$ such that the sampling field v in Eq. 243 satisfies*

$$\|v(\mathbf{x}, s)\| \le C_0(1 + \|\mathbf{x}\|), \qquad (\mathbf{x}, s) \in \mathbb{R}^d \times [0, S] \tag{247}$$

Furthermore, the following hold.

(a) ***Inversion.*** *For every $s \in (0, S]$ and $\mathbf{x} \in \mathbb{R}^d$, there is a unique solution on $[0, s]$ with $\mathbf{y}_s = \mathbf{x}$. The inversion map $B_s(\mathbf{x}) \coloneqq \mathbf{y}_0$ is Lipschitz on $\overline{B}_R$ for every $R > 0$.*

(b) ***Sampling.*** *There is a Lebesgue-null set $N \subset \mathbb{R}^d$ such that every seed $\mathbf{y}_0 \notin N$ determines a unique solution on $[0, S]$. The sampling map $\Phi(\mathbf{y}_0) \coloneqq \mathbf{y}_S$ is thus defined on $\mathbb{R}^d \setminus N$, and it is Lebesgue measurable.*

(c) ***Seed recovery.*** *Inverting the sample $\Phi(\mathbf{y}_0)$ returns the seed for every $\mathbf{y}_0 \notin N$, so the sampling map is injective on $\mathbb{R}^d \setminus N$,*

$$B_S(\Phi(\mathbf{y}_0)) = \mathbf{y}_0, \qquad \mathbf{y}_0 \notin N \tag{248}$$

(d) ***Density transport.*** *Suppose ν has a density, and let $\mathbf{y}_0 \sim \nu$. Then the distribution μ_s of $\mathbf{y}_s$ has a density for every $s \in [0, S]$, and the curve $s \mapsto \mu_s$ solves the continuity equation*

$$\partial_s \mu + \nabla \cdot (v\mu) = 0, \qquad \mu\big|_{s=0} = \nu \tag{249}$$

as distributions on $(0, S) \times \mathbb{R}^d$. This solution is unique among narrowly continuous curves in $\mathcal{P}(\mathbb{R}^d)$.

Proof. The sampling field v is jointly Borel, since G is jointly Borel by Lemma F.2 and μ, g, and ξ are continuous. Steps 1 and 2 establish linear growth and a local one-sided Lipschitz bound for v. Steps 3 and 4 truncate v and apply the results of Lions & Seeger (2024) to prove (a), (b), and (c). Step 5 proves (d).

Step 1: Linear growth. By Eq. 203 and Part 1 of Lemma F.2, ξ_t and G_t have linear growth uniformly in $t \in [\epsilon, T']$. Since μ and g are bounded on $[\epsilon, T']$, every term of v in Eq. 243 has linear growth, and the triangle inequality proves Eq. 247.

Step 2: One-sided Lipschitz bound. By Eq. 185 and the Lipschitz bound in Assumption E.2,

$$\|\xi_t(\mathbf{x}) - \xi_t(\mathbf{y})\| \le \frac{\alpha_t C_J + 1}{\sigma_t^2} \|\mathbf{x} - \mathbf{y}\|, \qquad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d \tag{250}$$

Fix $R > 0$ and $\mathbf{x}, \mathbf{y} \in \overline{B}_R$, and write $\boldsymbol{\delta} \coloneqq \mathbf{x} - \mathbf{y}$. By Eq. 243,

$$\begin{aligned} \langle v(\mathbf{x}, s) - v(\mathbf{y}, s), \boldsymbol{\delta} \rangle = & -\mu(t)\|\boldsymbol{\delta}\|^2 + \frac{g(t)^2}{2} \underbrace{\langle \xi_t(\mathbf{x}) - \xi_t(\mathbf{y}), \boldsymbol{\delta} \rangle}_{\text{Term 1}} \\ & - \frac{g(t)^2 w}{2} \underbrace{\langle G(\mathbf{x}, t) - G(\mathbf{y}, t), \boldsymbol{\delta} \rangle}_{\text{Term 2}} \end{aligned} \tag{251}$$

Term 1 is bounded below by the Cauchy–Schwarz inequality and Eq. 250; Term 2 is bounded above by Part 3 of Lemma F.2,

$$\text{Term 1} \ge -\frac{\alpha_t C_J + 1}{\sigma_t^2} \|\boldsymbol{\delta}\|^2, \qquad \text{Term 2} \le c_{t,R} \|\boldsymbol{\delta}\|^2 \tag{252}$$

Since $w \ge 0$, the coefficient of Term 2 in Eq. 251 is nonpositive. Substitute Eq. 252 and $-\mu(t) \ge -|\mu(t)|$ into Eq. 251, then maximize the coefficient of $\|\boldsymbol{\delta}\|^2$ over t,

$$\begin{aligned} \langle v(\mathbf{x}, s) - v(\mathbf{y}, s), \mathbf{x} - \mathbf{y} \rangle \ge -L_R \|\mathbf{x} - \mathbf{y}\|^2, \qquad & \mathbf{x}, \mathbf{y} \in \overline{B}_R, \quad s \in [0, S] \\ L_R \coloneqq \max_{t \in [\epsilon, T']} & \left[|\mu(t)| + \frac{g(t)^2}{2} \left(\frac{\alpha_t C_J + 1}{\sigma_t^2} + w c_{t,R} \right) \right] \end{aligned} \tag{253}$$

L_R is finite because μ, g, α, σ, and $t \mapsto c_{t,R}$ are continuous on $[\epsilon, T']$, and $\sigma_t > 0$.

Step 3: Truncated fields. For each integer $n \ge 1$, choose a smooth cutoff $\chi_n : \mathbb{R}^d \to [0, 1]$ equal to one on $\overline{B}_n$ and zero outside B_{2n}. Truncate the sampling field by setting

$$v_n(\mathbf{x}, s) \coloneqq \chi_n(\mathbf{x}) v(\mathbf{x}, s) \tag{254}$$

v_n is jointly Borel, and it satisfies Eq. 247 because $0 \le \chi_n \le 1$. Unlike v, v_n has a one-sided Lipschitz bound on all of $\mathbb{R}^d$,

$$\langle v_n(\mathbf{x}, s) - v_n(\mathbf{y}, s), \mathbf{x} - \mathbf{y} \rangle \ge -\widetilde{L}_n \|\mathbf{x} - \mathbf{y}\|^2, \qquad \mathbf{x}, \mathbf{y} \in \mathbb{R}^d, \quad s \in [0, S]$$
$$\widetilde{L}_n \coloneqq L_{2n} + C_0(1 + 2n)\|\nabla \chi_n\|_\infty \tag{255}$$

To prove Eq. 255, fix $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ and write $\boldsymbol{\delta} \coloneqq \mathbf{x} - \mathbf{y}$. The cutoff satisfies

$$0 \le \chi_n \le 1, \qquad |\chi_n(\mathbf{x}) - \chi_n(\mathbf{y})| \le \|\nabla \chi_n\|_\infty \|\boldsymbol{\delta}\| \tag{256}$$

Case 1: $\mathbf{x}, \mathbf{y} \in \overline{B}_{2n}$. By definition of v_n,

$$\langle v_n(\mathbf{x}, s) - v_n(\mathbf{y}, s), \boldsymbol{\delta} \rangle = \chi_n(\mathbf{x}) \underbrace{\langle v(\mathbf{x}, s) - v(\mathbf{y}, s), \boldsymbol{\delta} \rangle}_{\text{Term 1}} + (\chi_n(\mathbf{x}) - \chi_n(\mathbf{y})) \underbrace{\langle v(\mathbf{y}, s), \boldsymbol{\delta} \rangle}_{\text{Term 2}} \tag{257}$$

Term 1 is bounded below by Eq. 253 with $R = 2n$; Term 2 is bounded above by the Cauchy–Schwarz inequality with Eq. 247,

$$\text{Term 1} \ge -L_{2n}\|\boldsymbol{\delta}\|^2, \qquad |\text{Term 2}| \le C_0(1 + 2n)\|\boldsymbol{\delta}\| \tag{258}$$

Substitute Eq. 258 and Eq. 256 into Eq. 257 to obtain Eq. 255.

Case 2: $\mathbf{x} \notin \overline{B}_{2n}$ *or* $\mathbf{y} \notin \overline{B}_{2n}$. The left-hand side of Eq. 255 is symmetric in $\mathbf{x}$ and $\mathbf{y}$, so assume $\mathbf{y} \notin \overline{B}_{2n}$. Then $\chi_n(\mathbf{y}) = 0$, and Eq. 256 reduces to $\chi_n(\mathbf{x}) \le \|\nabla \chi_n\|_\infty \|\boldsymbol{\delta}\|$. If $\chi_n(\mathbf{x}) \ne 0$, then $\mathbf{x} \in B_{2n}$, so $\|v(\mathbf{x}, s)\| \le C_0(1 + 2n)$ by Eq. 247. In either case,

$$\chi_n(\mathbf{x})\|v(\mathbf{x}, s)\| \le C_0(1 + 2n)\chi_n(\mathbf{x}) \le C_0(1 + 2n)\|\nabla \chi_n\|_\infty \|\boldsymbol{\delta}\| \tag{259}$$

Apply the Cauchy–Schwarz inequality and substitute Eq. 259 to obtain

$$\langle v_n(\mathbf{x}, s) - v_n(\mathbf{y}, s), \boldsymbol{\delta} \rangle = \chi_n(\mathbf{x}) \langle v(\mathbf{x}, s), \boldsymbol{\delta} \rangle \ge -\chi_n(\mathbf{x})\|v(\mathbf{x}, s)\|\|\boldsymbol{\delta}\|$$
$$\ge -C_0(1 + 2n)\|\nabla \chi_n\|_\infty \|\boldsymbol{\delta}\|^2 \ge -\widetilde{L}_n \|\boldsymbol{\delta}\|^2 \tag{260}$$

Cases 1 and 2 prove Eq. 255. By Eq. 247 and Eq. 255, the field v_n satisfies condition (2.1) of Lions & Seeger (2024) with the constants C_0 and $\widetilde{L}_n$. By Lions & Seeger (2024, Lemma 2.2), v_n has a unique backward trajectory on $[0, s]$ that ends at $\mathbf{x}$, for every $s \in (0, S]$ and $\mathbf{x} \in \mathbb{R}^d$. By Lions & Seeger (2024, Theorems 4.8 and 4.12), v_n has a unique forward trajectory $\mathbf{y}$ on $[0, S]$ from almost every seed $\mathbf{y}_0$, which satisfies

$$\mathbf{y}_s = \mathbf{y}_0 + \int_0^s v_n(\mathbf{y}_r, r)\, dr, \qquad s \in [0, S] \tag{261}$$

Write $\varphi^n_{r,s}(\mathbf{x}) \coloneqq \mathbf{y}_r$ for the flow of v_n, where $\mathbf{y}$ is the trajectory of v_n with $\mathbf{y}_s = \mathbf{x}$. For $r \ge s$, the map $\varphi^n_{r,s}$ is defined almost everywhere and Lebesgue measurable (Lions & Seeger, 2024, Theorem 4.8).

Step 4: Localization. Write $\mathcal{V}_n$ for the Filippov set of v_n, defined as in Eq. 245. A path $\mathbf{y}$ *visits* a set K if $\mathbf{y}_s \in K$ at some time s. We show that every Filippov trajectory visiting $\overline{B}_R$ stays in a larger ball, on which $v_n = v$ for large n. Thus, v_n and v have the same Filippov trajectories visiting $\overline{B}_R$; so the conclusions of Step 3 transfer from v_n to v and yield (a), (b), and (c).

Confinement. Fix $R > 0$, set

$$\widehat{R} \coloneqq (1 + R)\exp(C_0 S) - 1 \tag{262}$$

and fix an integer $n_R > \widehat{R}$. Let $\bar{v}$ be either v or v_{n_R}, with Filippov set $\bar{\mathcal{V}}$, and let $\mathbf{y}$ be a Filippov trajectory of $\bar{v}$ on $[0, \tau]$ for an arbitrary $\tau \in (0, S]$. The field $\bar{v}$ satisfies Eq. 247 by Steps 1 and 3. Limits and convex combinations preserve Eq. 247, so

$$\|\mathbf{k}\| \le C_0(1 + \|\mathbf{x}\|), \qquad \mathbf{k} \in \bar{\mathcal{V}}(\mathbf{x}, s), \quad (\mathbf{x}, s) \in \mathbb{R}^d \times [0, S] \tag{263}$$

Since $\dot{\mathbf{y}}_r \in \bar{\mathcal{V}}(\mathbf{y}_r, r)$ for almost every r, Gronwall's inequality gives

$$1 + \|\mathbf{y}_r\| \le (1 + \|\mathbf{y}_s\|)\exp(C_0|r - s|), \qquad r, s \in [0, \tau] \tag{264}$$

Suppose $\mathbf{y}$ visits $\overline{B}_R$. Then Eq. 262 gives

$$\|\mathbf{y}_r\| \le \widehat{R}, \qquad r \in [0, \tau] \tag{265}$$

By Eq. 254, v and v_{n_R} agree on the open ball $B_{n_R} \supset \overline{B}_{\widehat{R}}$. Then by Eq. 245, $\mathcal{V}$ and $\mathcal{V}_{n_R}$ agree on $\overline{B}_{\widehat{R}} \times [0, S]$. By Eq. 265, every Filippov trajectory of v or v_{n_R} visiting $\overline{B}_R$ stays in $\overline{B}_{\widehat{R}}$, where $\mathcal{V}$ and $\mathcal{V}_{n_R}$ agree. Thus v and v_{n_R} have the same Filippov trajectories visiting $\overline{B}_R$.

Part (a). Fix $s \in (0, S]$ and $\mathbf{x} \in \overline{B}_R$. Every backward trajectory on $[0, s]$ ending at $\mathbf{x}$ visits $\overline{B}_R$ at time s. By *Confinement*, v and v_{n_R} have the same backward trajectories ending at $\mathbf{x}$. By Step 3, v_{n_R} has a unique backward trajectory ending at $\mathbf{x}$. Hence $B_s(\mathbf{x}) = \varphi^{n_R}_{0,s}(\mathbf{x})$ for every $\mathbf{x} \in \overline{B}_R$, and the map $\varphi^{n_R}_{0,s}$ is Lipschitz by Lions & Seeger (2024, Lemma 2.2). Since R is arbitrary, (a) holds.

Part (b). Let $R \ge 1$ be an integer. By Step 3, v_{n_R} has a unique forward trajectory on $[0, S]$ from almost every seed. Let N_R be the Lebesgue-null set of remaining seeds. Fix a seed $\mathbf{y}_0 \in \overline{B}_R \setminus N_R$, and let $\mathbf{y}$ be the forward trajectory of v_{n_R} from $\mathbf{y}_0$. By Step 3, $\mathbf{y}$ satisfies Eq. 261, and $\mathbf{y}_s = \varphi^{n_R}_{s,0}(\mathbf{y}_0)$ for every $s \in [0, S]$. Every forward trajectory from $\mathbf{y}_0$ visits $\overline{B}_R$, so by *Confinement*, v and v_{n_R} have the same forward trajectories from $\mathbf{y}_0$. Thus $\mathbf{y}$ is also the unique forward trajectory of v from $\mathbf{y}_0$. By Eq. 265, $\mathbf{y}$ stays in $\overline{B}_{\widehat{R}}$, where $v_{n_R} = v$. So by Eq. 261,

$$\mathbf{y}_s = \mathbf{y}_0 + \int_0^s v(\mathbf{y}_r, r)\, dr, \qquad s \in [0, S] \tag{266}$$

Set

$$N \coloneqq \bigcup_{R \in \mathbb{N}} N_R \tag{267}$$

The set N is Lebesgue-null as a countable union of Lebesgue-null sets. Fix a seed $\mathbf{y}_0 \notin N$. $\mathbf{y}_0$ lies in $\overline{B}_R \setminus N_R$ for some $R \in \mathbb{N}$, so v has a unique forward trajectory $\mathbf{y}$ from $\mathbf{y}_0$, and $\mathbf{y}$ satisfies Eq. 266. Fix $s \in [0, S]$. By Step 3, the map $\varphi^{n_R}_{s,0}$ is Lebesgue measurable, and $\mathbf{y}_s = \varphi^{n_R}_{s,0}(\mathbf{y}_0)$ for every seed $\mathbf{y}_0 \in \overline{B}_R \setminus N_R$. Thus $\mathbf{y}_0 \mapsto \mathbf{y}_s$ is Lebesgue measurable on $\mathbb{R}^d \setminus N$. In particular, Φ is Lebesgue measurable, and (b) holds.

Part (c). Fix a seed $\mathbf{y}_0 \notin N$, and let $\mathbf{y}$ be the forward trajectory of v from $\mathbf{y}_0$. $\mathbf{y}$ ends at $\mathbf{y}_S = \Phi(\mathbf{y}_0)$. By (a), $\mathbf{y}$ is the unique backward trajectory on $[0, S]$ ending at $\Phi(\mathbf{y}_0)$, so $B_S(\Phi(\mathbf{y}_0)) = \mathbf{y}_0$. Thus Φ is injective on $\mathbb{R}^d \setminus N$, and (c) holds.

Step 5: Density transport. Suppose ν has a density, and let $\mathbf{y}_0 \sim \nu$. The set N from (b) is Lebesgue-null, so $\nu(N) = 0$. By (b), the map $\mathbf{y}_0 \mapsto \mathbf{y}_s$ is Lebesgue measurable, so the distribution μ_s of $\mathbf{y}_s$ is well defined for every $s \in [0, S]$.

Density. $\mu_0 = \nu$ has a density by assumption. Fix $s \in (0, S]$ and a Lebesgue-null set $A \subset \mathbb{R}^d$. Fix a seed $\mathbf{y}_0 \notin N$, and let $\mathbf{y}$ be the forward trajectory of v from $\mathbf{y}_0$. The restriction of $\mathbf{y}$ to $[0, s]$ is a backward trajectory ending at $\mathbf{y}_s$, so by (a), $\mathbf{y}_0 = B_s(\mathbf{y}_s)$. Thus every seed $\mathbf{y}_0$ with $\mathbf{y}_s \in A$ lies in N or in $B_s(A)$, so

$$\mu_s(A) = \mathbb{P}(\mathbf{y}_s \in A) \le \nu(N) + \nu(B_s(A)) \tag{268}$$

By (a), B_s is Lipschitz on every ball, so $B_s(A)$ is Lebesgue-null. Since ν has a density, $\nu(N) = \nu(B_s(A)) = 0$. Thus $\mu_s(A) = 0$ for every Lebesgue-null set A, and μ_s has a density.

Continuity equation. Fix $\varphi \in C_c^\infty((0, S) \times \mathbb{R}^d)$. Fix a seed $\mathbf{y}_0 \notin N$, and let $\mathbf{y}$ be the forward trajectory of v from $\mathbf{y}_0$. By Eq. 266, $\dot{\mathbf{y}}_s = v(\mathbf{y}_s, s)$ for almost every s, and the chain rule gives

$$\frac{d}{ds}\varphi(s, \mathbf{y}_s) = \partial_s \varphi(s, \mathbf{y}_s) + \langle \nabla \varphi(s, \mathbf{y}_s), v(\mathbf{y}_s, s) \rangle, \qquad \text{a.e. } s \in [0, S] \tag{269}$$

Integrate Eq. 269 over $[0, S]$. The left-hand side integrates to $\varphi(S, \mathbf{y}_S) - \varphi(0, \mathbf{y}_0) = 0$, because φ has compact support in $(0, S)$. The right-hand side is bounded, because φ has compact support and v satisfies Eq. 247. Since $\nu(N) = 0$, take expectations over $\mathbf{y}_0 \sim \nu$ and apply Fubini's theorem to obtain

$$\int_0^S \int_{\mathbb{R}^d} (\partial_s \varphi + \langle \nabla \varphi, v \rangle)\, d\mu_s\, ds = 0 \tag{270}$$

Thus $s \mapsto \mu_s$ solves the continuity equation in distributions, with $\mu_0 = \nu$. For every bounded continuous ψ, the map $s \mapsto \mathbb{E}[\psi(\mathbf{y}_s)]$ is continuous by dominated convergence, because each path $s \mapsto \mathbf{y}_s$ is continuous. Hence $s \mapsto \mu_s$ is narrowly continuous.

Uniqueness. Every solution of Eq. 244 is a Filippov trajectory of v, because

$$v(\mathbf{x}, s) \in \mathcal{V}(\mathbf{x}, s), \qquad (\mathbf{x}, s) \in \mathbb{R}^d \times [0, S] \tag{271}$$

To prove Eq. 271, fix $(\mathbf{x}, s)$, a unit vector $\mathbf{e}$, and $R > \|\mathbf{x}\| + 1$. For $r \in (0, 1]$, apply Eq. 253 at the points $\mathbf{x} + r\mathbf{e}$ and $\mathbf{x}$, then divide by r to obtain

$$\langle v(\mathbf{x} + r\mathbf{e}, s) - v(\mathbf{x}, s), \mathbf{e} \rangle \geq -L_R r \tag{272}$$

$v(\mathbf{x} + r\mathbf{e}, s)$ is bounded by Eq. 247, so $v(\mathbf{x} + r_j\mathbf{e}, s) \to \mathbf{k}$ along some sequence $r_j \downarrow 0$. By Eq. 245, the limit $\mathbf{k}$ lies in $\mathcal{V}(\mathbf{x}, s)$. Let $r = r_j \downarrow 0$ in Eq. 272 to obtain $\langle v(\mathbf{x}, s), \mathbf{e} \rangle \leq \langle \mathbf{k}, \mathbf{e} \rangle$. Since $\mathbf{e}$ is arbitrary, no hyperplane separates $v(\mathbf{x}, s)$ from the compact convex set $\mathcal{V}(\mathbf{x}, s)$, so Eq. 271 holds.

Let $(\widetilde{\mu}_s)_{s \in [0,S]}$ be a narrowly continuous solution of the continuity equation in (d). By Eq. 247,

$$\int_0^S \int_{\mathbb{R}^d} \frac{\|v(\mathbf{x}, s)\|}{1 + \|\mathbf{x}\|} \, \widetilde{\mu}_s(d\mathbf{x}) \, ds \leq C_0 S \tag{273}$$

The superposition principle (Ambrosio, 2008, Theorem 3.2) gives a random path $\boldsymbol{\gamma}$ in $C([0, S], \mathbb{R}^d)$ such that $\boldsymbol{\gamma}_s$ has distribution $\widetilde{\mu}_s$ for every $s \in [0, S]$. Almost surely, $\boldsymbol{\gamma}$ is an absolutely continuous solution of Eq. 244. By Eq. 271, $\boldsymbol{\gamma}$ is then a Filippov trajectory of v. Since $\boldsymbol{\gamma}_0$ has distribution ν and $\nu(N) = 0$, almost surely $\boldsymbol{\gamma}_0 \notin N$. By (b), $\boldsymbol{\gamma}$ is the unique forward trajectory from $\boldsymbol{\gamma}_0$, so $\boldsymbol{\gamma}_s$ is the image of $\boldsymbol{\gamma}_0$ under the map $\mathbf{y}_0 \mapsto \mathbf{y}_s$. The seeds $\boldsymbol{\gamma}_0$ and $\mathbf{y}_0$ both have distribution ν, so $\widetilde{\mu}_s = \mu_s$ for every $s \in [0, S]$, and (d) holds. □

G PROXIMAL CORRECTION BIAS

For pure constraints $f = \chi_{\mathcal{K}}$, ProxDG's final proximal correction (Section 4.3) is the Euclidean projection onto $\mathcal{K}$. We prove that this projection never inflates the sampler's TV error (Proposition G.3). We then prove that the corrected distribution converges to the target posterior in TV; in the surface regime, the uncorrected path cannot converge in TV (Proposition D.15).

G.1 SAMPLING BIAS OF PROXIMAL CORRECTION

We fix a Borel selection of the nearest-point projection.

Lemma G.1. *Let $\mathcal{K} \subseteq \mathbb{R}^d$ be nonempty and closed, and write $d_{\mathcal{K}}(\mathbf{x}) \coloneqq \inf_{\mathbf{z}\in\mathcal{K}} \|\mathbf{x}-\mathbf{z}\|$. There is a Borel map $P_{\mathcal{K}} : \mathbb{R}^d \to \mathcal{K}$ such that*

$$\|\mathbf{x} - P_{\mathcal{K}}(\mathbf{x})\| = d_{\mathcal{K}}(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^d \tag{274}$$

Furthermore, $P_{\mathcal{K}}$ is the identity on $\mathcal{K}$. If $d_{\mathcal{K}}(\mathbf{x}) < \operatorname{reach}(\mathcal{K})$, then $P_{\mathcal{K}}(\mathbf{x})$ is the unique nearest point of $\mathbf{x}$ in $\mathcal{K}$.

Proof. Define the nearest-point mapping

$$M(\mathbf{x}) \coloneqq \{\mathbf{z} \in \mathcal{K} : \|\mathbf{x}-\mathbf{z}\| = d_{\mathcal{K}}(\mathbf{x})\}, \qquad \mathbf{x} \in \mathbb{R}^d \tag{275}$$

Each $M(\mathbf{x})$ is nonempty and compact, since $\mathcal{K}$ is nonempty and closed and $\mathbf{z} \mapsto \|\mathbf{x}-\mathbf{z}\|$ is continuous and coercive. Suppose $\mathbf{x}_j \to \mathbf{x}$ and $\mathbf{z}_j \to \mathbf{z}$ with $\mathbf{z}_j \in M(\mathbf{x}_j)$. Then $\mathbf{z} \in \mathcal{K}$ because $\mathcal{K}$ is closed, and continuity of $d_{\mathcal{K}}$ gives

$$\|\mathbf{x}-\mathbf{z}\| = \lim_{j\to\infty} \|\mathbf{x}_j - \mathbf{z}_j\| = \lim_{j\to\infty} d_{\mathcal{K}}(\mathbf{x}_j) = d_{\mathcal{K}}(\mathbf{x}) \tag{276}$$

Hence $\mathbf{z} \in M(\mathbf{x})$, so M is outer semicontinuous. A closed-valued outer semicontinuous mapping is Borel measurable (Rockafellar & Wets, 1998, Exercise 14.9), and a closed-valued measurable mapping admits a measurable selection (Rockafellar & Wets, 1998, Corollary 14.6). Let $P_{\mathcal{K}}$ be a Borel selection of M.

If $\mathbf{x} \in \mathcal{K}$, then $M(\mathbf{x}) = \{\mathbf{x}\}$ and $P_{\mathcal{K}}(\mathbf{x}) = \mathbf{x}$. If $d_{\mathcal{K}}(\mathbf{x}) < \operatorname{reach}(\mathcal{K})$, then $M(\mathbf{x})$ is a singleton by the definition of reach (Federer, 1959, Definition 4.1), so $P_{\mathcal{K}}(\mathbf{x})$ is the unique nearest point. □

Lemma G.2. *Let ν and p be probability measures on $\mathbb{R}^d$ and let $T : \mathbb{R}^d \to \mathbb{R}^d$ be Borel. Then $\|T_{\#}\nu - T_{\#}p\|_{\mathrm{TV}} \le \|\nu - p\|_{\mathrm{TV}}$.*

Proof. Fix a Borel set $A \subseteq \mathbb{R}^d$ and recall that $T_{\#}\nu(A) = \nu(T^{-1}(A))$. Since $T^{-1}(A)$ is itself Borel,

$$|T_{\#}\nu(A) - T_{\#}p(A)| = |\nu(T^{-1}(A)) - p(T^{-1}(A))| \le \|\nu - p\|_{\mathrm{TV}} \tag{277}$$

Take the supremum over A to obtain $\|T_{\#}\nu - T_{\#}p\|_{\mathrm{TV}} \le \|\nu - p\|_{\mathrm{TV}}$. □

In practice, the sampler stops at a small time $\epsilon > 0$. Due to score approximation and discretization error, the final iterate follows some distribution μ_ϵ rather than the ideal marginal p_ϵ^λ. By Lemma G.2, projecting both distributions onto $\mathcal{K}$ does not inflate the TV gap $\|\mu_\epsilon - p_\epsilon^\lambda\|_{\mathrm{TV}}$.

Proposition G.3. *Let ν, p, and π be probability measures on $\mathbb{R}^d$, and let $T : \mathbb{R}^d \to \mathbb{R}^d$ be Borel. Then*

$$\|T_{\#}\nu - \pi\|_{\mathrm{TV}} \le \|\nu - p\|_{\mathrm{TV}} + \|T_{\#}p - \pi\|_{\mathrm{TV}} \tag{278}$$

Proof. By the triangle inequality,

$$\|T_{\#}\nu - \pi\|_{\mathrm{TV}} \le \|T_{\#}\nu - T_{\#}p\|_{\mathrm{TV}} + \|T_{\#}p - \pi\|_{\mathrm{TV}} \tag{279}$$

Lemma G.2 bounds $\|T_{\#}\nu - T_{\#}p\|_{\mathrm{TV}}$ by $\|\nu - p\|_{\mathrm{TV}}$. □

Let $\pi^\dagger$ denote the target, which is p_{tar} in the volume regime and π_{tar} in the surface regime. Apply Eq. 278 with $\nu = \mu_\epsilon$, $p = p_\epsilon^\lambda$, $\pi = \pi^\dagger$, and $T = P_{\mathcal{K}}$ to obtain

$$\|(P_{\mathcal{K}})_{\#}\mu_\epsilon - \pi^\dagger\|_{\mathrm{TV}} \le \underbrace{\|\mu_\epsilon - p_\epsilon^\lambda\|_{\mathrm{TV}}}_{\text{sampler gap}} + \underbrace{\|(P_{\mathcal{K}})_{\#}p_\epsilon^\lambda - \pi^\dagger\|_{\mathrm{TV}}}_{\text{early stopping}} \tag{280}$$

The sampler gap is the implementation error accumulated before the correction. The early-stopping term compares the corrected ideal path to the target, and Appendix G.2 proves that it vanishes as $\epsilon \to 0$.

G.2 TV Convergence of the Corrected Distribution

Define the corrected distribution

$$\hat{p}_t \coloneqq (P_{\mathcal{K}})_{\#}p_t^\lambda, \qquad t \in (0, T] \tag{281}$$

In the volume regime, $\mathcal{K} = \mathcal{C}$ for a closed set $\mathcal{C}$ containing $\mathrm{dom}(f)$. In the surface regime, $\mathcal{K}$ is a closed set with $\mathrm{dom}(f) \subseteq \mathcal{K} \subseteq \mathcal{S}$.

Proposition G.4. *Assume the hypotheses of Proposition D.9. In the volume regime, $\|\hat{p}_t - p_{\mathrm{tar}}\|_{\mathrm{TV}} \to 0$ as $t \to 0$.*

Proof. The density $p_{\mathrm{tar}} \propto q_0 \exp(-f)$ vanishes outside $\mathrm{dom}(f)$, so $p_{\mathrm{tar}}(\mathcal{C}) = 1$. The projection $P_{\mathcal{C}}$ is the identity on $\mathcal{C}$ by Lemma G.1, which gives $(P_{\mathcal{C}})_{\#}p_{\mathrm{tar}} = p_{\mathrm{tar}}$. Lemma G.2 and Proposition D.9 then give

$$\|\hat{p}_t - p_{\mathrm{tar}}\|_{\mathrm{TV}} = \|(P_{\mathcal{C}})_{\#}p_t^\lambda - (P_{\mathcal{C}})_{\#}p_{\mathrm{tar}}\|_{\mathrm{TV}} \le \|p_t^\lambda - p_{\mathrm{tar}}\|_{\mathrm{TV}} \to 0 \tag{282}$$

□

Theorem G.5. *Assume the hypotheses of Theorem D.18. In the surface regime, $\|\hat{p}_t - \pi_{\mathrm{tar}}\|_{\mathrm{TV}} \to 0$ as $t \to 0$.*

Proof. Recall the tube notation of Appendix D.4.1. Set $A \coloneqq U_{r_0} \cap P_{\mathcal{S}}^{-1}(\mathcal{K})$, which is Borel because $P_{\mathcal{S}}$ is Borel and $\mathcal{K}$ is closed. Split $\hat{p}_t$ according to whether the sample lies in A,

$$\hat{p}_t = \hat{p}_t^{\mathrm{in}} + \hat{p}_t^{\mathrm{out}}, \qquad \hat{p}_t^{\mathrm{in}} \coloneqq (P_{\mathcal{K}})_{\#}(p_t^\lambda \llcorner A), \qquad \hat{p}_t^{\mathrm{out}} \coloneqq (P_{\mathcal{K}})_{\#}(p_t^\lambda \llcorner A^c) \tag{283}$$

where $p_t^\lambda \llcorner A$ denotes the restriction of p_t^λ to A. We show that the surface density of $\hat{p}_t^{\mathrm{in}}$ converges to the surface density of π_{tar}, while $\hat{p}_t^{\mathrm{out}}$ loses all mass.

Step 1: Surface density of the tube. By Lemma D.12, every $\mathbf{x} \in U_{r_0}$ decomposes uniquely as $\mathbf{x} = \mathbf{y} + \mathbf{v}$ with $\mathbf{y} = P_{\mathcal{S}}(\mathbf{x})$, $\mathbf{v} \in N_{\mathbf{y}}\mathcal{S}$, and $\|\mathbf{v}\| < r_0$. Suppose $\mathbf{x} \in A$, so that $\mathbf{y} \in \mathcal{K}$. Since $\mathcal{K} \subseteq \mathcal{S}$,

$$d_{\mathcal{S}}(\mathbf{x}) \le d_{\mathcal{K}}(\mathbf{x}) \le \|\mathbf{x} - \mathbf{y}\| = d_{\mathcal{S}}(\mathbf{x}) \tag{284}$$

Lemma G.1 gives $\|\mathbf{x} - P_{\mathcal{K}}(\mathbf{x})\| = d_{\mathcal{K}}(\mathbf{x})$, so Eq. 284 makes $P_{\mathcal{K}}(\mathbf{x})$ a nearest point of $\mathbf{x}$ in $\mathcal{S}$. The nearest point of $\mathbf{x}$ in $\mathcal{S}$ is unique because $d_{\mathcal{S}}(\mathbf{x}) < r_0 < \mathrm{reach}(\mathcal{S})$. Hence $P_{\mathcal{K}}(\mathbf{y} + \mathbf{v}) = \mathbf{y}$. $P_{\mathcal{K}}(\mathbf{y} + \mathbf{v})$ does not depend on $\mathbf{v}$, so $\hat{p}_t^{\mathrm{in}}$ is the $\mathbf{y}$-marginal of $p_t^\lambda \llcorner A$. The tube formula (Eq. 120) then computes the surface density of $\hat{p}_t^{\mathrm{in}}$ as $\mathbf{1}_{\mathcal{K}} h_t$, where

$$h_t(\mathbf{y}) \coloneqq \frac{1}{Z_t^{\mathrm{prox}}} \int_{\{\|\mathbf{v}\| < r_0\}} q_t(\mathbf{y} + \mathbf{v}) \exp(-f_t^\lambda(\mathbf{y} + \mathbf{v})) J(\mathbf{y}, \mathbf{v})\, \mathcal{L}^m(d\mathbf{v}), \qquad \mathbf{y} \in \mathcal{S} \tag{285}$$

Apply the tube formula (Eq. 120) to $p_t^\lambda \llcorner (U_{r_0} \setminus A)$ as well,

$$\hat{p}_t^{\mathrm{out}}(\mathbb{R}^d) = p_t^\lambda(A^c) = p_t^\lambda(U_{r_0}^c) + \int_{\mathcal{S} \setminus \mathcal{K}} h_t \, d\mathcal{H}^n \tag{286}$$

Step 2: Convergence of the surface densities. Let F_t be the rescaled tube integrand in Eq. 157 with $\varphi \equiv 1$, and let F_0 be the pointwise limit of F_t in Eq. 159. Set $c_t \coloneqq (2\pi\lambda_t)^{-m/2} Z_t^{\mathrm{prox}}$. Rescale the fibre in Eq. 285 by $\mathbf{v} = \sqrt{\lambda_t}\mathbf{w}$ to obtain

$$h_t(\mathbf{y}) = \frac{(2\pi)^{-m/2}}{c_t} \int_{N_{\mathbf{y}}\mathcal{S}} F_t(\mathbf{y}, \mathbf{w})\, \mathcal{L}^m(d\mathbf{w}), \qquad \mathbf{y} \in \mathcal{S} \tag{287}$$

Integrate the factor $\exp(-\|\mathbf{w}\|^2/2)$ of F_0 over $N_{\mathbf{y}}\mathcal{S}$ to obtain

$$h_\pi(\mathbf{y}) \coloneqq \frac{\exp(-g(\mathbf{y}))q_0(\mathbf{y})}{Z_{\mathrm{surf}}} = \frac{(2\pi)^{-m/2}}{Z_{\mathrm{surf}}} \int_{N_{\mathbf{y}}\mathcal{S}} F_0(\mathbf{y}, \mathbf{w})\, \mathcal{L}^m(d\mathbf{w}), \qquad \mathbf{y} \in \mathcal{S} \tag{288}$$

h_π is the surface density of π_{tar} in Eq. 136.

Part 1 of Theorem D.18 gives $c_t \to Z_{\mathrm{surf}} > 0$, so $F_t/c_t \to F_0/Z_{\mathrm{surf}}$ pointwise on $\mathcal{N}$. Furthermore, $c_t \geq Z_{\mathrm{surf}}/2$ for sufficiently small t, so Eq. 166 bounds $|F_t/c_t|$ and $|F_0/Z_{\mathrm{surf}}|$ by a constant times K_1. The dominator K_1 is μ-integrable by the proof of Proposition D.17; thus, dominated convergence on $\mathcal{N}$ yields

$$\int_{\mathcal{S}} |h_t - h_\pi|\, d\mathcal{H}^n \leq (2\pi)^{-m/2} \int_{\mathcal{N}} \left| \frac{F_t}{c_t} - \frac{F_0}{Z_{\mathrm{surf}}} \right| d\mu \to 0 \tag{289}$$

Step 3: Assembly. By Steps 1 and 2, the measures $\hat{p}_t^{\mathrm{in}}$ and π_{tar} have surface densities $\mathbf{1}_{\mathcal{K}} h_t$ and h_π, and $\hat{p}_t^{\mathrm{out}}$ has the total mass in Eq. 286. Since $\mathrm{dom}(f) \subseteq \mathcal{K}$, the density h_π vanishes on $\mathcal{S} \setminus \mathcal{K}$. Hence

$$\|\hat{p}_t - \pi_{\mathrm{tar}}\|_{\mathrm{TV}} \leq \underbrace{\int_{\mathcal{S}} |h_t - h_\pi|\, d\mathcal{H}^n}_{\text{Term 1}} + \underbrace{p_t^\lambda(U_{r_0}^c)}_{\text{Term 2}} \tag{290}$$

Term 1 tends to zero by Eq. 289. Term 2 tends to zero by Eq. 156 with $\varphi \equiv 1$, since $c_t \to Z_{\mathrm{surf}} > 0$. □

H Inexact Proximal Solves

Appendices C, D, E, and G assume exact proximal solves, but an iterative solver returns only an approximate proximal point. We show that their regularity and sampling guarantees still hold under inexact proximal solves, and we bound the resulting error of the guided SDE in path space.

H.1 Roadmap

A *solver* is a map $\hat{\mathbf{z}} : \mathbb{R}^d \times (0,\infty) \to \mathbb{R}^d$ that returns an approximate stationary point $\hat{\mathbf{z}}(\mathbf{u},\lambda)$ of the proximal objective

$$\varphi_{\mathbf{u},\lambda}(\mathbf{z}) \coloneqq f(\mathbf{z}) + \frac{1}{2\lambda}\|\mathbf{u}-\mathbf{z}\|^2 \tag{291}$$

The output $\hat{\mathbf{z}}(\mathbf{u},\lambda)$ can differ from the exact proximal point in two ways. *Branch error* occurs when f is nonconvex and the solver converges toward a stationary point of $\varphi_{\mathbf{u},\lambda}$ other than the proximal point. *Numerical error* occurs because the solver stops after finitely many inner iterations, so its output is only approximately stationary.

The two errors demand different tools. We bound the effect of numerical error in path space, and the bound shrinks with the solver tolerance (Theorem H.12). Branch error may not vanish with the solver tolerance and admits no analogous bound. Instead, we replace the Moreau envelope with the *branch envelope*: the proximal objective evaluated at the solver output. The branch envelope defines its own guidance function and probability path (Eq. 295). We prove the resulting path inherits the continuity and convergence guarantees of Appendix D, and it has no branch error by construction.

H.2 Assumptions

Write $\hat{V}(\mathbf{u},\lambda) \coloneqq \varphi_{\mathbf{u},\lambda}(\hat{\mathbf{z}}(\mathbf{u},\lambda))$ for the branch envelope, and set $\lambda_{\max} \coloneqq \sup_{t\in(0,T]} \lambda_t$, which is finite by Assumption C.1.

Assumption H.1. The map $(\mathbf{u},\lambda) \mapsto \hat{\mathbf{z}}(\mathbf{u},\lambda)$ is jointly Borel measurable, with feasible output $\hat{\mathbf{z}}(\mathbf{u},\lambda) \in \mathrm{dom}(f)$ for every $(\mathbf{u},\lambda)$.

Assumption H.2. There exists $L \geq 0$ such that

$$\|\mathbf{u}-\hat{\mathbf{z}}(\mathbf{u},\lambda)\| \leq \underbrace{\mathrm{dist}(\mathbf{u},\mathrm{dom}(f))}_{\text{feasibility}} + \underbrace{\lambda L(1+\|\mathbf{u}\|)}_{\text{optimization}} \tag{292}$$

for all $\mathbf{u} \in \mathbb{R}^d$ and $\lambda \in (0,\lambda_{\max}]$.

Assumption H.3. The restriction of f to $\mathrm{dom}(f)$ is continuous.

Assumption H.4. For each $t^* \in (0,T]$, the branch envelope $\hat{V}$ is continuous at $(\hat{\mathbf{x}}_0(\mathbf{x},t^*),\lambda_{t^*})$ for Lebesgue-a.e. $\mathbf{x} \in \mathbb{R}^d$.

Assumption H.5. There exists a bounded Borel measurable $\eta : (0,T] \to [0,\infty)$ such that the solver output is an η_t-approximate stationary point of the proximal objective. That is, for every $\mathbf{u} \in \mathbb{R}^d$ and $t \in (0,T]$,

$$\mathrm{dist}\left(\mathbf{0}, \partial\varphi_{\mathbf{u},\lambda_t}(\hat{\mathbf{z}}(\mathbf{u},\lambda_t))\right) \leq \eta_t \tag{293}$$

where $\partial\varphi_{\mathbf{u},\lambda}(\mathbf{z}) = \partial f(\mathbf{z}) + \frac{1}{\lambda}(\mathbf{z}-\mathbf{u})$ and ∂f is the limiting subdifferential (Rockafellar & Wets, 1998).

Assumption H.6. For each $\lambda \in (0,\lambda_{\max}]$, the map $\hat{\mathbf{z}}(\cdot,\lambda)$ is differentiable at Lebesgue-a.e. $\mathbf{u} \in \mathbb{R}^d$, with $\|\mathbf{J}_{\hat{\mathbf{z}}}(\mathbf{u},\lambda)\|_{\mathrm{op}} \leq M_{\hat{\mathbf{z}}}$ wherever the Jacobian exists.

The feasibility requirement in Assumption H.1 is trivial for soft objectives and enforced by projections for hard constraints. Assumption H.2 states that the solver output lies at most $\lambda L(1+\|\mathbf{u}\|)$ farther from $\mathbf{u}$ than the nearest feasible point. Every guarantee except surface-regime convergence also holds under the weaker bound

$$\|\mathbf{u}-\hat{\mathbf{z}}(\mathbf{u},\lambda)\| \leq \underbrace{k\cdot\mathrm{dist}(\mathbf{u},\mathrm{dom}(f))}_{\text{feasibility}} + \underbrace{\lambda L(1+\|\mathbf{u}\|)}_{\text{optimization}} \tag{294}$$

for some $k \geq 1$. Assumption H.3 requires continuity along feasible sequences. If $\mathbf{z}_j \to \mathbf{z}$ with $\mathbf{z}_j, \mathbf{z} \in \mathrm{dom}(f)$, then $f(\mathbf{z}_j) \to f(\mathbf{z})$. Assumption H.4 permits discontinuities of the branch envelope,

provided their preimages under $\mathbf{x} \mapsto (\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t)$ are Lebesgue-null at each time t. For a piecewise continuous solver, Assumption H.4 holds when the preimages of the branch boundaries are Lebesgue-null.

Assumptions H.5 and H.6 control the gap between the implemented guidance and the gradient of the branch envelope. Assumption H.5 is the standard inexactness criterion for proximal methods (Ehrhardt et al., 2024). Assumption H.6 holds for a piecewise Lipschitz solver with constant $M_{\hat{\mathbf{z}}}$ on each piece, by Rademacher's theorem.

H.3 The Branch Envelope

We compose the branch envelope with Tweedie's estimator to obtain the realized guidance function $\hat{f}_t^\lambda$ and the realized probability path $\hat{p}_t^\lambda$,

$$\hat{f}_t^\lambda(\mathbf{x}) \coloneqq \hat{V}\left(\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t\right), \qquad \hat{p}_t^\lambda(\mathbf{x}) \propto q_t(\mathbf{x}) \exp\left(-\hat{f}_t^\lambda(\mathbf{x})\right) \tag{295}$$

The solver output also defines the realized displacement $\hat{\Delta}_t$ and the realized guidance field $\hat{G}_t$,

$$\hat{\mathbf{z}}_t \coloneqq \hat{\mathbf{z}}\left(\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t\right), \qquad \hat{\Delta}_t \coloneqq \frac{1}{\lambda_t}\left(\hat{\mathbf{x}}_0 - \hat{\mathbf{z}}_t\right), \qquad \hat{G}_t \coloneqq \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \hat{\Delta}_t \tag{296}$$

The realized sampler integrates the guided SDE (Eq. 12) with $\hat{G}_t$ in place of ∇f_t^λ.

Lemma H.7. *Assume C.6 and H.1. Then*

$$C_f \le f_t^\lambda(\mathbf{x}) \le \hat{f}_t^\lambda(\mathbf{x}) < \infty, \qquad (\mathbf{x}, t) \in \mathbb{R}^d \times (0, T] \tag{297}$$

Proof. The lower bound holds because the proximal objective is at least C_f at every point. The middle inequality holds because $f_t^\lambda(\mathbf{x})$ is the minimum of the proximal objective, whereas $\hat{f}_t^\lambda(\mathbf{x})$ evaluates the objective at $\hat{\mathbf{z}}_t$. The upper bound holds because $f(\hat{\mathbf{z}}_t) < \infty$ by Assumption H.1. □

Suppose $f = g + \chi_{\mathcal{C}}$ for a finite soft objective g and a closed constraint set $\mathcal{C}$; set $\mathcal{C} = \mathbb{R}^d$ if f has no constraints. Let $N_{\mathcal{C}}$ and $T_{\mathcal{C}}$ be the limiting normal cone and the tangent cone of $\mathcal{C}$. The set $\mathcal{C}$ is *Clarke regular* at a point $\mathbf{z} \in \mathcal{C}$ if $N_{\mathcal{C}}(\mathbf{z})$ and $T_{\mathcal{C}}(\mathbf{z})$ are polar (Rockafellar & Wets, 1998). Convex sets and C^2 submanifolds are Clarke regular everywhere.

Lemma H.8. *Fix $\lambda > 0$ and $\mathbf{u} \in \mathbb{R}^d$, and write $\hat{\mathbf{z}} \coloneqq \hat{\mathbf{z}}(\mathbf{u}, \lambda)$. Suppose $\hat{\mathbf{z}}(\cdot, \lambda)$ takes values in $\mathcal{C}$ and is differentiable at $\mathbf{u}$ with Jacobian $\mathbf{J}_{\hat{\mathbf{z}}}$, and g is differentiable at $\hat{\mathbf{z}}$. Then $\hat{V}(\cdot, \lambda)$ is differentiable at $\mathbf{u}$ with*

$$\nabla_{\mathbf{u}} \hat{V}(\mathbf{u}, \lambda) = \frac{1}{\lambda}(\mathbf{u} - \hat{\mathbf{z}}) + \mathbf{J}_{\hat{\mathbf{z}}}^\top \mathbf{s}, \qquad \mathbf{s} \coloneqq \nabla g(\hat{\mathbf{z}}) + \frac{1}{\lambda}(\hat{\mathbf{z}} - \mathbf{u}) \tag{298}$$

If $\mathcal{C}$ is Clarke regular at $\hat{\mathbf{z}}$, then $\mathbf{J}_{\hat{\mathbf{z}}}^\top \mathbf{n} = \mathbf{0}$ for every $\mathbf{n} \in N_{\mathcal{C}}(\hat{\mathbf{z}})$. Furthermore, if g is continuously differentiable on a neighborhood of $\hat{\mathbf{z}}$, then

$$\left\| \nabla_{\mathbf{u}} \hat{V}(\mathbf{u}, \lambda) - \frac{1}{\lambda}(\mathbf{u} - \hat{\mathbf{z}}) \right\| \le \|\mathbf{J}_{\hat{\mathbf{z}}}\|_{\mathrm{op}} \mathrm{dist}\left(\mathbf{0}, \partial \varphi_{\mathbf{u}, \lambda}(\hat{\mathbf{z}})\right) \tag{299}$$

Proof. **Decomposition.** The solver output lies in $\mathcal{C}$, so the indicator vanishes and

$$\hat{V}(\mathbf{u}', \lambda) = g(\hat{\mathbf{z}}(\mathbf{u}', \lambda)) + \frac{1}{2\lambda}\|\mathbf{u}' - \hat{\mathbf{z}}(\mathbf{u}', \lambda)\|^2, \qquad \mathbf{u}' \in \mathbb{R}^d \tag{300}$$

Since g is differentiable at $\hat{\mathbf{z}}$ and $\hat{\mathbf{z}}(\cdot, \lambda)$ is differentiable at $\mathbf{u}$, the chain rule gives

$$\nabla_{\mathbf{u}} \hat{V}(\mathbf{u}, \lambda) = \mathbf{J}_{\hat{\mathbf{z}}}^\top \nabla g(\hat{\mathbf{z}}) + \frac{1}{\lambda}(\mathbf{I} - \mathbf{J}_{\hat{\mathbf{z}}})^\top (\mathbf{u} - \hat{\mathbf{z}}) \tag{301}$$

which rearranges to Eq. 298.

Normal cone. Fix $\mathbf{w} \in \mathbb{R}^d$. The points $\hat{\mathbf{z}}(\mathbf{u} + h\mathbf{w}, \lambda)$ lie in $\mathcal{C}$ for every $h > 0$, so the directional derivative

$$\mathbf{J}_{\hat{\mathbf{z}}} \mathbf{w} = \lim_{h \downarrow 0} \frac{\hat{\mathbf{z}}(\mathbf{u} + h\mathbf{w}, \lambda) - \hat{\mathbf{z}}}{h} \tag{302}$$

lies in the tangent cone $T_{\mathcal{C}}(\hat{\mathbf{z}})$. Hence the range of $\mathbf{J}_{\hat{\mathbf{z}}}$ is a linear subspace contained in $T_{\mathcal{C}}(\hat{\mathbf{z}})$. By Clarke regularity, every $\mathbf{n} \in N_{\mathcal{C}}(\hat{\mathbf{z}})$ satisfies $\langle \mathbf{n}, \mathbf{v} \rangle \leq 0$ for all $\mathbf{v} \in T_{\mathcal{C}}(\hat{\mathbf{z}})$. Take $\mathbf{v} = \pm \mathbf{J}_{\hat{\mathbf{z}}} \mathbf{w}$ to obtain $\langle \mathbf{n}, \mathbf{J}_{\hat{\mathbf{z}}} \mathbf{w} \rangle = 0$. Since $\mathbf{w}$ is arbitrary, $\mathbf{J}_{\hat{\mathbf{z}}}^\top \mathbf{n} = \mathbf{0}$.

Residual bound. Since g is continuously differentiable near $\hat{\mathbf{z}}$, the sum rule (Rockafellar & Wets, 1998, Exercise 8.8) gives $\partial \varphi_{\mathbf{u},\lambda}(\hat{\mathbf{z}}) = \mathbf{s} + N_{\mathcal{C}}(\hat{\mathbf{z}})$. The cone $N_{\mathcal{C}}(\hat{\mathbf{z}})$ is closed, so some $\mathbf{n} \in N_{\mathcal{C}}(\hat{\mathbf{z}})$ attains $\|\mathbf{s} + \mathbf{n}\| = \operatorname{dist}(\mathbf{0}, \partial \varphi_{\mathbf{u},\lambda}(\hat{\mathbf{z}}))$. Since $\mathcal{C}$ is Clarke regular at $\hat{\mathbf{z}}$, we have $\mathbf{J}_{\hat{\mathbf{z}}}^\top \mathbf{n} = \mathbf{0}$, so $\mathbf{J}_{\hat{\mathbf{z}}}^\top \mathbf{s} = \mathbf{J}_{\hat{\mathbf{z}}}^\top (\mathbf{s} + \mathbf{n})$. Substitute into Eq. 298 to obtain

$$\left\| \nabla_{\mathbf{u}} \hat{V}(\mathbf{u}, \lambda) - \frac{1}{\lambda}(\mathbf{u} - \hat{\mathbf{z}}) \right\| = \|\mathbf{J}_{\hat{\mathbf{z}}}^\top (\mathbf{s} + \mathbf{n})\| \leq \|\mathbf{J}_{\hat{\mathbf{z}}}\|_{\mathrm{op}} \|\mathbf{s} + \mathbf{n}\| \tag{303}$$

which is Eq. 299. □

Corollary H.9. *In the setting of Lemma H.8, suppose Eq. 299 holds and $\mathbf{0} \in \partial \varphi_{\mathbf{u},\lambda}(\hat{\mathbf{z}})$. Then*

$$\nabla_{\mathbf{u}} \hat{V}(\mathbf{u}, \lambda) = \frac{1}{\lambda}(\mathbf{u} - \hat{\mathbf{z}}) \tag{304}$$

In particular, suppose $\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable at $\mathbf{x}$ and take

$$(\mathbf{u}, \lambda) = (\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t) \tag{305}$$

and suppose $\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable at $\mathbf{x}$. Then

$$\hat{G}_t(\mathbf{x}) = \nabla \hat{f}_t^{\lambda}(\mathbf{x}) \tag{306}$$

Proof. The right-hand side of Eq. 299 vanishes when $\mathbf{0} \in \partial \varphi_{\mathbf{u},\lambda}(\hat{\mathbf{z}})$, so Eq. 304 holds. At $(\mathbf{u}, \lambda) = (\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t)$, the right-hand side of Eq. 304 is $\hat{\Delta}_t$ by Eq. 296; hence the chain rule through $\hat{\mathbf{x}}_0(\cdot, t)$ yields

$$\nabla \hat{f}_t^{\lambda} = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \nabla_{\mathbf{u}} \hat{V} = \mathbf{J}_{\hat{\mathbf{x}}_0}^\top \hat{\Delta}_t = \hat{G}_t \tag{307}$$

□

H.4 PATH-SPACE KL

The branch envelope removes branch error; the remaining error is numerical, namely the gap between the implemented guidance $\hat{G}_t$ and the envelope gradient $\nabla \hat{f}_t^{\lambda}$. We bound the path-space KL divergence between the two guided SDEs. A pair $(\mathbf{x}, t)$ is a *regular point* if $\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable at $\mathbf{x}$ and $\hat{\mathbf{z}}(\cdot, \lambda_t)$ is differentiable at $\mathbf{u} \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}, t)$. If $\mathbf{J}_{\hat{\mathbf{x}}_0}(\cdot, t)$ is nonsingular almost everywhere, then $(\mathbf{x}, t)$ is a regular point for almost every $\mathbf{x}$. Indeed, $\hat{\mathbf{x}}_0(\cdot, t)$ is differentiable almost everywhere by Assumption E.2 and Rademacher's theorem. Assumption H.6 makes $\hat{\mathbf{z}}(\cdot, \lambda_t)$ differentiable off a null set, and Lemma H.10 pulls this null set back to a null set of $\mathbf{x}$.

Lemma H.10. *Let $T : \mathbb{R}^d \to \mathbb{R}^d$ be Lipschitz, with Jacobian $\mathbf{J}_T$ nonsingular almost everywhere. Then $T^{-1}(N)$ is Lebesgue-null for every Lebesgue-null set $N \subseteq \mathbb{R}^d$.*

Proof. We may assume N is Borel, since every null set lies in a Borel null set. Then $T^{-1}(N)$ is Borel. Since N is Lebesgue-null, the area formula (Evans & Gariepy, 2015, Section 3.3) gives

$$\int_{T^{-1}(N)} |\det \mathbf{J}_T(\mathbf{x})| \, d\mathbf{x} = \int_N \# T^{-1}\{\mathbf{u}\} \, d\mathbf{u} = 0 \tag{308}$$

The integrand on the left is positive almost everywhere, so $T^{-1}(N)$ is Lebesgue-null. □

Lemma H.11. *Assume E.2, H.1, H.5, and H.6. Suppose the soft objective g is continuously differentiable and $\mathcal{C}$ is Clarke regular everywhere. Then, at every regular point $(\mathbf{x}, t)$,*

$$\|\hat{G}_t(\mathbf{x}) - \nabla \hat{f}_t^{\lambda}(\mathbf{x})\| \leq C_J M_{\hat{\mathbf{z}}} \eta_t \tag{309}$$

Proof. Fix a regular point $(\mathbf{x}, t)$, so that Lemma H.8 applies at $(\mathbf{u}, \lambda_t)$. Assumptions H.5 and H.6 bound the right-hand side of Eq. 299 by $M_{\hat{\mathbf{z}}}\eta_t$, so

$$\|\nabla_{\mathbf{u}}\hat{V}(\mathbf{u}, \lambda_t) - \hat{\Delta}_t\| \le M_{\hat{\mathbf{z}}}\eta_t \tag{310}$$

The chain rule through $\hat{\mathbf{x}}_0(\cdot, t)$ yields

$$\nabla \hat{f}_t^{\lambda}(\mathbf{x}) - \hat{G}_t(\mathbf{x}) = \mathbf{J}_{\hat{\mathbf{x}}_0}^{\top}\left(\nabla_{\mathbf{u}}\hat{V}(\mathbf{u}, \lambda_t) - \hat{\Delta}_t\right) \tag{311}$$

Assumption E.2 gives $\|\mathbf{J}_{\hat{\mathbf{x}}_0}\|_{\mathrm{op}} \le C_J$. Take norms and apply Eq. 310 to obtain $\|\nabla \hat{f}_t^{\lambda}(\mathbf{x}) - \hat{G}_t(\mathbf{x})\| \le C_J M_{\hat{\mathbf{z}}}\eta_t$. □

Theorem H.12. *Fix $0 < \epsilon < T$, $w \in \mathbb{R}$, and $\nu \in \mathcal{P}(\mathbb{R}^d)$. Assume the hypotheses of Theorem H.27 for both $\hat{G}$ and G^{env}. Let $\hat{\mathbb{P}}$ be the path distribution of the guided SDE in Eq. 206 on $[\epsilon, T]$, with guidance $\hat{G}_t$ and terminal distribution ν. Let $\mathbb{P}^{\mathrm{env}}$ be the corresponding path distribution with G_t^{env} in place of $\hat{G}_t$. Then*

$$\mathrm{KL}\left(\mathbb{P}^{\mathrm{env}} \,\|\, \hat{\mathbb{P}}\right) \le \frac{w^2 C_J^2 M_{\hat{\mathbf{z}}}^2}{2} \int_{\epsilon}^{T} g(t)^2 \eta_t^2 \, dt \tag{312}$$

Proof. Let $\hat{b}_t$ and b_t^{env} be the guided drifts in Eq. 184, with $\hat{G}_t$ and G_t^{env} in place of G_t. The two drifts share the terms $\mu(t)\mathbf{x}$ and $g(t)^2\xi_t(\mathbf{x})$, so

$$b_t^{\mathrm{env}} - \hat{b}_t = w g(t)^2 \left(G_t^{\mathrm{env}} - \hat{G}_t\right) \tag{313}$$

Lemma H.11 holds at every regular point, and regular points have full Lebesgue measure at each time by hypothesis. Hence Eq. 313 gives

$$\|b_t^{\mathrm{env}}(\mathbf{x}) - \hat{b}_t(\mathbf{x})\| \le |w| g(t)^2 C_J M_{\hat{\mathbf{z}}}\eta_t, \qquad \text{a.e. } \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{314}$$

By Theorem H.27, the time marginals of $\mathbb{P}^{\mathrm{env}}$ have Lebesgue densities for $t < T$. By Tonelli's theorem, Eq. 314 then holds at $\mathbf{x}_t$ for almost every t, $\mathbb{P}^{\mathrm{env}}$-almost surely.

By Eq. 314, the normalized drift difference $(b_t^{\mathrm{env}} - \hat{b}_t)/g(t)$ is bounded by the deterministic quantity $|w| g(t) C_J M_{\hat{\mathbf{z}}}\eta_t$. Since g and η are bounded on $[\epsilon, T]$, Novikov's condition holds. Girsanov's theorem and uniqueness in distribution identify the resulting path distribution with $\hat{\mathbb{P}}$ and yield

$$\mathrm{KL}\left(\mathbb{P}^{\mathrm{env}} \,\|\, \hat{\mathbb{P}}\right) = \mathbb{E}_{\mathbb{P}^{\mathrm{env}}} \int_{\epsilon}^{T} \frac{\|b_t^{\mathrm{env}}(\mathbf{x}_t) - \hat{b}_t(\mathbf{x}_t)\|^2}{2g(t)^2} \, dt \tag{315}$$

$$\le \frac{w^2 C_J^2 M_{\hat{\mathbf{z}}}^2}{2} \int_{\epsilon}^{T} g(t)^2 \eta_t^2 \, dt \tag{316}$$

□

H.5 Solver Cost

Theorem H.12 states that the path-space error scales quadratically with the solver tolerance η_t. However, the theorem does not say how many solver iterations reach a given tolerance. To bound the iteration count, we consider a smooth soft objective $f = g$ solved by gradient descent; specifically, we assume the proximal objective $\varphi_{\mathbf{u},\lambda}$ is strongly convex, which holds once λ_t is sufficiently small. Under this assumption, gradient descent converges linearly to any target tolerance at a rate uniform in the dimension d and the time t (Proposition H.14). Hence a number of gradient steps logarithmic in the target accuracy controls the path-space error at every sampling step (Corollary H.15).

Assumption H.13. The objective is $f = g$ for a differentiable $g : \mathbb{R}^d \to \mathbb{R}$ with $\|\nabla g\| \le B_g$ on $\mathbb{R}^d$. The gradient ∇g is L_g-Lipschitz, and $\lambda_{\max} L_g < 1$.

Under Assumption H.13, the proximal objective $\varphi_{\mathbf{u},\lambda}$ is strongly convex with modulus $1/\lambda - L_g$. Fix a step count $K \ge 1$. The K-step solver runs gradient descent on $\varphi_{\mathbf{u},\lambda}$ with step size λ,

$$\mathbf{z}_0 \coloneqq \mathbf{u}, \qquad \mathbf{z}_{j+1} \coloneqq \mathbf{z}_j - \lambda \nabla \varphi_{\mathbf{u},\lambda}(\mathbf{z}_j) = \mathbf{u} - \lambda \nabla g(\mathbf{z}_j), \qquad \hat{\mathbf{z}}(\mathbf{u}, \lambda) \coloneqq \mathbf{z}_K \tag{317}$$

The map $\mathbf{z} \mapsto \mathbf{u} - \lambda \nabla g(\mathbf{z})$ is a contraction with constant $q_\lambda \coloneqq \lambda L_g \le \lambda_{\max} L_g < 1$.

Proposition H.14. *Under Assumption H.13, fix $\lambda \in (0, \lambda_{\max}]$. The K-step solver in Eq. 317 satisfies the following.*

1. *The residual of the solver output satisfies*
$$\|\nabla \varphi_{\mathbf{u},\lambda}(\hat{\mathbf{z}}(\mathbf{u},\lambda))\| \le q_\lambda^K B_g, \qquad \mathbf{u} \in \mathbb{R}^d \tag{318}$$
so Assumption H.5 holds with $\eta_t = q_{\lambda_t}^K B_g$.
2. *$\hat{\mathbf{z}}(\cdot,\lambda)$ is Lipschitz with constant $1/(1-q_\lambda)$, so Assumption H.6 holds with $M_{\hat{\mathbf{z}}} = 1/(1-\lambda_{\max} L_g)$.*

Proof. **Part 1.** By Eq. 317,
$$\nabla \varphi_{\mathbf{u},\lambda}(\mathbf{z}_j) = \frac{1}{\lambda}(\mathbf{z}_j - \mathbf{z}_{j+1}), \qquad \mathbf{z}_{j+1} - \mathbf{z}_{j+2} = \lambda(\nabla g(\mathbf{z}_{j+1}) - \nabla g(\mathbf{z}_j)) \tag{319}$$
Since ∇g is L_g-Lipschitz, Eq. 319 gives
$$\|\nabla \varphi_{\mathbf{u},\lambda}(\mathbf{z}_{j+1})\| = \|\nabla g(\mathbf{z}_{j+1}) - \nabla g(\mathbf{z}_j)\| \le L_g \|\mathbf{z}_{j+1} - \mathbf{z}_j\| = q_\lambda \|\nabla \varphi_{\mathbf{u},\lambda}(\mathbf{z}_j)\| \tag{320}$$
Since $\nabla \varphi_{\mathbf{u},\lambda}(\mathbf{z}_0) = \nabla g(\mathbf{u})$, induction on Eq. 320 gives
$$\|\nabla \varphi_{\mathbf{u},\lambda}(\mathbf{z}_K)\| \le q_\lambda^K \|\nabla g(\mathbf{u})\| \le q_\lambda^K B_g \tag{321}$$

Part 2. Fix $\mathbf{u}, \mathbf{u}' \in \mathbb{R}^d$ with iterates $\mathbf{z}_j$ and $\mathbf{z}'_j$. By Eq. 317,
$$\|\mathbf{z}_{j+1} - \mathbf{z}'_{j+1}\| \le \|\mathbf{u} - \mathbf{u}'\| + q_\lambda \|\mathbf{z}_j - \mathbf{z}'_j\| \tag{322}$$
Since $\mathbf{z}_0 - \mathbf{z}'_0 = \mathbf{u} - \mathbf{u}'$, induction on Eq. 322 gives
$$\|\mathbf{z}_K - \mathbf{z}'_K\| \le \sum_{j=0}^{K} q_\lambda^j \|\mathbf{u} - \mathbf{u}'\| \le \frac{1}{1-q_\lambda} \|\mathbf{u} - \mathbf{u}'\| \tag{323}$$
Hence Assumption H.6 holds by Rademacher's theorem and $q_\lambda \le \lambda_{\max} L_g$. □

The K-step solver also satisfies Assumptions H.1, H.2, H.3, and H.4. Assumptions H.1, H.3, and H.4 hold because the map $(\mathbf{u},\lambda) \mapsto \hat{\mathbf{z}}(\mathbf{u},\lambda)$ is continuous and $\mathrm{dom}(f) = \mathbb{R}^d$. Assumption H.2 holds with $L = B_g$, since $\|\mathbf{u} - \mathbf{z}_K\| = \lambda \|\nabla g(\mathbf{z}_{K-1})\| \le \lambda B_g$.

Corollary H.15. *Assume H.13 and the hypotheses of Theorem H.12. Let $\hat{\mathbf{z}}$ be the K-step solver in Eq. 317. Then there is a constant C independent of K such that*
$$\mathrm{KL}\left(\mathbb{P}^{\mathrm{env}} \,\|\, \hat{\mathbb{P}}\right) \le C(\lambda_{\max} L_g)^{2K} \tag{324}$$
In particular, $\mathrm{KL}(\mathbb{P}^{\mathrm{env}} \,\|\, \hat{\mathbb{P}}) \le \delta$ after $K = O(\log(1/\delta))$ gradient steps.

Proof. By Proposition H.14, Assumptions H.5 and H.6 hold with $\eta_t \le (\lambda_{\max} L_g)^K B_g$ and $M_{\hat{\mathbf{z}}} = 1/(1-\lambda_{\max} L_g)$. Hence Theorem H.12 gives Eq. 324 with
$$C := \frac{w^2 C_J^2 B_g^2}{2(1-\lambda_{\max} L_g)^2} \int_\epsilon^T g(t)^2 \, dt \tag{325}$$
The bound $C(\lambda_{\max} L_g)^{2K}$ is at most δ exactly when $K \ge \log(C/\delta)/(2\log(1/(\lambda_{\max} L_g)))$. □

By Eq. 324, each gradient step shrinks the error bound by $(\lambda_{\max} L_g)^2$, so halving δ costs a fixed number of extra steps. Part 1 of Proposition H.14 also gives a sufficient step count to reach a tolerance $\eta_t > 0$. If $L_g B_g = 0$, a single gradient step has zero residual. Otherwise, K_t gradient steps suffice to reach η_t, where
$$K_t := \max\left\{1, \left\lceil \frac{\log(B_g/\eta_t)}{\log(1/(\lambda_t L_g))} \right\rceil \right\} \tag{326}$$
For a fixed tolerance, the step count K_t is nonincreasing as λ_t decreases, and $K_t = 1$ once $\lambda_t \le \eta_t/(L_g B_g)$.

H.6 $\hat{f}_t^\lambda$ PROPERTIES

We now verify that $\hat{f}_t^\lambda$ inherits the properties of f_t^λ from Appendix C. By Lemma H.7, $\hat{f}_t^\lambda$ is finite. By Corollary H.9, $\hat{f}_t^\lambda$ has gradient $\hat{G}_t$ when the solver output is exactly stationary. By Assumption H.4 and Lemma C.8, $\hat{f}_t^\lambda$ is continuous almost everywhere at each positive time; almost-everywhere continuity suffices to prove TV- and W_p-continuity of $(\hat{p}_t^\lambda)_t$ in Section H.7. We prove pointwise convergence and the critical-point property below.

Proposition H.16. *Assume H.1, H.2, H.3, and the hypotheses of Proposition C.11. Then $\hat{f}_t^\lambda(\mathbf{x}) \to f(\mathbf{x})$ as $t \to 0$ for every $\mathbf{x} \in \mathbb{R}^d$.*

Proof. Lemma H.7 and Proposition C.11 give

$$\liminf_{t\to 0} \hat{f}_t^\lambda(\mathbf{x}) \geq \lim_{t\to 0} f_t^\lambda(\mathbf{x}) = f(\mathbf{x}) \tag{327}$$

which handles $f(\mathbf{x}) = \infty$. Now fix $\mathbf{x} \in \mathrm{dom}(f)$, write $\mathbf{u}_t \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}, t)$, and decompose

$$\hat{f}_t^\lambda(\mathbf{x}) = \underbrace{f(\hat{\mathbf{z}}_t)}_{\text{Term 1}} + \underbrace{\frac{1}{2\lambda_t}\|\mathbf{u}_t - \hat{\mathbf{z}}_t\|^2}_{\text{Term 2}} \tag{328}$$

We show that Term 1 converges to $f(\mathbf{x})$ and Term 2 converges to 0. Both terms rest on the displacement bound

$$\|\mathbf{u}_t - \hat{\mathbf{z}}_t\| \leq \|\mathbf{u}_t - \mathbf{x}\| + \lambda_t L(1 + \|\mathbf{u}_t\|) \tag{329}$$

which follows from Assumption H.2 and $\mathrm{dist}(\mathbf{u}_t, \mathrm{dom}(f)) \leq \|\mathbf{u}_t - \mathbf{x}\|$, since $\mathbf{x}$ is feasible.

Term 1. $\mathbf{u}_t \to \mathbf{x}$ by Lemma C.8, and $\lambda_t \to 0$ by Assumption C.1. Thus, the right-hand side of Eq. 329 vanishes and $\hat{\mathbf{z}}_t \to \mathbf{x}$. The output $\hat{\mathbf{z}}_t$ lies in $\mathrm{dom}(f)$ by Assumption H.1, so $f(\hat{\mathbf{z}}_t) \to f(\mathbf{x})$ by Assumption H.3.

Term 2. Square Eq. 329 with $(a+b)^2 \leq 2a^2 + 2b^2$ to obtain

$$\frac{\|\mathbf{u}_t - \hat{\mathbf{z}}_t\|^2}{2\lambda_t} \leq \frac{\|\mathbf{u}_t - \mathbf{x}\|^2}{\lambda_t} + \lambda_t L^2(1 + \|\mathbf{u}_t\|)^2 \tag{330}$$

The first summand vanishes by Lemma C.10, and the second since $\lambda_t \to 0$ and $\mathbf{u}_t$ is bounded. Hence $\hat{f}_t^\lambda(\mathbf{x}) \to f(\mathbf{x})$. □

Suppose the solver returns an exact stationary point of the proximal objective. Then $\hat{\mathbf{x}}_0$ maps every critical point of $\hat{f}_t^\lambda$ to a critical point of f, regardless of which stationary point the solver selects.

Proposition H.17. *Assume H.1. Suppose the soft objective g is continuously differentiable and $\mathcal{C}$ is Clarke regular everywhere. Let $(\mathbf{x}^*, t)$ be a regular point with $\mathbf{u}^* \coloneqq \hat{\mathbf{x}}_0(\mathbf{x}^*, t)$, and assume the following.*

1. *The Jacobian $\mathbf{J}_{\hat{\mathbf{x}}_0}(\mathbf{x}^*, t)$ is nonsingular.*
2. *The solver output $\hat{\mathbf{z}}_t$ is exactly stationary, $\mathbf{0} \in \partial\varphi_{\mathbf{u}^*,\lambda_t}(\hat{\mathbf{z}}_t)$.*

If $\nabla \hat{f}_t^\lambda(\mathbf{x}^*) = \mathbf{0}$, then $\mathbf{0} \in \partial f(\mathbf{u}^*)$.

Proof. Corollary H.9 gives

$$\nabla \hat{f}_t^\lambda(\mathbf{x}^*) = \frac{1}{\lambda_t}\mathbf{J}_{\hat{\mathbf{x}}_0}^\top (\mathbf{u}^* - \hat{\mathbf{z}}_t) \tag{331}$$

Since $\nabla \hat{f}_t^\lambda(\mathbf{x}^*) = \mathbf{0}$ and $\mathbf{J}_{\hat{\mathbf{x}}_0}$ is nonsingular, Eq. 331 gives $\hat{\mathbf{z}}_t = \mathbf{u}^*$. The sum rule then gives

$$\partial\varphi_{\mathbf{u}^*,\lambda_t}(\hat{\mathbf{z}}_t) = \partial f(\hat{\mathbf{z}}_t) + \frac{1}{\lambda_t}(\hat{\mathbf{z}}_t - \mathbf{u}^*) = \partial f(\mathbf{u}^*) \tag{332}$$

Since $\hat{\mathbf{z}}_t$ is exactly stationary, $\mathbf{0} \in \partial\varphi_{\mathbf{u}^*,\lambda_t}(\hat{\mathbf{z}}_t) = \partial f(\mathbf{u}^*)$. □

H.7 $\hat{p}_t^\lambda$ Continuity and Convergence

We now re-establish the guarantees of Appendix D for $\hat{p}_t^\lambda$. Interior continuity holds in both regimes, and convergence at $t = 0$ splits by regime as before.

Lemma H.18. *Assume C.1, C.2, C.4, C.6, D.1, H.1, and H.4. Then for every $t^* \in (0, T]$,*

$$\lim_{t \to t^*} \|\hat{p}_t^\lambda - \hat{p}_{t^*}^\lambda\|_{\mathrm{TV}} = 0 \tag{333}$$

Proof. Replace f_t^λ with $\hat{f}_t^\lambda$ in the proof of Lemma D.8. The proof uses Proposition C.9 only for the bounds $C_f \le f_t^\lambda < \infty$ and the pointwise convergence $f_{t_n}^\lambda \to f_{t^*}^\lambda$. Lemma H.7 gives $C_f \le \hat{f}_t^\lambda < \infty$. Assumption H.4 makes $\hat{V}$ continuous at $(\hat{\mathbf{x}}_0(\mathbf{x}, t^*), \lambda_{t^*})$ for almost every $\mathbf{x}$, and Lemma C.8 makes $(\mathbf{x}, t) \mapsto (\hat{\mathbf{x}}_0(\mathbf{x}, t), \lambda_t)$ continuous. Hence $\hat{f}_{t_n}^\lambda \to \hat{f}_{t^*}^\lambda$ at almost every $\mathbf{x}$. Dominated convergence and Scheffé's lemma require only almost-everywhere convergence, so every other step is unchanged. □

H.7.1 Volume Regime

Proposition H.19. *Assume D.1, D.3, and the hypotheses of Proposition H.16. Then*

$$\lim_{t \to 0} \|\hat{p}_t^\lambda - p_{\mathrm{tar}}\|_{\mathrm{TV}} = 0 \tag{334}$$

Proof. Replace f_t^λ with $\hat{f}_t^\lambda$ in the proof of Proposition D.9, and write $\hat{Z}_t^{\mathrm{prox}}$ for the normalization constant of $\hat{p}_t^\lambda$. The proof uses f_t^λ only through the lower bound $f_t^\lambda \ge C_f$, the positivity $Z_t^{\mathrm{prox}} > 0$, and the pointwise convergence $f_t^\lambda \to f$. Lemma H.7 gives $\hat{f}_t^\lambda \ge C_f$. The proof of Lemma H.18 gives $\hat{Z}_t^{\mathrm{prox}} > 0$. Proposition H.16 gives $\hat{f}_t^\lambda(\mathbf{x}) \to f(\mathbf{x})$ for every $\mathbf{x} \in \mathbb{R}^d$, including points where $f(\mathbf{x}) = \infty$. Hence every other step is unchanged. □

Finally, the volume regime's main result.

Theorem H.20. *Assume the hypotheses of Lemma H.18 and Proposition H.19. Then the map $t \mapsto \hat{p}_t^\lambda$ is continuous on $[0, T]$ in total variation, with the convention $\hat{p}_0^\lambda := p_{\mathrm{tar}}$.*

H.7.2 Surface Regime

The proof of Proposition D.17 uses f_t^λ in two places, the lower bound Eq. 151 and the fibrewise limit of Lemma D.16. The lower bound transfers through Lemma H.7. The fibrewise limit requires a new proof, since the branch envelope evaluates the proximal objective at $\hat{\mathbf{z}}_t$ and not at a minimizer over $\mathcal{S}$.

Lemma H.21. *Assume H.1, H.2, H.3, and the hypotheses of Lemma D.16. Fix $\mathbf{y} \in \mathcal{S}$ and $\mathbf{w} \in N_{\mathbf{y}}\mathcal{S}$, and set*

$$\mathbf{x}_t := \mathbf{y} + \sqrt{\lambda_t}\mathbf{w}, \qquad \mathbf{u}_t := \hat{\mathbf{x}}_0(\mathbf{x}_t, t) \tag{335}$$

As $t \to 0$,

$$\exp(-\hat{V}(\mathbf{u}_t, \lambda_t)) \to \exp(-g(\mathbf{y}) - \|\mathbf{w}\|^2/2) \tag{336}$$

Proof. Lemma H.7 and Eq. 148 give

$$\liminf_{t \to 0} \hat{V}(\mathbf{u}_t, \lambda_t) \ge \liminf_{t \to 0} \Psi_f(\mathbf{u}_t, \lambda_t) \ge g(\mathbf{y}) + \frac{\|\mathbf{w}\|^2}{2} \tag{337}$$

which handles $g(\mathbf{y}) = \infty$. Now suppose $g(\mathbf{y}) < \infty$, so that $\mathbf{y} \in \mathrm{dom}(f)$, and decompose

$$\hat{V}(\mathbf{u}_t, \lambda_t) = \underbrace{g(\hat{\mathbf{z}}_t)}_{\text{Term 1}} + \underbrace{\frac{1}{2\lambda_t}\|\mathbf{u}_t - \hat{\mathbf{z}}_t\|^2}_{\text{Term 2}} \tag{338}$$

We show that Term 1 converges to $g(\mathbf{y})$ and Term 2 has limit superior at most $\|\mathbf{w}\|^2/2$. Both terms rest on the displacement bound

$$\|\mathbf{u}_t - \hat{\mathbf{z}}_t\| \le \sqrt{\lambda_t}\left(\|\mathbf{w}\| + c_1\varepsilon_t + \sqrt{\lambda_t}L(1 + \|\mathbf{u}_t\|)\right) \tag{339}$$

Step 1 of the proof of Lemma D.16 gives $\|\mathbf{u}_t - \mathbf{x}_t\| \le c_1\varepsilon_t\sqrt{\lambda_t}$, so $\operatorname{dist}(\mathbf{u}_t, \operatorname{dom}(f)) \le \|\mathbf{u}_t - \mathbf{y}\| \le \sqrt{\lambda_t}(\|\mathbf{w}\| + c_1\varepsilon_t)$. Assumption H.2 then gives Eq. 339.

Term 1. The right-hand side of Eq. 339 vanishes and $\mathbf{u}_t \to \mathbf{y}$, so $\hat{\mathbf{z}}_t \to \mathbf{y}$. The output $\hat{\mathbf{z}}_t$ lies in $\operatorname{dom}(f)$ by Assumption H.1, so $g(\hat{\mathbf{z}}_t) = f(\hat{\mathbf{z}}_t) \to f(\mathbf{y}) = g(\mathbf{y})$ by Assumption H.3.

Term 2. Square Eq. 339 and divide by $2\lambda_t$ to obtain

$$\frac{\|\mathbf{u}_t - \hat{\mathbf{z}}_t\|^2}{2\lambda_t} \le \frac{1}{2}\left(\|\mathbf{w}\| + c_1\varepsilon_t + \sqrt{\lambda_t}L(1 + \|\mathbf{u}_t\|)\right)^2 \to \frac{\|\mathbf{w}\|^2}{2} \tag{340}$$

since $\varepsilon_t \to 0$, $\lambda_t \to 0$, and $\mathbf{u}_t$ is bounded. Hence $\limsup_{t\to 0} \hat{V}(\mathbf{u}_t, \lambda_t) \le g(\mathbf{y}) + \|\mathbf{w}\|^2/2$, which matches Eq. 337. □

Proposition H.22. *Assume H.1, H.2, H.3, and the hypotheses of Proposition D.17. Let $\varphi : \mathbb{R}^d \to \mathbb{R}$ be continuous with $|\varphi(\mathbf{x})| \le C_\varphi(1 + \|\mathbf{x}\|^p)$ for some $C_\varphi < \infty$ and $p \in [1, \infty)$. As $t \to 0$,*

$$(2\pi\lambda_t)^{-m/2} \int_{\mathbb{R}^d} \varphi(\mathbf{x}) q_t(\mathbf{x}) \exp(-\hat{f}_t^\lambda(\mathbf{x}))\, d\mathbf{x} \to \int_{\mathcal{S}} \varphi(\mathbf{y}) \exp(-g(\mathbf{y})) q_0(\mathbf{y})\, \mathcal{H}^n(d\mathbf{y}) \tag{341}$$

Proof. Replace f_t^λ with $\hat{f}_t^\lambda$ in the proof of Proposition D.17. The proof uses f_t^λ only through the lower bound Eq. 151, the fibrewise limit of Lemma D.16, and the Borel measurability of f_t^λ. Lemma H.7 gives $\hat{f}_t^\lambda \ge f_t^\lambda$, so Eq. 151 holds for $\hat{f}_t^\lambda$ and Steps 1 and 3 are unchanged. Lemma H.21 gives the limit of $\exp(-\hat{f}_t^\lambda(\mathbf{x}_t))$ in Step 2. Assumption H.1 and lower semicontinuity of f make $\hat{f}_t^\lambda$ Borel, as the tube formula in Step 2 requires. □

Theorem H.23. *Assume D.7 and the hypotheses of Proposition H.22. As $t \to 0$,*

1. $(2\pi\lambda_t)^{-m/2}\hat{Z}_t^{\mathrm{prox}} \to Z_{\mathrm{surf}} \in (0, \infty)$.
2. $\int \varphi\, d\hat{p}_t^\lambda \to \int \varphi\, d\pi_{\mathrm{tar}}$ *for every continuous $\varphi : \mathbb{R}^d \to \mathbb{R}$ satisfying $|\varphi(\mathbf{x})| \le C_\varphi(1 + \|\mathbf{x}\|^p)$ for some $C_\varphi < \infty$ and $p \in [1, \infty)$.*
3. $W_p(\hat{p}_t^\lambda, \pi_{\mathrm{tar}}) \to 0$ *for every $p \in [1, \infty)$.*

Proof. The proof of Theorem D.18 applies with $\hat{f}_t^\lambda$ in place of f_t^λ, once Proposition H.22 replaces Proposition D.17. The ratio in Eq. 173 requires $0 < \hat{Z}_t^{\mathrm{prox}} < \infty$. Assumption D.1 gives $q_t > 0$, and Lemma H.7 gives $C_f \le \hat{f}_t^\lambda < \infty$, so $0 < \hat{Z}_t^{\mathrm{prox}} \le \exp(-C_f)$. Every other step is unchanged. □

H.7.3 W_p-Continuity on $[0, T]$

Theorem H.24. *Assume D.2 and the hypotheses of Lemma H.18. In the volume regime, assume the hypotheses of Proposition H.19 and set $\hat{p}_0^\lambda := p_{\mathrm{tar}}$. In the surface regime, assume the hypotheses of Theorem H.23 and set $\hat{p}_0^\lambda := \pi_{\mathrm{tar}}$. Then $t \mapsto \hat{p}_t^\lambda$ is continuous on $[0, T]$ in W_p for every $p \in [1, \infty)$.*

Proof. Replace f_t^λ with $\hat{f}_t^\lambda$ in the proof of Theorem D.19. The proof uses f_t^λ only through the bounds $C_f \le f_t^\lambda < \infty$, the pointwise convergence of f_t^λ as $t \to t^*$, and Part 3 of Theorem D.18. Lemma H.7 gives $C_f \le \hat{f}_t^\lambda < \infty$, so the dominating function is unchanged and each Z_t lies in $(0, \infty)$. For $t^* > 0$, Assumption H.4 and Lemma C.8 give $\hat{f}_t^\lambda(\mathbf{x}) \to \hat{f}_{t^*}^\lambda(\mathbf{x})$ at almost every $\mathbf{x}$, which suffices for dominated convergence. For $t^* = 0$ in the volume regime, Proposition H.16 gives $\hat{f}_t^\lambda(\mathbf{x}) \to f(\mathbf{x})$ at every $\mathbf{x}$. In the surface regime, Part 3 of Theorem H.23 gives continuity at $t = 0$. □

H.8 Well-Posedness

Theorem E.6 uses the guidance field only through two properties: joint Borel measurability and uniform linear growth on $[\epsilon, T]$. We verify both properties for $\hat{G}_t$ and, under the hypotheses of Lemma H.11, for $\nabla\hat{f}_t^\lambda$. Neither property requires continuity, so branch switching does not affect well-posedness.

H.8.1 Regularity of the Realized Drift

Lemma H.25. *Assume C.6, E.1, E.2, H.1, and H.2. Fix $\mathbf{z}^\dagger \in \mathrm{dom}(f)$, and define*

$$K_{\hat{G}}(\mathbf{x}) \coloneqq \frac{C_J}{\lambda_{\min}} \left[\|\mathbf{z}^\dagger\| + \lambda_{\max} L + (1 + \lambda_{\max} L)(H_0 + C_J \|\mathbf{x}\|) \right] \tag{342}$$

There is a jointly Borel map $\hat{G} : \mathbb{R}^d \times [\epsilon, T] \to \mathbb{R}^d$ with slices $\hat{G}_t \coloneqq \hat{G}(\cdot, t)$ such that

$$\|\hat{G}_t(\mathbf{x})\| \le K_{\hat{G}}(\mathbf{x}), \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{343}$$

For every $t \in [\epsilon, T]$, the slice $\hat{G}_t$ equals $\mathbf{J}_{\hat{\mathbf{x}}_0}^\top \hat{\Delta}_t$ at almost every $\mathbf{x} \in \mathbb{R}^d$.

Proof. **Construction.** Apply the construction in Lemma E.4 to the coordinates of $\hat{\mathbf{x}}_0$, with the operator-norm bound C_J in place of K_G. The construction gives a jointly Borel $\mathcal{J}$ with $\|\mathcal{J}\|_{\mathrm{op}} \le C_J$ and $\mathcal{J}(\cdot, t) = \mathbf{J}_{\hat{\mathbf{x}}_0}(\cdot, t)$ almost everywhere, for each t. Assumption H.1 makes the displacement $\hat{\Delta}_t$ in Eq. 296 jointly Borel. Set $\hat{G} \coloneqq \mathcal{J}^\top \hat{\Delta}_t$.

Linear growth. Replace the proximal point $\mathbf{z}$ with $\hat{\mathbf{z}}_t$ in the proof of Lemma E.3. The proof uses $\mathbf{z}$ only through the displacement bound Eq. 194. Since $\mathbf{z}^\dagger \in \mathrm{dom}(f)$, Assumption H.2 gives

$$\|\mathbf{u} - \hat{\mathbf{z}}(\mathbf{u}, \lambda_t)\| \le \|\mathbf{u} - \mathbf{z}^\dagger\| + \lambda_{\max} L(1 + \|\mathbf{u}\|) \le \|\mathbf{z}^\dagger\| + \lambda_{\max} L + (1 + \lambda_{\max} L)\|\mathbf{u}\| \tag{344}$$

in place of Eq. 194. Every other step is unchanged, with $K_{\hat{G}}$ in place of K_G. □

Lemma H.26. *Assume the hypotheses of Lemmas H.11 and H.25, and suppose regular points have full Lebesgue measure for each $t \in [\epsilon, T]$. There is a jointly Borel map $G^{\mathrm{env}} : \mathbb{R}^d \times [\epsilon, T] \to \mathbb{R}^d$ with slices $G_t^{\mathrm{env}} \coloneqq G^{\mathrm{env}}(\cdot, t)$ such that*

$$\|G_t^{\mathrm{env}}(\mathbf{x})\| \le K_{\hat{G}}(\mathbf{x}) + C_J M_{\hat{\mathbf{z}}} \eta_t, \qquad \mathbf{x} \in \mathbb{R}^d, \quad t \in [\epsilon, T] \tag{345}$$

For every $t \in [\epsilon, T]$, the slice G_t^{env} equals $\nabla \hat{f}_t^\lambda$ at almost every $\mathbf{x} \in \mathbb{R}^d$.

Proof. At every regular point, Lemma H.11 and Eq. 343 bound $\|\nabla \hat{f}_t^\lambda(\mathbf{x})\|$ by the right-hand side of Eq. 345. Apply the construction in Lemma E.4 to $\hat{f}_t^\lambda$, with the right-hand side of Eq. 345 in place of K_G. The construction requires only a jointly Borel function and a jointly Borel bound. Assumption H.1 makes $\hat{f}_t^\lambda$ jointly Borel, and η Borel makes the bound jointly Borel. The slice G_t^{env} equals $\nabla \hat{f}_t^\lambda$ at every regular point, hence almost everywhere. □

H.8.2 Strong Well-Posedness

Theorem H.27. *Assume the hypotheses of Lemma H.25, and fix $w \in \mathbb{R}$. Then the conclusions of Theorem E.6 and Corollary E.7 hold with guidance $\hat{G}$. Under the hypotheses of Lemma H.26, the same conclusions hold with guidance G^{env}.*

Proof. The bounds Eq. 343 and Eq. 345 are linear growth bounds uniform in t, since η is bounded. Hence $\hat{G}$ and G^{env} satisfy the hypotheses of Lemma E.5. □

H.9 Proximal Correction Bias

Proposition G.3 applies to every terminal distribution, including the output distribution of the realized sampler, with the same projection $P_{\mathcal{K}}$. We show that the corrected realized path $(P_{\mathcal{K}})_{\#}\hat{p}_t^\lambda$ converges to the target in each regime, with $\mathcal{K}$ as in Appendix G.2.

Proposition H.28. *Assume the hypotheses of Proposition H.19. In the volume regime, $\|(P_{\mathcal{C}})_{\#}\hat{p}_t^\lambda - p_{\mathrm{tar}}\|_{\mathrm{TV}} \to 0$ as $t \to 0$.*

Proof. Replace p_t^λ with $\hat{p}_t^\lambda$ in the proof of Proposition G.4. The proof uses p_t^λ only through Proposition D.9, and Proposition H.19 gives the same convergence for $\hat{p}_t^\lambda$. □

Theorem H.29. *Assume the hypotheses of Theorem H.23. In the surface regime, $\|(P_{\mathcal{K}})_{\#}\hat{p}_t^\lambda - \pi_{\mathrm{tar}}\|_{\mathrm{TV}} \to 0$ as $t \to 0$.*

Proof. Replace f_t^λ with $\hat{f}_t^\lambda$ in the proof of Theorem G.5. The proof uses f_t^λ only through the pointwise limit Eq. 159, the domination Eq. 166, the tube complement Eq. 156, and Part 1 of Theorem D.18. The proof of Proposition H.22 establishes Eq. 159, Eq. 166, and Eq. 156 for $\hat{f}_t^\lambda$. Part 1 of Theorem H.23 gives the limit of the normalization constant. □

The final correction is one more proximal solve at $\lambda_{\Delta t}$. For a pure constraint $f = \chi_{\mathcal{K}}$, an exact solve is a Borel nearest-point selection onto $\mathcal{K}$, by Assumption H.1. The proofs in Appendix G use $P_{\mathcal{K}}$ only as a Borel nearest-point selection. Hence Proposition H.28 and Theorem H.29 hold with the exact solve in place of $P_{\mathcal{K}}$.

I EXPERIMENTS

I.1 PDE SOLUTION GENERATION

I.1.1 EXPERIMENTAL SETUP

Our experimental setup follows Utkarsh et al. (2025), which details FFM training, data generation, and the residual formulas.

Table 4: The five PDE solution-generation tasks. Each ground-truth ensemble fixes the listed condition and varies the free parameter. IC/BC are the initial/boundary condition, CL is the conservation law, and grid is $N_x(\times N_y) \times N_t$.

Task	Dim	Grid	Fixed	Free	Constraints
Heat	1D	100×100	IC	decay rate	IC, mass CL (linear)
Navier–Stokes	2D	$64 \times 64 \times 50$	IC (vort.)	forcing	IC, vorticity CL (linear)
Reaction–Diffusion	1D	128×100	IC	BC	IC, mass CL (nonlinear)
Burgers (IC)	1D	101×101	IC	BC	IC, mass CL (nonlinear), local residual
Burgers (BC)	1D	101×101	BC	IC	BC, mass CL (nonlinear)

Shared prior and projection. All methods guide the same pretrained Functional Flow Matching prior (Kerrigan et al., 2024) with an FNO velocity field. All proximal guidance methods (PCFM, PDM, ProxDG) use Gauss–Newton projection onto the constraint set $h(\mathbf{u}) = 0$. Given a candidate $\mathbf{u}_1$, the projection solves the constrained least-squares problem

$$\min_{\mathbf{u}'} \ \tfrac{1}{2}\|\mathbf{u}' - \mathbf{u}_1\|^2 \quad \text{s.t.} \quad h(\mathbf{u}') = 0 \tag{346}$$

Linearizing h at $\mathbf{u}_1$ gives the closed-form step

$$\mathbf{u}_{\text{proj}} = \mathbf{u}_1 - J^\top (JJ^\top)^{-1} h(\mathbf{u}_1), \qquad J = \nabla h(\mathbf{u}_1)^\top \tag{347}$$

This step is exact for affine h, and we iterate it to tolerance for nonlinear h. The Burgers (IC) local residual enforces a Godunov-flux collocation step, defined in Utkarsh et al. (2025).

Evaluation protocol. Each task holds a condition fixed and varies a free parameter (Table 4), and the resulting set of solutions forms the ground-truth ensemble. We draw K samples per condition to match the ground-truth size, with $K = 512$ for the 1D tasks and $K = 100$ for Navier–Stokes. We average metrics over 20 conditions for Reaction–Diffusion and both Burgers tasks, and over a single condition for Heat and Navier–Stokes. Test data are regenerated with PCFM's released scripts, using the analytic solution for Heat.

Metrics. We adopt ECI's metrics (Cheng et al., 2025). CE is the ℓ_2 norm of the flattened constraint residual, averaged over samples and conditions. MMSE and SMSE compare the first two moments of the generated and ground-truth ensembles,

$$\text{MMSE} = \text{mean}_\Omega \left((\bar{u}_{\text{gen}} - \bar{u}_{\text{gt}})^2 \right), \qquad \text{SMSE} = \text{mean}_\Omega \left((\sigma_{\text{gen}} - \sigma_{\text{gt}})^2 \right) \tag{348}$$

where $\bar{u}$ and σ are the per-grid-point ensemble mean and standard deviation, and mean_Ω averages over the grid. FPD is the Fréchet distance between Poseidon-B embeddings (Herde et al., 2024) of the two ensembles.

I.1.2 HYPERPARAMETER TUNING

We tune each method's free hyperparameters per-problem on a held-out validation split. We generate this split with PCFM's data-generation script[1] and an independent random seed. To reduce per-trial cost, we use $K = 128$ for the 1D problems, $K = 32$ for Navier–Stokes, and 5 of the 20 conditions for Burgers and Reaction–Diffusion. We keep all other settings from Utkarsh et al. (2025) fixed. Every sweep jointly minimizes MMSE and CE, with final projections enabled for PCFM, PDM,

[1] https://github.com/cpfpengfei/PCFM

and ProxDG. Table 5 lists each method's search space and trial count. The threshold $t_{\min}$ disables guidance at earlier times, and the optional $(1-t)$ scaling decays guidance over time. The weights w (ProxDG), η_{sparse} and η_{pde} (FunDPS), and w_{pde} (D-Flow) scale the guidance or its residual loss terms, and the D-Flow learning rate sets the L-BFGS step size over the initial noise.

For each problem, we select a configuration from each method's (MMSE, CE) Pareto front. For PCFM, PDM, and ProxDG, every Pareto-front configuration reaches zero constraint violation, so we take the one with the lowest MMSE. No FunDPS or D-Flow configuration reaches zero violation, so we take the one with the lowest CE.

Table 5: Hyperparameter search spaces and trial counts.

Method	Hyperparameter	Range / Distribution	Trials
PCFM, PDM	$t_{\min}$	$\{0.00, 0.02, \ldots, 0.98\}$	50
FunDPS	η_{sparse}	log-uniform $[10^{-2}, 10^{2}]$	150
	η_{pde}	log-uniform $[10^{-2}, 10^{2}]$	
	$t_{\min}$	uniform $[0, 1]$	
	$(1-t)$ scaling	{on, off}	
D-Flow	learning rate	log-uniform $[10^{-2}, 1]$	100
	w_{pde}	log-uniform $[10^{-3}, 10^{3}]$	
ProxDG	w	log-uniform $[10^{-2}, 10^{2}]$	100
	$t_{\min}$	uniform $[0, 1]$	
	$(1-t)$ scaling	{on, off}	

I.1.3 Results

Sampling cost. Table 6 reports end-to-end sampling time, and Figure 8 reports the time of each guidance operation; Table 7 reports the operations each method uses. Each operation takes a few milliseconds except the Navier–Stokes projection, which runs one sample at a time because its batched Jacobian does not fit in memory. On Navier–Stokes, the projection dominates the cost of all three proximal guidance methods. D-Flow is the slowest method on the 1D tasks.

Table 6: End-to-end sampling time on the five PDE tasks, in seconds per sample. Times are amortized over batches of $K = 128$ samples for the 1D tasks and $K = 32$ for Navier–Stokes.

Method	Heat	Navier–Stokes	Reaction–Diffusion	Burgers (IC)	Burgers (BC)
PCFM	0.765	48.3	0.498	0.347	0.213
PDM	0.047	52.0	0.943	1.82	0.479
FunDPS	0.136	1.02	0.351	0.801	0.693
D-Flow	3.27	32.3	8.28	7.43	7.62
ProxDG (Ours)	0.117	52.8	1.09	2.35	0.843

Table 7: Guidance operations used by each method in Figure 8.

Method	Clean-sample estimate	Denoiser VJP	Clean-space loss gradient	Gauss–Newton projection
PCFM	✓			✓
PDM				✓
FunDPS	✓	✓	✓	
D-Flow		✓	✓	
ProxDG (Ours)	✓	✓		✓

On Burgers (BC), ProxDG places second to PCFM on all three fidelity metrics. The per-timestep diagnostics (Fig. 9) show what drives this performance gap: the two methods produce nearly identical mean fields at every step, so the gap is driven by the ensemble spread (SMSE) rather than the mean.

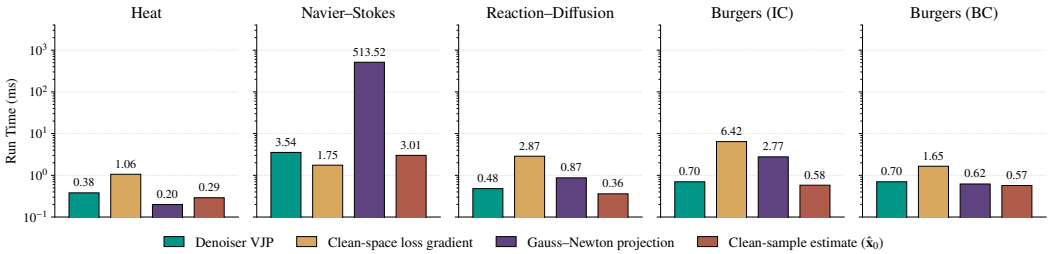

Figure 8: Run time of each guidance operation on the five PDE tasks, in milliseconds per sample (log scale). Times are amortized over batches of $K = 128$ samples for the 1D tasks and $K = 32$ for Navier–Stokes. The Gauss–Newton projection reports a single iteration.

This difference in spread stems from the prior. The unconstrained FFM prior reproduces GT's within-condition spread to within 5% (Table 9), whereas on the other four tasks it is over-spread by factors of 2 to 25. PCFM's hard projection is active only over the final $\approx 16\%$ of the trajectory ($t_{\min} = 0.84$), while ProxDG corrects at every step ($t_{\min} \approx 0$). PCFM therefore leaves the prior's already-accurate spread largely intact (covariance-trace ratio 0.95), while ProxDG's correction across the full trajectory reduces the spread further below GT (0.84).

Table 8: PDE solution generation across all five tasks. We report constraint residuals (CE IC/BC, CE CL), ensemble mean and standard-deviation errors (MMSE, SMSE), and the Fréchet PDE distance (FPD). All values $\times 10^{-2}$ except FPD. Best in bold, second-best underlined. "div." marks runs with NaN samples, left unscored.

Dataset	Method	CE (IC/BC)	CE (CL)	MMSE	SMSE	FPD
Heat	PCFM	**0**	**0**	0.107	0.0742	0.100
	PDM	**0**	**0**	0.470	0.118	0.291
	FunDPS	<u>6.43</u>	<u>2.92</u>	<u>0.0331</u>	**0.0325**	**0.0302**
	D-Flow	19.9	3.18	0.0421	0.0840	0.0922
	ProxDG (Ours)	**0**	**0**	**0.00637**	<u>0.0546</u>	<u>0.0316</u>
Navier–Stokes	PCFM	**0**	**0**	17.8	12.7	2.47
	PDM	**0**	**0**	12.2	6.01	1.96
	FunDPS	<u>579</u>	9.32	<u>2.98</u>	<u>1.57</u>	1.91
	D-Flow	1020	<u>3.02</u>	5.94	4.06	<u>0.823</u>
	ProxDG (Ours)	**0**	**0**	**2.02**	**0.821**	**0.451**
Reaction–Diffusion	PCFM	**0**	**0**	0.290	0.120	50.5
	PDM	**0**	**0**	2.56	0.685	158
	FunDPS	<u>7.26</u>	34.6	0.180	**0.0687**	25.6
	D-Flow	93.1	<u>6.76</u>	<u>0.109</u>	0.681	**14.8**
	ProxDG (Ours)	**0**	**0**	**0.0860**	<u>0.0693</u>	<u>23.5</u>
Burgers (IC)	PCFM	**0**	**0**	3.74	2.67	20.2
	PDM	div.	div.	div.	div.	div.
	FunDPS	<u>35.3</u>	<u>14.1</u>	<u>0.397</u>	<u>0.651</u>	**0.583**
	D-Flow	68.2	16.9	0.693	1.75	3.00
	ProxDG (Ours)	**0**	**0**	**0.180**	**0.268**	<u>0.897</u>
Burgers (BC)	PCFM	**0**	**0**	**0.437**	**0.108**	**1.26**
	PDM	**0**	**0**	0.840	0.442	1.86
	FunDPS	<u>14.5</u>	2.83	0.789	0.307	3.86
	D-Flow	20.2	<u>2.09</u>	0.852	0.469	1.47
	ProxDG (Ours)	**0**	**0**	<u>0.496</u>	<u>0.238</u>	<u>1.31</u>

Table 9: Within-condition covariance trace ratio to ground truth (GT), averaged over conditions. FFM is the unconstrained prior. Values near 1.0 match GT width.

Task	FFM	PCFM	ProxDG
Heat	5.95	0.68	0.54
Navier–Stokes	1.98	3.04	1.08
Reaction–Diffusion	25.00	2.42	0.20
Burgers (BC fixed)	1.05	0.95	0.84
Burgers (IC fixed)	5.27	2.17	0.70

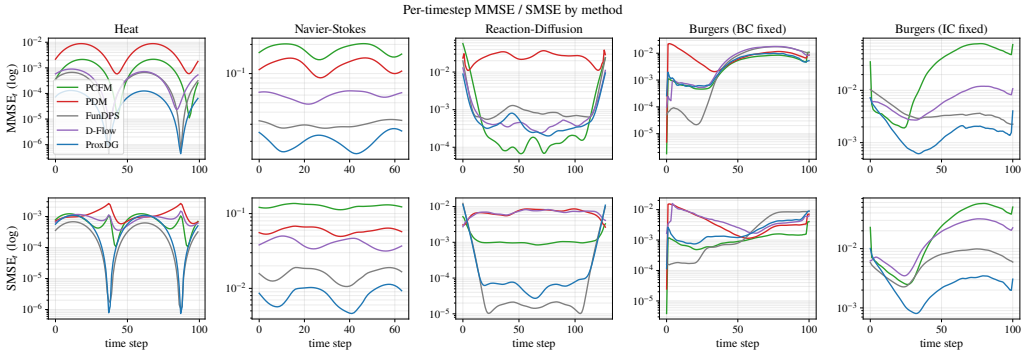

Figure 9: MMSE (top) and SMSE (bottom) versus physical timestep, across the five PDE tasks. Each metric is computed over the spatial domain at each timestep, then averaged over conditions.

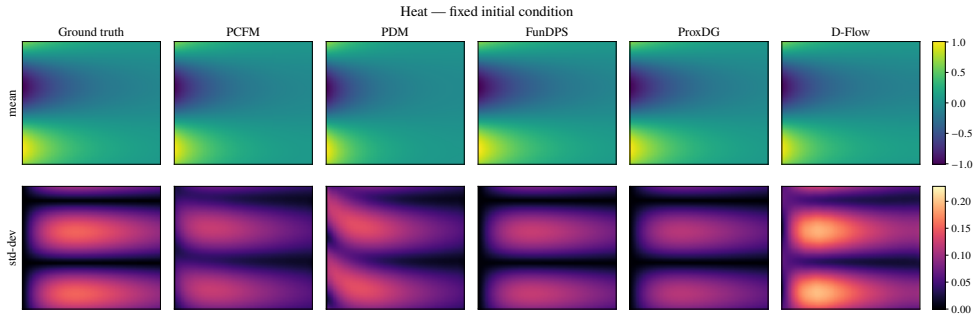

Figure 10: Heat (fixed initial condition). Per-grid-point ensemble mean (top row) and standard deviation (bottom row) of the generated solutions for each method (PCFM, PDM, FunDPS, ProxDG), alongside the ground-truth ensemble. Axes are the spatial coordinate x and physical time t.

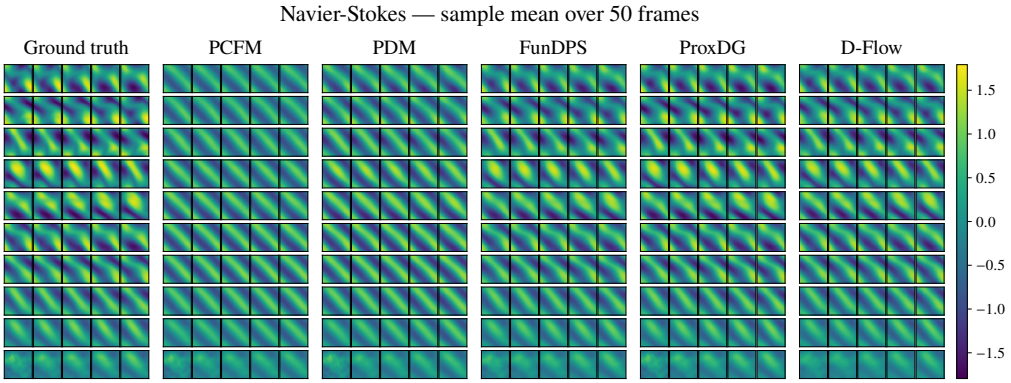

Figure 11: Navier–Stokes (fixed initial vorticity). Per-pixel ensemble mean across 100 generated solutions for each method (PCFM, PDM, FunDPS, ProxDG), alongside the ground-truth ensemble. Each method's 50 time frames are tiled as a 10×5 grid (rows $\times$ columns), with time increasing left-to-right then top-to-bottom.

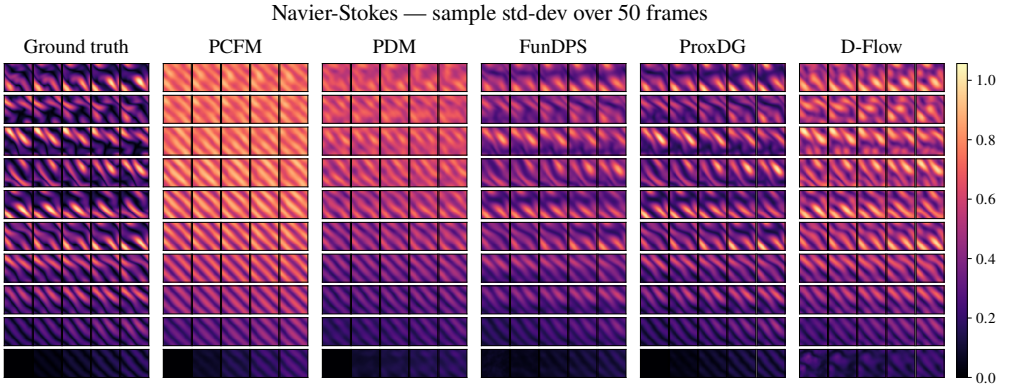

Figure 12: Navier–Stokes (fixed initial vorticity). Per-pixel ensemble standard deviation across 100 generated solutions for each method (PCFM, PDM, FunDPS, ProxDG), alongside the ground-truth ensemble. Each method's 50 time frames are tiled as a 10×5 grid (rows $\times$ columns), with time increasing left-to-right then top-to-bottom.

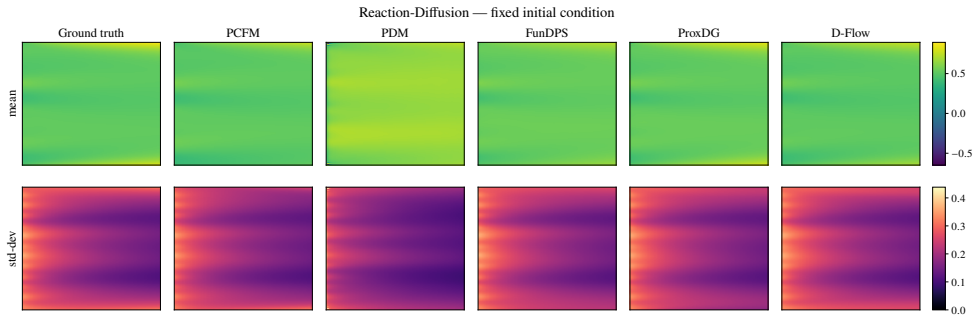

Figure 13: Reaction–Diffusion (fixed initial condition). Per-grid-point ensemble mean (top row) and standard deviation (bottom row) of the generated solutions for each method (PCFM, PDM, FunDPS, ProxDG), alongside the ground-truth ensemble. Axes are the spatial coordinate x and physical time t.

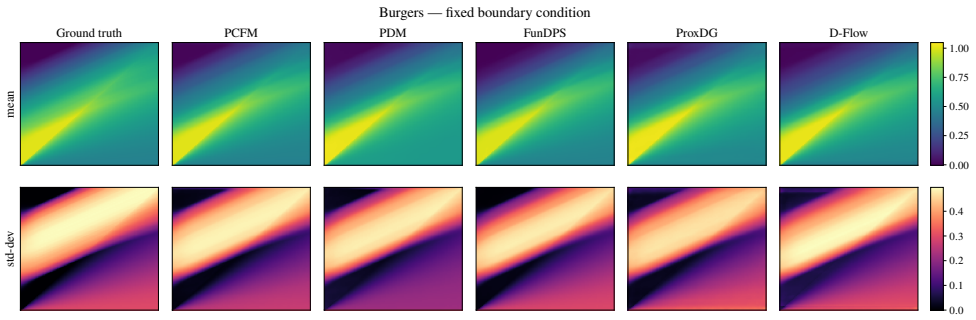

Figure 14: Burgers (fixed boundary condition). Per-grid-point ensemble mean (top row) and standard deviation (bottom row) of the generated solutions for each method (PCFM, PDM, FunDPS, ProxDG), alongside the ground-truth ensemble. Axes are the spatial coordinate x and physical time t.

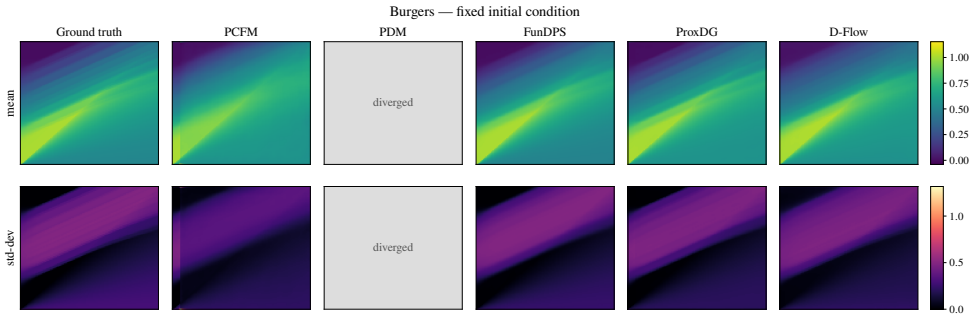

Figure 15: Burgers (fixed initial condition). Per-grid-point ensemble mean (top row) and standard deviation (bottom row) of the generated solutions for each method (PCFM, PDM, FunDPS, ProxDG), alongside the ground-truth ensemble. Axes are the spatial coordinate x and physical time t.

I.1.4 Ablations

Higher-order sampling. Multistep integrators such as DPM-Solver++ (Lu et al., 2025) preserve sample quality under guidance with few steps. They do not readily transfer to proximal guidance methods such as PCFM and PDM, which do not integrate a continuous sampling trajectory. ProxDG, FunDPS, and D-Flow, by contrast, can adopt such integrators directly. We replace the Euler integrator of ProxDG, FunDPS, and D-Flow with DPM-Solver++ and cut the number of sampling steps by $4\times$, from 100 to 25 on Heat and Navier–Stokes and from 200 to 50 on Reaction–Diffusion and both Burgers tasks (Table 10). ProxDG remains exactly feasible on every task, whereas FunDPS and D-Flow leave constraint residuals between 1.72 and 518. ProxDG also leads on MMSE and SMSE on the three nonlinear 1D tasks, but trails on Heat and Navier–Stokes.

Table 10: PDE solution generation with DPM-Solver++ at $4\times$ fewer sampling steps than Euler. Each cell reports DPM-Solver++ / Euler, with Euler values from Table 8. All values $\times 10^{-2}$ except FPD.

Dataset	Method	CE (IC/BC)	CE (CL)	MMSE	SMSE	FPD
Heat	FunDPS	3.14 / 6.43	4.25 / 2.92	0.0028 / 0.0331	0.0537 / 0.0325	0.035 / 0.0302
	D-Flow	25.2 / 19.9	1.72 / 3.18	0.0045 / 0.0421	0.0553 / 0.0840	0.075 / 0.0922
	ProxDG	0 / 0	0 / 0	0.0096 / 0.00637	0.0559 / 0.0546	0.038 / 0.0316
Navier–Stokes	FunDPS	334 / 579	4.42 / 9.32	2.13 / 2.98	1.16 / 1.57	0.511 / 1.91
	D-Flow	518 / 1020	5.01 / 3.02	0.89 / 5.94	0.599 / 4.06	0.248 / 0.823
	ProxDG	0 / 0	0 / 0	2.46 / 2.02	1.07 / 0.821	0.441 / 0.451
Reaction–Diffusion	FunDPS	10.5 / 7.26	9.90 / 34.6	0.075 / 0.180	0.048 / 0.0687	25.4 / 25.6
	D-Flow	104 / 93.1	6.13 / 6.76	0.150 / 0.109	0.768 / 0.681	21.0 / 14.8
	ProxDG	0 / 0	0 / 0	0.043 / 0.0860	0.046 / 0.0693	32.9 / 23.5
Burgers (IC)	FunDPS	28.4 / 35.3	9.44 / 14.1	0.294 / 0.397	0.512 / 0.651	0.575 / 0.583
	D-Flow	62.8 / 68.2	19.8 / 16.9	0.656 / 0.693	1.73 / 1.75	2.89 / 3.00
	ProxDG	0 / 0	0 / 0	0.202 / 0.180	0.391 / 0.268	0.654 / 0.897
Burgers (BC)	FunDPS	32.1 / 14.5	4.50 / 2.83	0.312 / 0.789	0.198 / 0.307	2.38 / 3.86
	D-Flow	76.2 / 20.2	3.14 / 2.09	0.531 / 0.852	3.27 / 0.469	1.86 / 1.47
	ProxDG	0 / 0	0 / 0	0.209 / 0.496	0.119 / 0.238	1.27 / 1.31

I.2 COMPRESSED-SENSING MRI

I.2.1 PROBLEM DETAILS

Compressed-sensing MRI shortens scan time by measuring only a small fraction of the image's Fourier coefficients. Our setup follows InverseBench (Zheng et al., 2025), which describes dataset preparation and diffusion model training. We use the multi-coil fastMRI knee dataset (Zbontar et al., 2019; Knoll et al., 2020) and evaluate acceleration factors $R = 4$ and $R = 8$, which retain $1/4$ and $1/8$ of the k-space lines. We evaluate each factor on both raw and simulated (noiseless) measurements (Table 14).

Forward model. A multi-coil MRI scanner surrounds the patient with J receiver coils. Each coil j records sensitivity-weighted Fourier samples of the same complex-valued image $\mathbf{z} \in \mathbb{C}^n$,

$$\mathbf{y}_j = \mathbf{P}\mathbf{F}(\mathbf{S}_j\mathbf{z}) + \mathbf{n}_j, \qquad j = 1, \dots, J \tag{349}$$

where $\mathbf{F}$ is the 2D Fourier transform, $\mathbf{S}_j$ is the known sensitivity map of coil j, $\mathbf{P}$ is a binary subsampling mask in k-space, and $\mathbf{n}_j$ is Gaussian thermal noise. We write $\mathbf{y}$ for the stacked measurements and $\mathbf{A}$ for the stacked per-coil operators $\mathbf{PFS}_j$.

Mask and acceleration factor. The mask $\mathbf{P}$ retains $1/R$ of the k-space lines along the phase-encode axis. It fully samples a central block of low-frequency lines (8% of all lines at $R = 4$ and 4% at $R = 8$) and draws the remaining lines uniformly at random from the higher frequencies. Following InverseBench, all test slices share one fixed mask.

Measurement regimes. *Raw* measurements use the original fastMRI k-space, which carries real thermal noise. *Simulated* measurements $\mathbf{A}\mathbf{z}^*$ are generated without noise from the ground-truth image $\mathbf{z}^*$, which isolates reconstruction quality from measurement noise. The ground truth is the minimum-variance unbiased estimate (MVUE) that InverseBench computes from the fully sampled k-space.

Evaluation. We use the InverseBench test set of 96 knee slices, each 320×320. We score each reconstruction $\hat{\mathbf{z}}$ by the PSNR and SSIM of its magnitude image and by the data-fit residual $\|\mathbf{y} - \mathbf{A}\hat{\mathbf{z}}\|_2$. Finally, we report the mean (std) of these metrics over the test set.

I.2.2 SOLVER DETAILS

Proximal subproblem. At each sampling step, ProxDG solves the proximal subproblem in Eq. 9 for the data-fit objective,

$$\mathbf{z}^\star = \arg\min_{\mathbf{z}} \left\{ \frac{\lambda_t}{2} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \frac{\rho}{2} \|\mathbf{z} - \hat{\mathbf{x}}_0\|_2^2 \right\} \tag{350}$$

where $\hat{\mathbf{x}}_0$ is the clean estimate. The forward operator $\mathbf{A}$ is linear, so the minimizer solves the normal equations

$$\left(\lambda_t \mathbf{A}^{\mathsf{H}}\mathbf{A} + \rho\mathbf{I}\right) \mathbf{z} = \lambda_t \mathbf{A}^{\mathsf{H}}\mathbf{y} + \rho\hat{\mathbf{x}}_0 \tag{351}$$

A single linear solve therefore returns the exact proximal point, with no Gauss–Newton iterations.

Row decoupling. The mask structure reduces the system in Eq. 351 to 320 independent 320×320 systems, one per readout row. Because $\mathbf{P}$ subsamples only the phase-encode direction, $\mathbf{F}^{\mathsf{H}}\mathbf{P}^\top\mathbf{PF}$ applies the same 1D circular convolution $\mathbf{\Psi}$ to each row. The coil sensitivity maps act pointwise, so the rows stay independent. For row x, the system matrix is $\lambda_t(\mathbf{\Psi} \circ \mathbf{C}_x) + \rho\mathbf{I}$, where $\circ$ denotes the Hadamard product and $\mathbf{C}_x[w, w'] = \sum_j \overline{s}_j(x, w) s_j(x, w')$ is the row's coil correlation matrix. Both $\mathbf{\Psi}$ and $\mathbf{C}_x$ are Hermitian positive semidefinite, so their Hadamard product is too by the Schur product theorem. Adding $\rho\mathbf{I}$ with $\rho > 0$ makes each system Hermitian positive definite, so we factor all 320 systems in parallel with batched Cholesky.

Implementation. Since all test slices share one mask, we compute $\mathbf{\Psi}$ once and reuse it across all 1000 diffusion steps. We compute $\mathbf{\Psi}$ by passing a unit impulse through the forward operator's FFT-and-mask chain, so $\mathbf{\Psi}$ matches the operator's FFT convention by construction. At each diffusion step, we form the coil correlations and factor the row systems, at a cost of a few GFLOPs. The resulting solve is exact up to floating-point error.

I.2.3 HYPERPARAMETER TUNING

Following InverseBench (Zheng et al., 2025), we tune ProxDG with a coarse grid search followed by Bayesian optimization on each problem's validation split. The validation splits contain 6 knee slices for CS-MRI, 5 images for BHI, and 1 velocity map for FWI. Each sweep runs 100 trials over the four hyperparameters in Table 11 and maximizes PSNR. We disable the final proximal correction (Section 4.3) during tuning. Unless stated otherwise, we also *fully retune* every baseline and ablation variant with the same procedure and trial budget.

Table 11: ProxDG search space for the InverseBench problems.

Hyperparameter	Search space
β	log-uniform $[10^{-3}, 10^{7}]$
w	log-uniform $[10^{-2}, 10^{2}]$
α_t scaling	{on, off}
$t_{\min}$	uniform $[0, T]$

Here, β sets the annealing schedule $\lambda_t = \beta\sigma_t^2$ (Section 4.1), and w is the guidance scale. The optional α_t scaling multiplies the guidance weight by the time-varying factor α_t (Yu et al., 2023; Shen et al., 2024). Below the timestep $t_{\min}$, the λ_t schedule stops decaying.

I.2.4 RESULTS

Sampling cost. Table 12 reports end-to-end sampling time on all three InverseBench problems. Figures 16, 19, and 25 report the time of each guidance operation, and Table 13 lists the operations each method uses.

Table 12: End-to-end sampling time per sample on the InverseBench problems. MRI uses raw measurements at $R = 4$, and BHI uses the 100% observation-time ratio. We time ProxDG (PP), and ProxDG (DS) differs by one proximal solve.

Problem	DPS	DiffPIR	RED-diff	PnP-DM	DAPS	ProxDG (Ours)
MRI (s)	43.5	20.7	20.6	81.1	45.8	49.2
BHI (s)	37.7	13.0	13.2	188.7	17.4	273.0
FWI (min)	93.3	97.0	99.5	157.5	732.7	29.9

Table 13: Guidance operations used by each method, as timed in Figures 16, 19, and 25. The number of calls per timestep varies by method.

Method	Clean-sample estimate	Denoiser VJP	Clean-space loss gradient	Proximal solve
DPS	✓	✓	✓	
DiffPIR	✓		✓	
RED-diff	✓		✓	
PnP-DM	✓		✓	
DAPS	✓		✓	
ProxDG (Ours)	✓	✓		✓

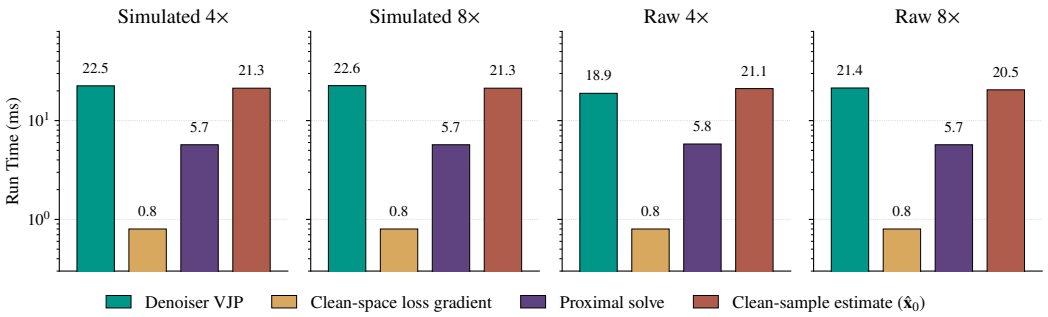

Figure 16: Per-step cost (ms per sample) of the four guidance operations on compressed sensing MRI, at $4\times$ and $8\times$ acceleration on simulated and raw measurements. Values are per-trajectory medians over all sampling timesteps.

Table 14 compares ProxDG against nine training-free guidance methods. Seven are InverseBench baselines, including the MRI-specific CSGM (Jalal et al., 2021) and ScoreMRI (Chung & Ye, 2022); we also evaluate SITCOM (Alkhouri et al., 2025) and DMPlug (Wang et al., 2024).

ProxDG attains the best PSNR and SSIM at both acceleration factors and in both measurement regimes. Its PSNR margin over the strongest baseline is 1.49 dB at $R = 4$ and 1.51 dB at $R = 8$ on simulated measurements, and 0.61 dB and 0.70 dB on raw measurements. ProxDG (PP) also reaches the lowest simulated data misfit at $R = 8$ (0.232), below RED-diff (0.485), which optimizes the data fit directly. ProxDG (PP) reduces the simulated data misfit of DS by $4\times$ at $R = 4$ and $8\times$ at $R = 8$. SITCOM achieves the lowest raw data misfit at $R = 8$ and the second-lowest at $R = 4$ without leading on PSNR or SSIM, so lower measurement residuals do not imply better reconstructions.

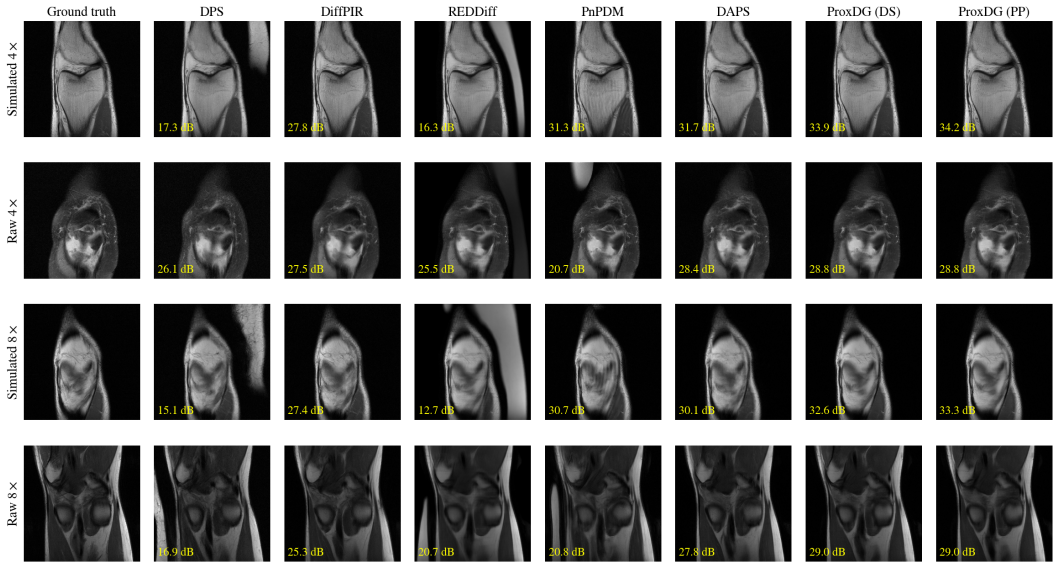

Figure 17: Compressed-sensing MRI. Reconstructions of selected knee slices from each method, alongside ground truth. Rows show simulated (noiseless) and raw measurements at $4\times$ and then $8\times$ acceleration. Each reconstruction is annotated with its PSNR in dB.

Table 14: Compressed-sensing MRI results. PSNR, SSIM, and data misfit of training-free guidance methods at acceleration factors $R = 4$ and $R = 8$, on simulated (noiseless) and raw measurements. Mean (std) over the 96-slice fastMRI knee test set. Best per column in bold, second-best underlined; ties at the reported precision share a rank. Baseline numbers other than SITCOM and DMPlug are reproduced from Zheng et al. (2025).

Subsampling ratio	×4						×8					
Measurement type	Simulated (noiseless)			Raw			Simulated (noiseless)			Raw		
Method	PSNR↑	SSIM↑	Data misfit↓	PSNR↑	SSIM↑	Data misfit↓	PSNR↑	SSIM↑	Data misfit↓	PSNR↑	SSIM↑	Data misfit↓
CSGM	28.78 (6.173)	0.710 (0.147)	1.518 (0.433)	25.17 (6.246)	0.582 (0.167)	31.642 (15.382)	26.15 (6.383)	0.625 (0.158)	1.142 (1.078)	21.17 (8.314)	0.425 (0.192)	22.088 (10.740)
ScoreMRI	25.97 (1.681)	0.468 (0.087)	10.828 (1.731)	25.60 (1.618)	0.463 (0.086)	33.697 (15.209)	25.01 (1.526)	0.405 (0.079)	8.360 (1.381)	24.74 (1.481)	0.403 (0.080)	24.028 (10.663)
RED-diff	29.36 (7.710)	0.733 (0.131)	**0.509** (0.077)	28.71 (2.755)	0.626 (0.126)	31.591 (15.368)	26.76 (6.696)	0.647 (0.124)	0.485 (0.068)	27.33 (2.441)	0.563 (0.117)	22.336 (10.838)
DiffPIR	28.31 (1.598)	0.632 (0.107)	10.545 (2.466)	27.60 (1.470)	0.624 (0.111)	34.015 (15.522)	26.78 (1.556)	0.588 (0.113)	7.787 (1.741)	26.26 (1.458)	0.590 (0.113)	24.208 (10.922)
DPS	26.13 (4.247)	0.620 (0.105)	9.900 (2.925)	25.83 (2.197)	0.548 (0.116)	35.095 (15.967)	20.82 (4.777)	0.536 (0.111)	6.737 (1.928)	23.00 (3.205)	0.507 (0.109)	24.842 (11.263)
DAPS	31.48 (1.988)	0.762 (0.089)	1.566 (0.390)	28.61 (2.197)	0.689 (0.102)	**31.115** (15.497)	29.01 (1.712)	0.681 (0.098)	1.280 (0.301)	27.10 (2.034)	0.629 (0.107)	22.729 (10.926)
PnP-DM	31.80 (3.473)	0.780 (0.096)	4.701 (0.675)	27.62 (3.425)	0.679 (0.117)	32.261 (15.169)	29.33 (3.081)	0.704 (0.105)	3.421 (0.504)	25.28 (3.102)	0.607 (0.117)	22.879 (10.712)
SITCOM	30.88 (2.920)	0.733 (0.096)	2.937 (0.203)	27.16 (1.604)	0.603 (0.088)	31.311 (15.442)	29.14 (2.037)	0.642 (0.085)	3.279 (0.289)	24.90 (2.568)	0.508 (0.133)	**21.949** (10.764)
DMPlug	24.81 (4.816)	0.530 (0.117)	11.098 (2.733)	26.12 (2.566)	0.539 (0.129)	34.337 (15.532)	24.31 (4.186)	0.527 (0.122)	8.407 (2.001)	23.89 (3.744)	0.526 (0.128)	24.278 (10.983)
ProxDG (DS)	33.05 (2.678)	0.787 (0.113)	2.112 (0.298)	**29.32** (1.981)	**0.696** (0.123)	31.989 (15.330)	30.56 (2.216)	0.711 (0.117)	1.818 (0.318)	27.98 (1.924)	**0.650** (0.131)	22.353 (10.773)
ProxDG (PP)	**33.29** (2.782)	**0.794** (0.115)	0.525 (0.069)	**29.32** (1.947)	0.695 (0.126)	31.986 (15.334)	**30.84** (2.217)	**0.725** (0.120)	**0.232** (0.041)	**28.03** (1.756)	0.649 (0.134)	22.348 (10.779)

I.2.5 Ablations

Denoiser backpropagation. ProxDG backpropagates the proximal step through the denoiser, whereas DiffPIR substitutes the proximal point for the clean estimate. To compare the two under the same solver, we equip DiffPIR with our proximal solver (Section I.2.2), which we call DiffPIR (P). ProxDG leads DiffPIR (P) on every metric in every setting (Figure 18). Its PSNR margin is 0.92 dB at $R = 4$ and 0.73 dB at $R = 8$ on simulated measurements, and 0.30 and 0.26 dB on raw measurements. ProxDG (PP) also keeps a 5 to 10× lower simulated data misfit (0.525 vs. 2.77 at $R = 4$ and 0.232 vs. 2.26 at $R = 8$). Our solver *does* account for most of DiffPIR's original deficit, raising its PSNR by 3.33 to 4.06 dB on simulated measurements and 1.42 to 1.51 dB on raw measurements.

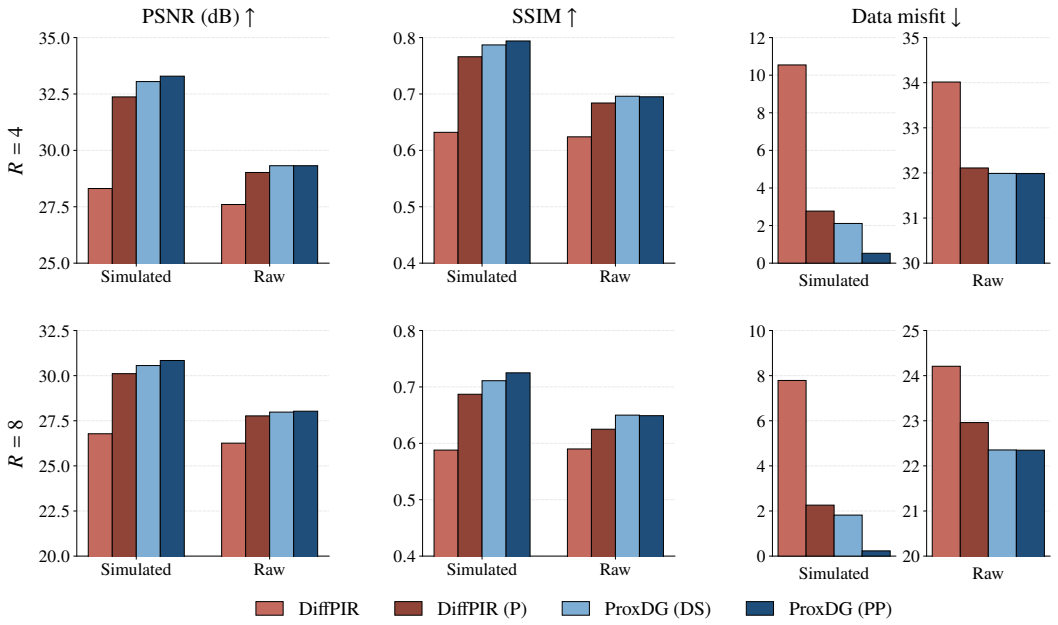

Figure 18: Effect of denoiser backpropagation on CS-MRI. PSNR, SSIM, and data misfit at $R = 4$ (top) and $R = 8$ (bottom), on simulated and raw measurements. DiffPIR (P) is DiffPIR with our proximal solver.

I.3 Black Hole Imaging

I.3.1 Problem Details

Black hole imaging reconstructs the sky-plane intensity of a supermassive black hole from sparse interferometric measurements taken by the Event Horizon Telescope (EHT), a global array of radio telescopes. Our setup follows InverseBench (Zheng et al., 2025), which details dataset preparation and diffusion model training. The model is trained on $50{,}000$ images from general relativistic magnetohydrodynamic (GRMHD) simulations (Mizuno, 2022). We evaluate observation-time ratios of 3%, 10%, and 100%.

Forward model. The unknown is a non-negative sky-plane intensity image $\mathbf{z} \in \mathbb{R}_{+}^{64\times 64}$, observed by radio telescopes spread across Earth via very long baseline interferometry (VLBI) (Chael et al., 2019). Each telescope pair samples one spatial Fourier coefficient of $\mathbf{z}$ at a frequency set by the pair's projected separation. At observation time t, pair (a, b) records this sample as the complex visibility

$$V_{a,b,t} = g_{a,t} g_{b,t} \exp(-i(\phi_{a,t} - \phi_{b,t}))\tilde{I}_{a,b,t}(\mathbf{z}) + \eta_{a,b,t} \tag{352}$$

where $\tilde{I}_{a,b,t}(\mathbf{z})$ is the ideal Fourier coefficient. Three error sources corrupt each visibility: unknown per-station amplitude gains $g_{\cdot,t}$, unknown per-station phases $\phi_{\cdot,t}$, and complex Gaussian thermal noise $\eta_{a,b,t}$ with estimated variance $\sigma^2_{a,b,t}$. The array samples only a sparse set of spatial frequencies, so recovering $\mathbf{z}$ is a super-resolution problem.

Closure quantities. Because the gains and phases are unknown, EHT imaging works with *closure quantities*: combinations of visibilities in which the calibration errors cancel. The closure phase of a telescope triangle (a, b, c) at time t is

$$\mathcal{C}_{a,b,c,t}(\mathbf{z}) = \angle\left(V_{a,b,t} V_{b,c,t} V_{c,a,t}\right) \tag{353}$$

where the station phases cancel around the triangle and the angle discards the gains. The log closure amplitude of a telescope quadruple (a, b, c, d) at time t is

$$\mathcal{A}_{a,b,c,d,t}(\mathbf{z}) = \log \frac{|V_{a,b,t}|\,|V_{c,d,t}|}{|V_{a,c,t}|\,|V_{b,d,t}|} \tag{354}$$

where each gain appears once in the numerator and once in the denominator, and the modulus discards the phases. Closure quantities also discard the global image scale, which the total flux $\mathbf{1}^\top \mathbf{z}$ supplies as an additional measurement. The full measurement $\mathbf{y} = (\mathbf{y}_{\text{cp}}, \mathbf{y}_{\text{logca}}, y_{\text{flux}})$ stacks the closure phases, the log closure amplitudes, and the total flux. The forward map $\mathbf{z} \mapsto \mathbf{y}$ is nonlinear in $\mathbf{z}$ through $\angle(\cdot)$ and $\log|\cdot|$.

Data-fit objective. Following InverseBench, the data-fit objective is the weighted χ^2 sum

$$\chi^2(\mathbf{z}) = w_{\text{cp}} \chi^2_{\text{cp}}(\mathbf{z}) + w_{\text{logca}} \chi^2_{\text{logca}}(\mathbf{z}) + w_{\text{flux}} \chi^2_{\text{flux}}(\mathbf{z}) \tag{355}$$

with weights $w_{\text{cp}} = w_{\text{logca}} = 1$ and $w_{\text{flux}} = 0.5$. Since closure phases are defined only modulo 2π, the closure-phase term takes the periodic form

$$\chi^2_{\text{cp}}(\mathbf{z}) = \frac{2}{N_{\text{cp}}} \sum_{i=1}^{N_{\text{cp}}} \frac{1 - \cos\left(\mathcal{C}_i(\mathbf{z}) - y_{\text{cp},i}\right)}{\sigma^2_{\text{cp},i}} \tag{356}$$

which reduces to the standard squared residual for small errors. The log closure amplitude and flux terms are normalized squared residuals,

$$\chi^2_{\text{logca}}(\mathbf{z}) = \frac{1}{N_{\text{logca}}} \sum_{i=1}^{N_{\text{logca}}} \frac{\left(\mathcal{A}_i(\mathbf{z}) - y_{\text{logca},i}\right)^2}{\sigma^2_{\text{logca},i}}, \qquad \chi^2_{\text{flux}}(\mathbf{z}) = \frac{\left(\mathbf{1}^\top \mathbf{z} - y_{\text{flux}}\right)^2}{\sigma^2_{\text{flux}}} \tag{357}$$

Here i indexes the closure measurements across all scans and telescope combinations. Each σ is the estimated standard deviation of its measurement.

Evaluation. We use the InverseBench test set of 100 synthetic 64×64 images. The GRMHD training set is not publicly released, so InverseBench samples these test images from a separately pre-trained diffusion model. The posterior is multimodal because of phase ambiguity and the need for super-resolution (Zheng et al., 2025). Each method therefore draws a batch of posterior samples per test image, and each metric scores the best sample in the batch. Closure quantities are translation-invariant, so we align each sample to the ground truth over a grid of sub-pixel shifts before computing PSNR. Blur PSNR repeats this protocol after blurring both images to the EHT's effective beam of $20\,\mu$as. For the χ^2 statistics, we rescale each sample to the observed total flux and fold each statistic as $\tilde{\chi}^2 = \max(\chi^2, 1/\chi^2)$, so under- and over-fitting both appear as values above the ideal 1.

I.3.2 Solver Details

Proximal subproblem. At each sampling step, ProxDG solves the proximal subproblem in Eq. 9 for the data-fit objective in Eq. 355,

$$\mathbf{z}^\star = \arg\min_{\mathbf{z}} \left\{ \frac{\lambda_t}{2}\chi^2(\mathbf{z}) + \frac{\rho}{2}\left\|\mathbf{z} - \hat{\mathbf{x}}_0\right\|_2^2 \right\} \tag{358}$$

where $\hat{\mathbf{x}}_0$ is the clean estimate. The forward map is ill-conditioned, so a single gradient step (as in standard gradient guidance) makes little progress on Eq. 358. We instead run a Gauss–Newton conjugate-gradient (GN-CG) inner loop.

Sum-of-squares reformulation. Gauss–Newton requires a residual vector $\mathbf{r}(\mathbf{z})$ with $\chi^2(\mathbf{z}) = \|\mathbf{r}(\mathbf{z})\|_2^2$. The log closure amplitude and flux terms in Eq. 355 are already sums of squares, with residuals

$$\mathbf{r}_{\text{logca}}(\mathbf{z}) = \sqrt{\tfrac{w_{\text{logca}}}{N_{\text{logca}}}}\,\frac{\mathcal{A}(\mathbf{z}) - \mathbf{y}_{\text{logca}}}{\sigma_{\text{logca}}}, \qquad r_{\text{flux}}(\mathbf{z}) = \sqrt{w_{\text{flux}}}\,\frac{\mathbf{1}^\top\mathbf{z} - y_{\text{flux}}}{\sigma_{\text{flux}}} \tag{359}$$

The closure-phase term in Eq. 356 is a sum of cosines rather than squares. Let $\mathcal{B}(\mathbf{z})$ collect the bispectra $V_{a,b,t}V_{b,c,t}V_{c,a,t}$ of all telescope triangles, so that $\mathcal{C}(\mathbf{z}) = \angle\mathcal{B}(\mathbf{z})$. We define the complex residual

$$\mathbf{r}_{\text{cp}}(\mathbf{z}) = \sqrt{\tfrac{w_{\text{cp}}}{N_{\text{cp}}}}\,\frac{1}{\sigma_{\text{cp}}}\left(\exp(i\mathbf{y}_{\text{cp}})\frac{\overline{\mathcal{B}(\mathbf{z})}}{|\mathcal{B}(\mathbf{z})|} - 1\right) \tag{360}$$

with all operations elementwise, and real and imaginary parts stacked into one real vector. Since $\overline{\mathcal{B}}/|\mathcal{B}| = \exp(-i\mathcal{C})$ and $|\exp(i\delta) - 1|^2 = 2(1 - \cos\delta)$,

$$\|\mathbf{r}_{\text{cp}}(\mathbf{z})\|_2^2 = \frac{w_{\text{cp}}}{N_{\text{cp}}}\sum_{k=1}^{N_{\text{cp}}}\frac{|\exp(i(y_{\text{cp},k} - \mathcal{C}_k(\mathbf{z}))) - 1|^2}{\sigma_{\text{cp},k}^2} = w_{\text{cp}}\chi_{\text{cp}}^2(\mathbf{z}) \tag{361}$$

The full residual $\mathbf{r}(\mathbf{z}) = (\mathbf{r}_{\text{cp}}, \mathbf{r}_{\text{logca}}, r_{\text{flux}})$ therefore satisfies $\|\mathbf{r}(\mathbf{z})\|_2^2 = \chi^2(\mathbf{z})$ exactly, with no small-angle approximation.

Gauss–Newton normal equations. Substituting the linearization $\mathbf{r}(\mathbf{z} + d\mathbf{z}) \approx \mathbf{r}(\mathbf{z}) + \mathbf{J}(\mathbf{z})\,d\mathbf{z}$ into Eq. 358 and minimizing over $d\mathbf{z}$ yields the damped normal equations

$$\left(\lambda_t\mathbf{J}(\mathbf{z})^\top\mathbf{J}(\mathbf{z}) + \rho\mathbf{I}\right)\,d\mathbf{z} = -\left(\lambda_t\mathbf{J}(\mathbf{z})^\top\mathbf{r}(\mathbf{z}) + \rho(\mathbf{z} - \hat{\mathbf{x}}_0)\right) \tag{362}$$

The proximal term $\rho\mathbf{I}$ acts as built-in Levenberg–Marquardt damping, so the system matrix $\mathbf{Q} = \lambda_t\mathbf{J}^\top\mathbf{J} + \rho\mathbf{I}$ is symmetric positive definite for any $\rho > 0$. The system has one unknown per pixel, so we solve it with conjugate gradient (CG) and never form $\mathbf{J}^\top\mathbf{J}$. Each CG iteration only needs the product

$$\mathbf{Q}\mathbf{p} = \lambda_t\mathbf{J}(\mathbf{z})^\top\left(\mathbf{J}(\mathbf{z})\mathbf{p}\right) + \rho\mathbf{p} \tag{363}$$

which costs one forward-mode Jacobian-vector product (JVP) and one reverse-mode vector-Jacobian product (VJP) through $\mathbf{r}$.

Algorithm. The solver runs one setup step, then K Gauss–Newton iterations of steps 2–4.

1. **Observation cache.** At the first diffusion step, cache the closure components of $\mathbf{y}$, the closure-phase phasor $\exp(i\mathbf{y}_{\text{cp}})$, and the estimated standard deviations and residual weights. All later diffusion steps reuse this cache.

2. **Residual and gradient.** Evaluate the full residual $\mathbf{r}(\mathbf{z})$, then form the right-hand side of Eq. 362 with one VJP through $\mathbf{r}$.
3. **CG solve.** Run up to M CG iterations on the system in Eq. 362, each costing one JVP and one VJP. Stop early once the relative CG residual falls below a tolerance ϵ.
4. **Update and clamp.** Set $\mathbf{z} \leftarrow \mathbf{z} + d\mathbf{z}$ and clamp $\mathbf{z}$ at zero to keep it non-negative.

Implementation. We tune ProxDG as in Section I.2.3 and add two solver hyperparameters to the search space. The first is the number of Gauss–Newton iterations $K \in \{1, \ldots, 5\}$. The second is the CG initialization, which either resets $d\mathbf{z}$ to $\mathbf{0}$ at each diffusion step or warm-starts it from the previous step's solution. All results in Table 15 fix $M = 5$ and $\epsilon = 10^{-4}$ and use the K and CG initialization selected by the sweep.

I.3.3 Results

Sampling cost. Figure 19 reports the time of each guidance operation on black hole imaging. Appendix I.2.4 reports end-to-end sampling time (Table 12) and the operations each method uses (Table 13). ProxDG is the slowest method on BHI because we give it a large solver budget. We choose this budget deliberately, to show that ProxDG turns extra solver compute into better reconstructions (Appendix I.3.4). The budget alone does not explain ProxDG's lead. Other samplers given the same solver and budget trail ProxDG at the 10% and 100% ratios and on data fit throughout (Figure 22). On FWI, we restrict ProxDG to a very small compute budget, and it still leads every baseline on every metric while running fastest (Appendix I.4.3).

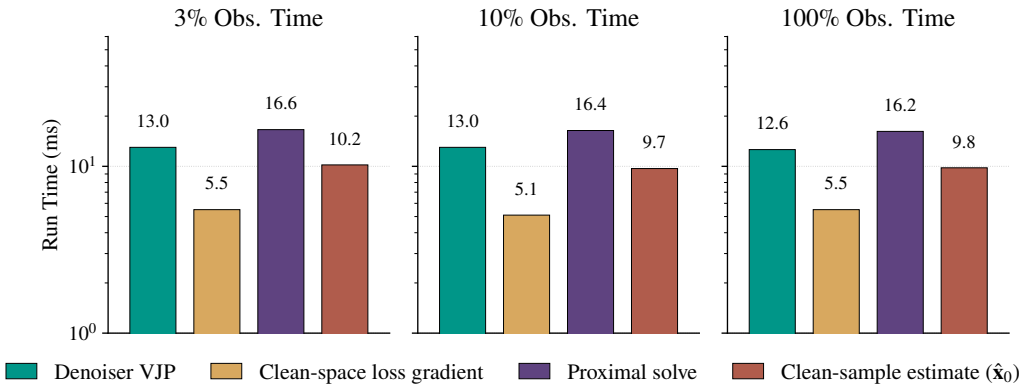

Figure 19: Per-step cost (ms per sample) of the four guidance operations on black hole imaging, at EHT observation-time ratios of 3%, 10%, and 100%. Values are per-trajectory medians over all sampling timesteps. The proximal solve is timed at $K = M = 1$.

Table 15: Full black hole imaging results at all three EHT observation-time ratios, mean (std) over the 100-image InverseBench test set. Higher is better for PSNR and Blur PSNR, while $\tilde{\chi}^2_{\text{cp}}$ and $\tilde{\chi}^2_{\text{logca}}$ are ideal near 1. Best per column bold, second-best underlined. Baseline numbers other than SITCOM and DMPlug are reproduced from Zheng et al. (2025).

Observation time ratio	3%				10%				100%			
Method	PSNR↑	Blur PSNR↑	$\tilde{\chi}^2_{\text{cp}}$	$\tilde{\chi}^2_{\text{logca}}$	PSNR↑	Blur PSNR↑	$\tilde{\chi}^2_{\text{cp}}$	$\tilde{\chi}^2_{\text{logca}}$	PSNR↑	Blur PSNR↑	$\tilde{\chi}^2_{\text{cp}}$	$\tilde{\chi}^2_{\text{logca}}$
DPS	24.20 (3.72)	30.83 (5.58)	8.024 (24.336)	5.007 (5.750)	24.36 (3.72)	30.79 (5.75)	13.052 (43.087)	6.614 (26.789)	25.86 (3.90)	32.94 (6.19)	8.759 (37.784)	5.456 (24.185)
LGD	22.51 (3.76)	28.50 (5.49)	15.825 (16.838)	12.862 (12.663)	22.08 (3.75)	27.48 (5.09)	10.775 (21.684)	13.375 (56.397)	21.22 (3.64)	26.06 (4.98)	13.239 (17.231)	13.233 (39.107)
RED-diff	20.74 (2.62)	26.10 (3.35)	6.713 (6.925)	9.128 (19.052)	22.53 (3.02)	27.67 (4.53)	2.488 (2.925)	4.916 (13.221)	23.77 (4.13)	29.13 (6.22)	1.853 (0.938)	2.050 (2.361)
PnP-DM	24.25 (3.45)	30.49 (4.93)	2.201 (1.352)	1.668 (0.551)	24.57 (3.47)	30.80 (5.22)	1.433 (0.417)	1.336 (0.478)	26.07 (3.70)	32.88 (6.02)	1.311 (0.195)	**1.199** (0.221)
DAPS	23.54 (3.28)	29.48 (4.88)	3.647 (3.287)	2.329 (1.354)	23.99 (3.56)	30.10 (5.13)	1.545 (0.705)	2.253 (9.903)	25.60 (3.64)	32.78 (5.68)	1.300 (0.324)	1.229 (0.532)
DiffPIR	24.12 (3.25)	30.45 (4.88)	14.085 (14.105)	10.545 (8.860)	23.84 (3.59)	30.04 (5.03)	5.374 (3.733)	5.205 (5.556)	25.01 (4.64)	31.86 (6.56)	3.271 (1.623)	2.970 (1.202)
SITCOM	25.12 (3.19)	32.20 (4.76)	1.648 (0.682)	1.918 (0.864)	25.49 (3.69)	32.56 (5.28)	1.415 (0.477)	2.179 (2.055)	27.68 (3.04)	35.49 (3.91)	1.285 (1.378)	1.385 (0.743)
DMPlug	25.11 (3.32)	31.65 (5.33)	2.201 (1.597)	1.955 (0.893)	25.40 (3.97)	32.16 (6.17)	1.392 (0.537)	1.348 (0.299)	25.87 (4.22)	32.82 (6.05)	1.277 (1.528)	1.504 (1.780)
ProxDG (DS)	25.19 (3.59)	**32.28** (5.27)	**1.555** (0.689)	**1.595** (0.599)	**26.16** (3.63)	**33.52** (5.48)	**1.207** (0.265)	**1.324** (0.243)	**28.01** (2.91)	36.00 (4.39)	1.234 (0.928)	1.259 (0.242)
ProxDG (PP)	**25.59** (3.55)	32.09 (5.28)	1.686 (1.295)	1.891 (0.812)	26.10 (3.24)	33.32 (5.28)	1.268 (0.255)	1.557 (0.314)	27.95 (3.00)	**36.16** (4.73)	**1.130** (0.401)	1.385 (0.167)

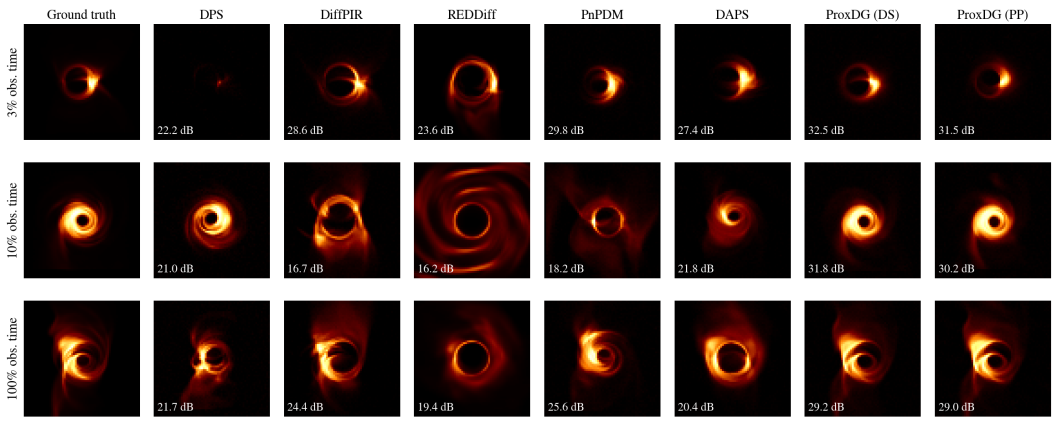

Figure 20: Black hole imaging. Reconstructions from each method, alongside ground truth. Rows show selected test images at EHT observation-time ratios of 3%, 10%, and 100%, from top to bottom. Each reconstruction is annotated with its PSNR in dB.

I.3.4 Ablations

λ_t schedules. Our schedule $\lambda_t = \beta\sigma_t^2$ satisfies Assumption C.1 but is not optimal. As shown in Appendix B.4.2, the optimal λ_t is the posterior covariance $\mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t]$, which adapts the proximal penalty to local manifold geometry. We therefore hypothesize that sample quality improves as λ_t more closely approximates $\mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t]$. To test our hypothesis, we compare against two alternative schedules. The first, adopted from DAPS (Zhang et al., 2025a), sets $\lambda_t = \beta\sigma_t^2/\alpha_t^2$: the conditional variance of the rescaled sample $\mathbf{x}_t/\alpha_t$. This schedule anneals more aggressively than ours and diverges as $\alpha_t \to 0$. The second, adopted from DTMPD (Boys et al., 2024), replaces the isotropic penalty with a per-pixel weighting $d_t \approx \mathrm{diag}(\mathrm{Cov}[\mathbf{x}_0|\mathbf{x}_t])$.

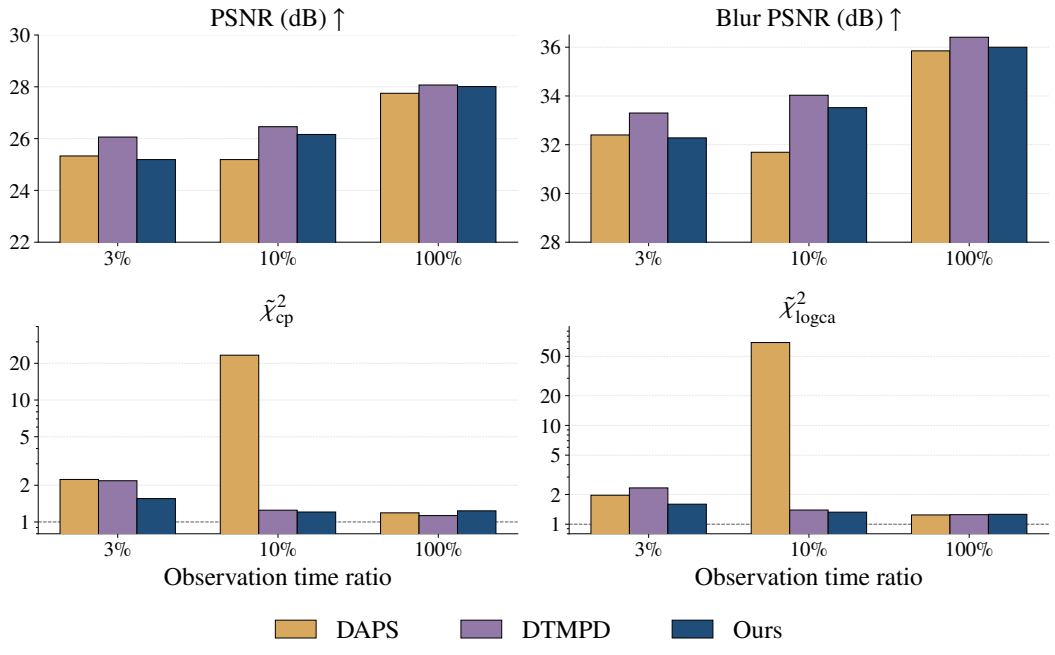

Figure 21: Comparison of three λ_t schedules for ProxDG (DS) across three EHT observation-time ratios. Metrics are averaged over the 100-image InverseBench test set. The $\tilde{\chi}^2$ panels (bottom) use a log scale, with the dashed line at 1.

The DAPS schedule never improves on ours and breaks down at 10%, where PSNR drops by 0.97 dB and $\tilde{\chi}^2_{\text{cp}}$ rises to 23.32. DTMPD matches or improves PSNR and Blur PSNR at every ratio, with gains of up to 0.87 dB PSNR and 1.60 dB Blur PSNR at 3%. It also attains the best $\tilde{\chi}^2_{\text{cp}}$ (1.128) at 100%, although its $\tilde{\chi}^2$ statistics sit slightly above 1 at 3% and 10%. As hypothesized, closer approximations of $\text{Cov}[\mathbf{x}_0|\mathbf{x}_t]$ yield better samples.

Equal-compute comparisons. At equal compute, we ablate two factors behind ProxDG's performance. To isolate denoiser backpropagation, DiffPIR (P) equips DiffPIR with our GN-CG solver from Section I.3.2. To isolate solver quality, DiffPIR (GD) and DPS (GD) replace the solver with gradient descent capped at ProxDG's gradient budget. With $K = M = 5$, our solver evaluates up to $K(M + 1) = 30$ gradients per diffusion step, so we tune and test the DiffPIR (GD) and DPS (GD) with up to 30 gradient steps. The cap favors DPS (GD), which uses up to $30\times$ more denoiser calls than ProxDG. Figure 22 reports the results.

The solver, not the sampler, determines data fit. With gradient descent, DiffPIR (GD) and DPS (GD) attain $\tilde{\chi}^2$ between 2.2 and 76.5. The additional gradient descent steps often worsen data fit, since Eq. 358 is an ill-conditioned objective. With our solver, DiffPIR (P) and ProxDG (PP) keep $\tilde{\chi}^2$ below 1.9 at every ratio. Our solver also improves the PSNR of DiffPIR (P) by 0.62 to 2.12 dB. However, DiffPIR (P) still trails ProxDG (PP) by 0.45 to 0.85 dB PSNR and 0.41 to 1.54 dB Blur PSNR. This gap in performance can only be explained by denoiser backpropagation, the one step that separates ProxDG (PP) from DiffPIR (P).

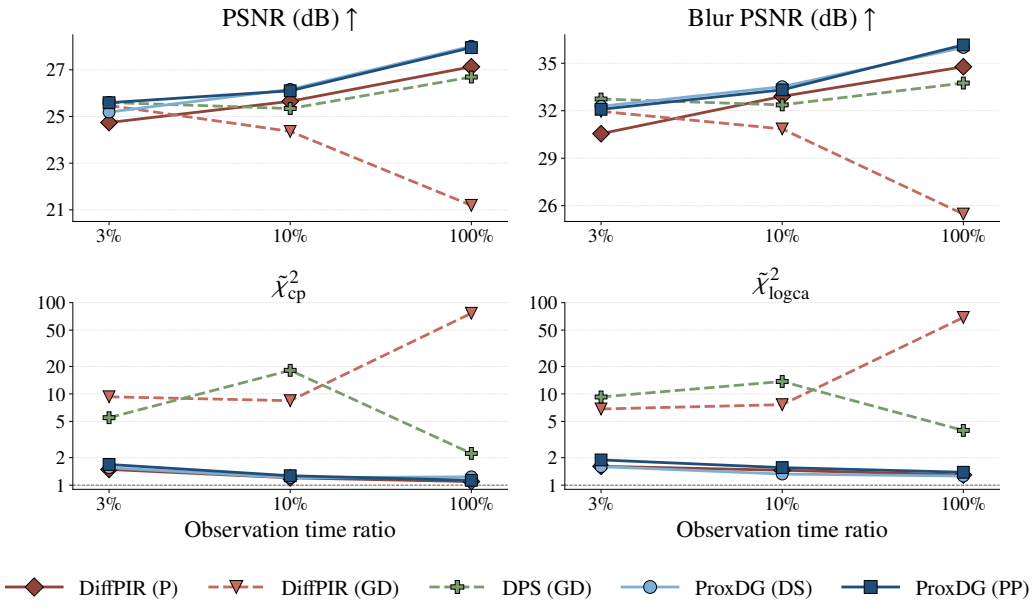

Figure 22: Equal-compute comparisons on black hole imaging across EHT observation-time ratios. Metrics are averaged over the 100-image InverseBench test set. The $\tilde{\chi}^2$ panels (bottom) use a log scale, with the dashed line at 1.

Solver strength. To isolate the effect of solver strength on sample quality, we sweep the number of Gauss–Newton iterations $K \in 1, 2, 3, 4, 5$ of our GN-CG solver (Section I.3.2) at a fixed $M = 1$ CG iterations. At the cheapest setting $K = M = 1$, the solver reduces to a single preconditioned gradient step on the proximal objective.

Stronger solvers fit the data better at every observation ratio. As K grows from 1 to 5, $\tilde{\chi}^2_{\text{cp}}$ and $\tilde{\chi}^2_{\text{logca}}$ decrease monotonically toward the ideal value of 1. PSNR rises by 1.57 and 1.97dB at 10% and 100%, but saturates by $K = 4$ at 3%. Setting $K = M = 5$, as in Table15, amplifies these gains. Relative to $K = M = 1$, PSNR rises by 0.99, 2.39, and 4.05 dB at the 3%, 10%, and 100% ratios, and $\tilde{\chi}^2_{\text{cp}}$ falls to between 1.21 and 1.56.

Note: At $K = M = 1$, ProxDG attains up to $4\times$ lower $\tilde{\chi}^2_{\text{cp}}$ than DPS, yet trails it in PSNR at the 10% and 100% ratios. Two factors explain this gap. First, the DPS results in Table 15 are hyperparameter-tuned, whereas we run the solver-strength ablation without tuning. Second, BHI is highly degenerate. Even at the 100% ratio, many distinct images fit the sparse EHT measurements equally well. A reconstruction can thus fit the data closely yet differ from the ground truth, which lowers its PSNR. Zheng et al. (2025) report the same effect.

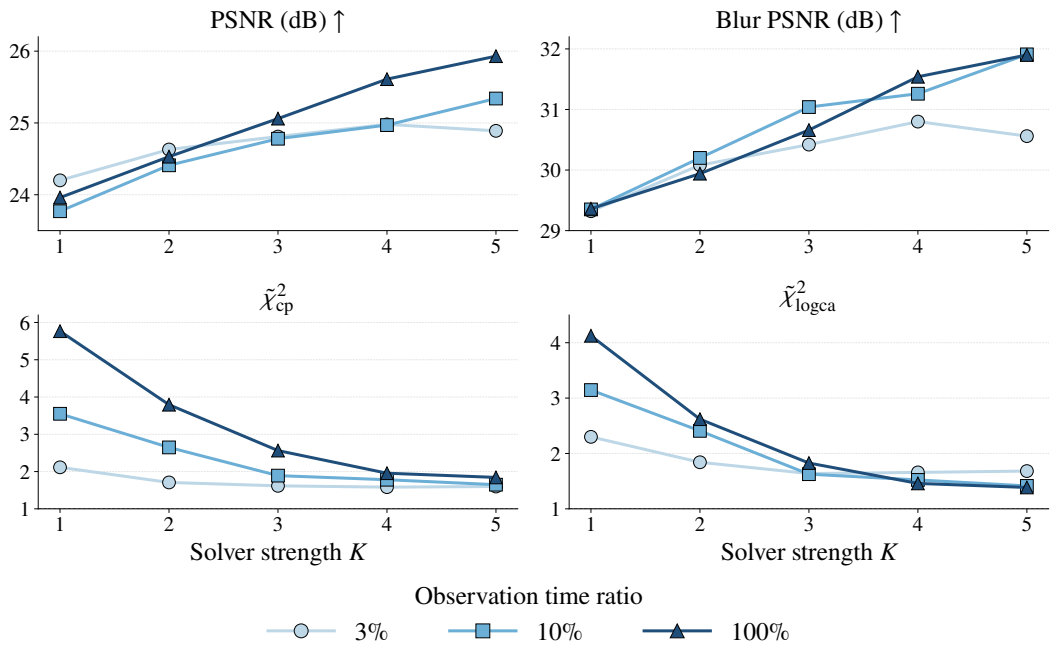

Figure 23: ProxDG (DS) under solver strengths $K \in \{1, 2, 3, 4, 5\}$ at fixed $M = 1$, across three EHT observation-time ratios. Metrics are averaged over the 100-image InverseBench test set.

λ_t floor.

Flooring the λ_t schedule. Our schedule floors λ_t at a positive $\lambda_{\min}$ once t falls below $t_{\min}$, which keeps guidance active through the final denoising steps. To isolate the floor's effects, we remove it and let λ_t anneal to zero (Figure 24). The floor improves PSNR and Blur PSNR at the 3% and 10% ratios and matches the unfloored schedule at 100%. It also lowers $\tilde{\chi}^2_{\text{cp}}$ at every ratio, although its effect on $\tilde{\chi}^2_{\text{logca}}$ is mixed. We therefore keep the floor at all three ratios.

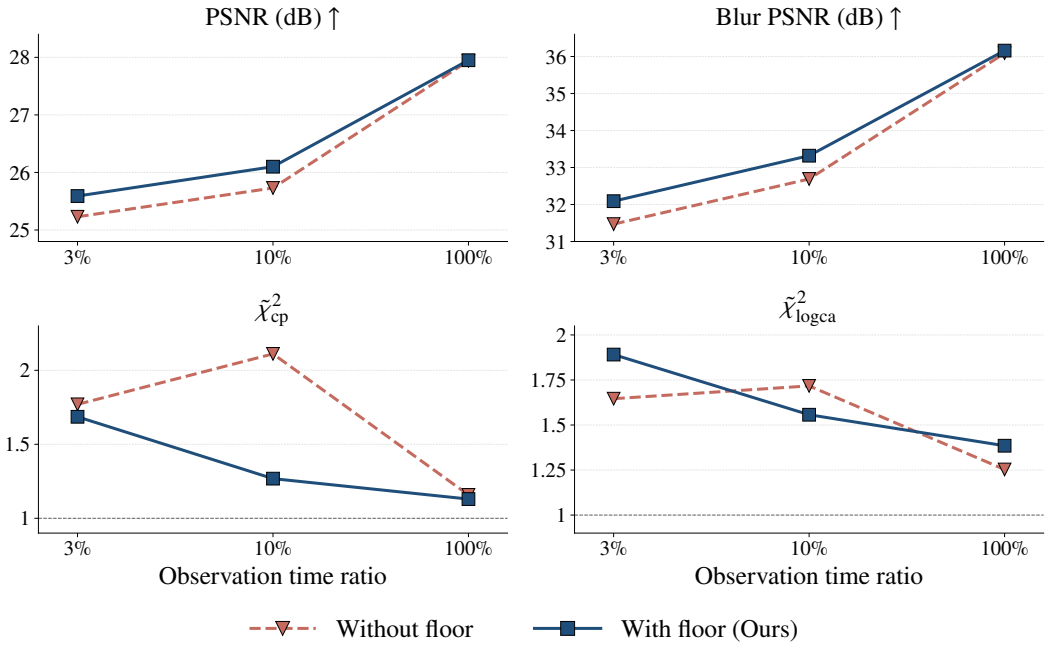

Figure 24: ProxDG (DS) with and without the λ_t floor across three EHT observation-time ratios. Metrics are averaged over the 100-image InverseBench test set.

I.4 Full Waveform Inversion

I.4.1 Problem Details

All problem-specific choices follow InverseBench (Zheng et al., 2025), including the training data and the pre-trained diffusion prior. We briefly summarize the forward model and our evaluation protocol.

Forward model. Full waveform inversion recovers a subsurface velocity field $v(x)$, discretized as $\mathbf{v} \in \mathbb{R}_+^{128\times128}$, from seismic waveforms recorded near the surface. Each shot s fires a source q_s, and the resulting pressure wavefield $u_s(x,t)$ satisfies the acoustic wave equation

$$\frac{1}{v(x)^2}\frac{\partial^2 u_s}{\partial t^2} - \nabla^2 u_s = q_s(x,t) \tag{364}$$

The receivers record $\mathbf{y}_s = \mathbf{P}_s u_s$, where $\mathbf{P}_s$ samples the wavefield at the receiver locations, and the full measurement stacks all S shots as $\mathbf{y} = (\mathbf{y}_s)_{s=1}^S$. We solve Eq. 364 by explicit finite-difference time stepping, which is stable only under the Courant–Friedrichs–Lewy (CFL) condition $\Delta t \le C\min(\Delta x, \Delta z)/\max_x v(x)$, where C is a constant set by the stencil.

Acquisition geometry. The 128×128 velocity grid uses $\Delta x = 20\,\text{m}$ horizontal and $\Delta z = 10\,\text{m}$ vertical spacing, covering a physical domain of $2.54\,\text{km} \times 1.27\,\text{km}$. The top edge is a free surface, and absorbing layers damp reflections from the other three edges. We place $R = 129$ receivers evenly at 10 m depth and $S = 16$ sources evenly at 1270 m depth, each emitting a 5 Hz Ricker wavelet. Each shot is recorded for 1000 ms at $\Delta t = 1\,\text{ms}$ time steps, starting from rest. This geometry follows InverseBench (Zheng et al., 2025).

Adjoint-state gradients. Guidance requires the gradient of the data-fit loss

$$\mathcal{L}(\mathbf{v}) = \frac{1}{2}\sum_{s=1}^{S} \|\mathbf{y}_s - \mathbf{P}_s u_s(\mathbf{v})\|_2^2 \tag{365}$$

with respect to the velocity. As in InverseBench, all wave solves use the finite-difference framework Devito (Louboutin et al., 2019), whose compiled kernels do not support automatic differentiation. We instead follow standard FWI practice and compute the gradient with the adjoint-state method (Virieux & Operto, 2010). Each shot requires one forward solve and one backward-in-time solve driven by the data residual, which together give the gradient $\nabla_m \mathcal{L}$ with respect to the squared slowness $m = 1/v^2$. The elementwise chain rule $\nabla_{\mathbf{v}}\mathcal{L} = -2\mathbf{v}^{-3} \circ \nabla_m \mathcal{L}$ then converts it to a gradient in velocity.

Evaluation. We use the InverseBench test set of 10 velocity maps drawn from the CurveFaultB dataset. For each test case, we report the relative ℓ_2 error $\|\hat{\mathbf{v}} - \mathbf{v}^*\|_2/\|\mathbf{v}^*\|_2$, PSNR and SSIM after rescaling velocities from $[1.5, 4.5]$ km/s to $[0, 1]$, and the data misfit $\sqrt{\mathcal{L}(\hat{\mathbf{v}})}$ between observed and predicted recordings. All numbers are mean (std) over the test set.

I.4.2 Solver Details

At each sampling step, ProxDG solves the proximal subproblem in Eq. 9 for the data-fit loss in Eq. 365,

$$\mathbf{v}^\star = \arg\min_{\mathbf{v}} \left\{ \lambda_t \mathcal{L}(\mathbf{v}) + \frac{\rho}{2}\|\mathbf{v} - \hat{\mathbf{x}}_0\|_2^2 \right\} \tag{366}$$

where $\hat{\mathbf{x}}_0$ is the clean estimate. Let $\mathbf{r}(\mathbf{v}) = \mathbf{y} - \mathbf{P}u(\mathbf{v})$ be the residual over all shots. Then $\mathcal{L}(\mathbf{v}) = \frac{1}{2}\|\mathbf{r}(\mathbf{v})\|_2^2$, and Eq. 366 is a smooth nonlinear least-squares problem. We solve it with the same GN-CG inner loop as in black hole imaging (Section I.3.2), which repeatedly linearizes $\mathbf{r}$ and solves the damped normal equations with CG. The rest of this section covers only the FWI-specific components.

Jacobian-vector products. Devito's compiled kernels do not support automatic differentiation (Section I.4.1), so we replace the BHI solver's forward- and reverse-mode autograd calls with wave-equation solves. As in black hole imaging, $\mathbf{J}$ denotes the Jacobian of the residual $\mathbf{r}$. The

vector-Jacobian product $\mathbf{J}^\top \mathbf{r}$ is then exactly the adjoint-state gradient $\nabla_{\mathbf{v}} \mathcal{L}$ of Section I.4.1, which costs one forward and one adjoint solve per shot. The forward solve also stores the background wavefield $u_{0,s}(x,t)$ of each shot for reuse. For a velocity perturbation $\mathbf{p}$, a linearized (Born) forward solve around $u_{0,s}$ propagates $\mathbf{p}$ to the receivers to form $\mathbf{Jp}$, and an adjoint solve applies $\mathbf{J}^\top$ to the result. The Hessian-vector product $\mathbf{J}^\top \mathbf{Jp}$ therefore costs one additional solve pair per shot.

Pseudo-Hessian preconditioning. The Gauss–Newton Hessian $\mathbf{J}^\top \mathbf{J}$ is ill-conditioned in FWI (Virieux & Operto, 2010), so we precondition CG with the pseudo-Hessian of Shin et al. (2001). This diagonal approximation of $\mathbf{J}^\top \mathbf{J}$ is built from the illumination

$$\mathrm{illum}(x) = \sum_{s=1}^{S} \sum_{t} u_{0,s}(x,t)^2 \tag{367}$$

which reuses the cached background wavefields $u_{0,s}$ at no extra cost. Converting to velocity and adding the proximal damping gives the preconditioner

$$\Pi = \left(\lambda_t (2\mathbf{v}^{-3})^2 \circ \mathrm{illum} + \rho + \varepsilon\right)^{-1} \tag{368}$$

with elementwise operations and $\varepsilon = 10^{-10}$ to avoid division by zero.

Stability and box constraints. The simulator runs at a fixed time step, so the CFL condition of Section I.4.1 caps the largest velocity the simulator can handle; a Gauss–Newton update that exceeds this cap makes the wave solve diverge. We therefore project the iterate onto a box $[v_{\min}, v_{\max}]$ after each Gauss–Newton iteration. The box brackets the training-set velocities and stays within the CFL-stable regime.

Implementation. We compute the background wavefields once per Gauss–Newton iteration and reuse them in every Hessian-vector product of the CG loop. We parallelize over the $S = 16$ shots with one Dask worker per shot, and each worker releases its wavefield $u_{0,s}$ once the CG solve completes. We tune hyperparameters with the same sweep procedure as in CS-MRI (Section I.2.3). To keep compute tractable, ProxDG uses 200 diffusion steps instead of 1000. All results in Table 16 use $K = M = 1$, which reduces each proximal solve to a single pseudo-Hessian-preconditioned gradient step on Eq. 366.

I.4.3 Results

Sampling cost. Figure 25 reports the time of each guidance operation on full waveform inversion. Appendix I.2.4 reports end-to-end sampling time (Table 12) and the operations each method uses (Table 13). On FWI, ProxDG is both the fastest and the best-performing method. A single proximal solve costs about $2.8\times$ a loss gradient, since it adds one Born and one adjoint solve per shot. We offset this cost by running only 200 diffusion steps with $K = M = 1$. ProxDG then samples in 29.9 minutes, about $3\times$ faster than the fastest baseline (DPS, 93.3 minutes).

Table 16: Full waveform inversion results. Relative ℓ_2, PSNR, SSIM, and data misfit, mean (std) over the 10-map InverseBench test set. Best per column in bold, second-best underlined. †: initialized from the ground truth blurred by a Gaussian filter with $\sigma = 20$. Baseline numbers other than SITCOM and DMPlug are reproduced from Zheng et al. (2025).

Method	Relative $\ell_2 \downarrow$	PSNR ↑	SSIM ↑	Data misfit ↓
DPS	0.250 (0.154)	14.111 (6.820)	0.491 (0.161)	155.08 (92.17)
LGD	0.244 (0.024)	12.288 (0.889)	0.341 (0.047)	258.47 (26.40)
DiffPIR	0.204 (0.129)	16.113 (6.962)	0.554 (0.191)	88.53 (56.91)
DAPS†	0.201 (0.103)	14.914 (4.184)	0.321 (0.067)	111.13 (71.33)
PnP-DM	0.259 (0.075)	11.983 (2.269)	0.431 (0.073)	308.84 (26.34)
RED-diff	0.319 (0.102)	10.372 (2.650)	0.280 (0.108)	94.67 (41.33)
SITCOM	0.387 (0.259)	12.643 (7.799)	0.405 (0.216)	99.21 (67.05)
DMPlug	0.256 (0.154)	13.942 (6.380)	0.505 (0.173)	133.20 (73.74)
ProxDG (DS)	<u>0.140</u> (0.050)	<u>17.527</u> (3.328)	<u>0.602</u> (0.095)	<u>76.02</u> (30.12)
ProxDG (PP)	**0.132** (0.054)	**18.202** (3.602)	**0.637** (0.096)	**55.17** (31.41)

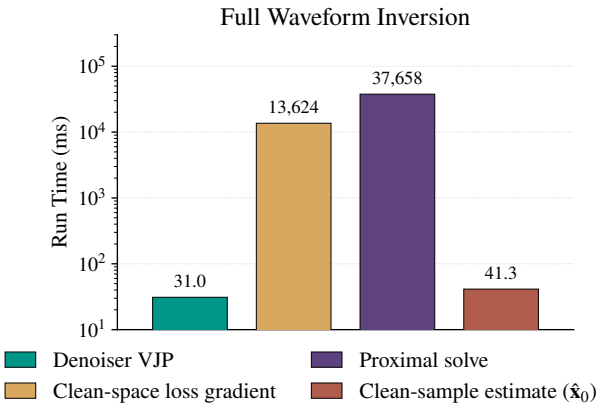

Figure 25: Per-step cost (ms per sample) of the four guidance operations on full waveform inversion. Values are per-trajectory medians over all sampling timesteps.

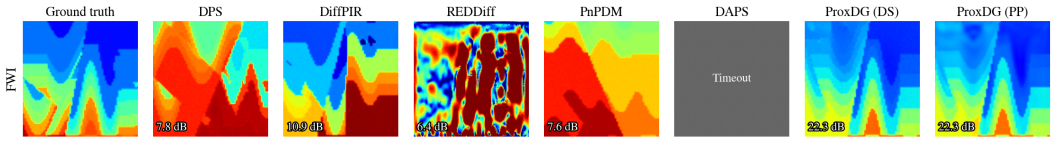

Figure 26: Full waveform inversion. Seismic velocity maps reconstructed by each method for a selected test example, alongside ground truth. Each reconstruction is annotated with its PSNR in dB. The gray panel indicates that DAPS timed out.

I.4.4 Ablations

Equal-compute comparisons. As in black hole imaging (Section I.3.4), we ablate two factors behind ProxDG's performance. To isolate denoiser backpropagation, DiffPIR (P) equips DiffPIR with our GN-CG solver from Section I.4.2 at $K = M = 1$. To isolate solver quality, DiffPIR (GD) and DPS (GD) replace the solver with gradient descent capped at ProxDG's per-step gradient budget. The baselines run 1000 diffusion steps against ProxDG's 200, which gives them five times ProxDG's total solver budget. Table 17 reports the results.

Table 17: Equal-compute comparisons on full waveform inversion, averaged over the 10-map InverseBench test set. Baselines run 1000 diffusion steps, and ProxDG runs 200. Best per column in bold, second-best underlined.

Method	Relative $\ell_2 \downarrow$	PSNR ↑	SSIM ↑	Data misfit ↓
DiffPIR (P)	0.169	16.49	0.539	127.5
DiffPIR (GD)	0.248	12.31	0.398	304.2
DPS (GD)	0.194	14.25	0.476	288.6
ProxDG (DS)	<u>0.140</u>	<u>17.53</u>	<u>0.602</u>	<u>76.0</u>
ProxDG (PP)	**0.132**	**18.20**	**0.637**	**55.2**

The solver determines data fit. With gradient descent, DiffPIR (GD) and DPS (GD) reach a data misfit of 304.2 and 288.6, four to five times that of either ProxDG variant. With our solver, DiffPIR (P) cuts the misfit to 127.5 but still trails both ProxDG variants on every metric, by 1.04 to 1.71 dB PSNR. As in CS-MRI and BHI, denoiser backpropagation improves reconstruction beyond what the solver alone provides.

I.5 Robot Motion Planning

I.5.1 Planning as Inference

We evaluate ProxDG as a trajectory sampler rather than a trajectory optimizer for three reasons.

Optimizers get trapped. Trajectory optimizers initialized from a straight line routinely converge to infeasible critical points (Carvalho et al., 2023). Sampling-based planners such as RRT and PRM avoid this through probabilistic completeness (Karaman & Frazzoli, 2011) and, in their optimal variants, asymptotic optimality (Gammell et al., 2014). Diffusion priors inherit a softer version of this property, spreading samples across the homotopy classes in their training demonstrations (Janner et al., 2022; Carvalho et al., 2023; Chi et al., 2023).

Deployment is closed-loop. A sampling-based MPC stack (Williams et al., 2018) replans at roughly 10 Hz as the scene changes (Janner et al., 2022; Chi et al., 2023). When the previous plan becomes infeasible, a planner that returned N candidates can switch to the best surviving one in $\mathcal{O}(N)$ collision checks, whereas a single-optimum planner must restart. For the same reason, we measure diversity with the Vendi score (Friedman & Dieng, 2023), which counts the effective number of distinct trajectories and does not reward near-duplicates of a single mode.

Selection criteria change at runtime. The planning objective rarely captures everything the deployed system cares about, such as the legibility of motion to a human collaborator (Dragan et al., 2013) or contingency margins against agents with uncertain intent (Chen et al., 2022). A sampler lets a downstream module re-rank its candidates by such criteria, whereas a single-optimum planner cannot.

We therefore evaluate constraint-handling methods on the joint quality of a population of N trajectories rather than on a single best one, following the planning-as-inference perspective of Toussaint (2009) and Mukadam et al. (2018); Dong et al. (2016; 2018).

I.5.2 Problem Details

Environments and robots. We evaluate on the six benchmarks of MPD (Carvalho et al., 2023). Simple2D and NarrowPassageDense2D use a 2-DoF planar point mass, Planar2Link and Planar4Link use 2- and 4-DoF revolute chains, and Spheres3D and Warehouse use the 7-DoF Franka Panda. In every environment, we test on the ExtraObjects scene, which adds obstacles the diffusion prior never saw during training. The prior alone cannot avoid these obstacles, so differences in validity reflect how each method handles constraints.

Trajectory parameterization and prior. Each method outputs a parametric trajectory τ, a vector of evenly spaced waypoints. A deterministic map Φ decodes τ into a continuous joint-space path $q(s) = \Phi(\tau)(s)$ on the phase variable $s \in [0, 1]$ over a fixed horizon $T = 10$ s. The map Φ is constructed so that $q(0) = q_{\text{start}}$, $q(1) = q_{\text{goal}}$, and $q'(0) = q'(1) = 0$ hold exactly. The diffusion prior is the conditional Temporal U-Net of MPD trained per environment on RRT-Connect demonstrations and conditioned on q_{start} together with either q_{goal} (2D environments and Planar2Link) or the goal end-effector pose $H^{\text{EE}}_{\text{goal}} \in SE(3)$ (Planar4Link and Panda environments). At inference we draw $N = 100$ trajectory samples per problem using DDIM with 15 denoising steps.

Hard constraints. All environments enforce workspace-collision avoidance and joint-position limits. The Panda environments also enforce self-collision avoidance and joint-velocity and joint-acceleration limits. We check every constraint at $H = 128$ phase-uniform points q_h along $q(s)$, with $h = 0, \dots, H - 1$.

One change from MPD affects comparison with its reported numbers. We replace MPD's hinge $\mathrm{ReLU}(-d)$ on the signed clearance d with the CHOMP hinge (Ratliff et al., 2009),

$$\phi_\epsilon(d) = \begin{cases} -d + \frac{1}{2}\epsilon, & d < 0 \\ \frac{1}{2\epsilon}(d - \epsilon)^2, & 0 \le d \le \epsilon \\ 0, & d > \epsilon \end{cases} \tag{369}$$

Unlike ReLU, ϕ_ϵ is C^1 and has a nonzero gradient within ϵ of an obstacle, so Newton-type solvers receive a signal before contact.

Let p_h^a and r^a be the center and radius of body sphere a at q_h. The hard constraints stack into $g(\tau) \leq 0$, with one group per constraint type evaluated at every dense point, sphere, and sphere pair,

$$g_{\text{obj}}(\tau) = \phi_\epsilon\left(\text{sdf}(p_h^a) - r^a - \epsilon\right) \tag{370}$$

$$g_{\text{self}}(\tau) = \phi_\epsilon\left(\|p_h^a - p_h^b\|_2 - r^a - r^b - \epsilon\right) \tag{371}$$

$$g_{\text{jpos}}(\tau) = \left(q_h - q^{\max}, q^{\min} - q_h\right) \tag{372}$$

$$g_{\text{vel}}(\tau) = \pm\dot{q}_h - \dot{q}^{\max} \tag{373}$$

$$g_{\text{acc}}(\tau) = \pm\ddot{q}_h - \ddot{q}^{\max} \tag{374}$$

The self-collision, velocity, and acceleration groups are absent on the lower-DoF environments.

Soft objectives. Every method minimizes the weighted cost $\sum_k w_k C_k(\tau)$, built from the cost terms of MPD (Carvalho et al., 2023), with per-environment weights $w_k \geq 0$ tuned by Bayesian optimization on a held-out validation set. Three smoothness terms penalize squared joint velocity, acceleration, and step length,

$$C_{\text{vel}}(\tau) = \tfrac{1}{2}\sum_h \|\dot{q}_h\|_2^2 \tag{375}$$

$$C_{\text{acc}}(\tau) = \tfrac{1}{2}\sum_h \|\ddot{q}_h\|_2^2 \tag{376}$$

$$C_{\text{path}}(\tau) = \tfrac{1}{2}\sum_h \|q_h - q_{h-1}\|_2^2 \tag{377}$$

On Planar4Link and the Panda environments, a fourth term pulls the final end-effector pose toward the goal pose,

$$C_{\text{ee}}(\tau) = \tfrac{1}{2}\left\|\text{Log}\left(H_{\text{goal}}^{\text{EE}} T(q_{H-1})^{-1}\right)\right\|_2^2 \tag{378}$$

with T the end-effector forward kinematics and Log the SE(3) matrix logarithm. Soft-guidance baselines enforce the constraints through penalties instead. They add the summed entries of g_{obj} and g_{self} to the cost, together with the joint-limit penalty $\sum_h[\phi_\epsilon(q^{\max} - q_h) + \phi_\epsilon(q_h - q^{\min})]$.

Evaluation procedure. In each environment, every method solves the same $M = 100$ problems and generates a batch of $N = 100$ trajectories per problem. A trajectory is *valid* if it satisfies every hard constraint at the $H = 128$ dense points and it ends within twice the robot body width of the goal (Anderson et al., 2018). On Planar4Link and the Panda environments, we measure this distance at the end-effector position. We report the following metrics, each averaged over the M problems.

- *Fraction valid* (FV) is the fraction of the N trajectories that are valid.
- *Fraction of problems solved* (FP) is the fraction of problems whose batch contains at least one valid trajectory.
- *Diversity* (Div) is the Vendi score (Friedman & Dieng, 2023) of the valid trajectories, the exponentiated Shannon entropy of the eigenvalues of the normalized RBF similarity matrix $\tilde{G}_{ij} \propto \exp(-\|\tau_i - \tau_j\|_2^2)$. Div ranges from 1, when all trajectories coincide, to N, when they are pairwise dissimilar. We prefer Div to sample variance, which can be high even when all trajectories cluster around a single mode.
- *Path length* (PathLen) is the joint-space arclength $\int_0^1 \|q'(s)\|_2\, ds$ of the selected trajectory, defined below.
- *End-effector error* (EE-pos, EE-ori) is the position error $\|p^{\text{EE}}(q(1)) - p_{\text{goal}}\|_2$ in meters and the rotation angle between $R^{\text{EE}}(q(1))$ and R_{goal} in degrees, measured on the selected trajectory. We report it on Planar4Link and the Panda environments.

We therefore compute PathLen and EE error on the trajectory a user would deploy. Following Carvalho et al. (2023), we deploy the valid trajectory with the lowest end-effector position error on Planar4Link and the Panda environments, and the shortest valid trajectory elsewhere. If a problem's batch contains no valid trajectory, there is nothing to deploy, so we average PathLen, EE-pos, and EE-ori only over the problems a method solves.

We tune the hyperparameters of every method, including ProxDG, per environment by Bayesian optimization on a held-out validation set. The set of hyperparameters to tune follows MPD (Carvalho et al., 2023).

I.5.3 Solver Details

Decision variable. The solver optimizes the n_w evenly spaced waypoints $\zeta = (\zeta_0, \dots, \zeta_{n_w-1}) \in \mathbb{R}^{n_w \times d}$ of a trajectory. The first waypoint is clamped to q_{start}. The last is clamped to q_{goal} when the goal is given in joint space and left free when it is an end-effector pose. The decision variable $z \in [-1, 1]^{n_z \times d}$ collects the remaining waypoints, each normalized to $[-1, 1]$ by a fixed joint-wise affine map. Costs and constraints are evaluated at the H dense points q_h of the decoded path $\Phi(\zeta(z))$, with velocities and accelerations from finite differences. Since Φ is linear, q, $\dot{q}$, and $\ddot{q}$ are all linear in z.

Problem. Given a proximal anchor $x_{\text{ref}} \in \mathbb{R}^{n_z \times d}$, a proximal weight $\rho_p > 0$, and an annealing parameter $\lambda > 0$, the solver tackles

$$\min_{z \in [-1,1]^{n_z \times d}} \tfrac{\rho_p}{2} \|z - x_{\text{ref}}\|_F^2 + \lambda f(z) \qquad \text{s.t.} \qquad g(z) \le 0 \tag{379}$$

Here $f = \sum_k w_k C_k$ sums the soft costs in Eq. 375–Eq. 378, and g stacks the hard constraints in Eq. 370–Eq. 374. The smoothness costs are quadratic in z and the joint limits are affine, so only C_{ee} and the collision constraints are nonlinear.

Augmented Lagrangian. We solve Eq. 379 with a Powell–Hestenes–Rockafellar (PHR) augmented Lagrangian. Each scalar constraint g_i has a dual $\nu_i \ge 0$. The constraints are split by type into at most five groups, and each group j has its own penalty $\rho_j > 0$. With $\gamma(i)$ the group of constraint i, the augmented Lagrangian is

$$\mathcal{L}(z; \nu, \rho) = \tfrac{\rho_p}{2} \|z - x_{\text{ref}}\|_F^2 + \lambda f(z) + \sum_i \frac{1}{2\rho_{\gamma(i)}} \big([\nu_i + \rho_{\gamma(i)} g_i(z)]_+^2 - \nu_i^2 \big) \tag{380}$$

An inner loop holds (ν, ρ) fixed and takes up to K_{inner} Newton steps on $\mathcal{L}$. An outer loop then applies the multiplier update $\nu_i \leftarrow [\nu_i + \rho_{\gamma(i)} g_i(z)]_+$ and repeats, for at most K_{outer} rounds or until every constraint violation falls below a tolerance ε_{tol}.

Newton steps. The kink in $[\cdot]_+$ makes $\mathcal{L}$ only C^1. To take Newton steps, we replace $[\cdot]_+$ by a softplus inside the gradient and Hessian. We also drop the second derivatives of forward kinematics, the SE(3) logarithm, and the constraints, which leaves a positive semidefinite Gauss–Newton Hessian. The proximal term adds $\rho_p I$, so the Newton system is positive definite and we solve it by Cholesky factorization. Armijo backtracking on the exact $\mathcal{L}$ then selects the step size, so the softplus shapes only the search direction. Finally, we project the iterate back onto the box.

Penalty updates. After each outer round, we multiply a group's penalty ρ_j by a fixed factor, up to a cap, only if the group's worst violation exceeds ε_{tol} and has not shrunk by a factor τ_d since the previous round. Separate penalties keep a stubborn collision violation from inflating the penalty on joint limits that are already satisfied.

Cost per step. The Hessian of the smoothness costs is constant in z, so we compute it once at setup. Each Newton step rebuilds only the Gauss–Newton terms from C_{ee} and the constraints. The solver runs batched over independent problems, each with its own (z, ν, ρ) and stopping test, so the per-step cost grows linearly with the batch size.

I.5.4 RESULTS

Table 18: Robot motion planning on all six *ExtraObjects* environments. For FV, FP, and Div, **best** is bold and second best underlined per environment. PathLen, EE-pos, and EE-ori are unranked because methods that miss problems aggregate over biased solution sets. †Averaged only over problems the method solves.

Env	Method	FV ↑	FP ↑	PathLen ↓	Div ↑	EE-pos ↓	EE-ori ↓
Simple2D	ALM	0.410	0.41	1.606†	1.0	—	—
	DPS	0.174	0.29	1.922†	1.1	—	—
	MPD	0.489	0.99	1.950†	2.7	—	—
	PDM	0.736	0.89	1.920†	1.5	—	—
	CPS	0.834	0.95	1.984†	2.0	—	—
	SafeDiffuser	0.451	0.76	1.996†	2.2	—	—
	DiRecT	0.624	0.63	1.654†	1.0	—	—
	ProxDG	**0.979**	**1.00**	1.826	**3.2**	—	—
NarrowPassageDense2D	ALM	**0.930**	0.99	2.167†	4.0	—	—
	DPS	0.131	0.21	2.015†	1.1	—	—
	MPD	0.198	0.58	2.019†	1.5	—	—
	PDM	0.776	0.95	2.138†	2.1	—	—
	CPS	0.650	0.92	2.188†	1.9	—	—
	SafeDiffuser	0.235	0.69	2.232†	2.9	—	—
	DiRecT	0.915	0.98	2.086†	2.9	—	—
	ProxDG	0.904	**1.00**	2.240	**5.3**	—	—
Planar2Link	ALM	0.879	0.97	5.407†	2.3	—	—
	DPS	0.784	0.94	5.252†	2.0	—	—
	MPD	0.730	0.96	5.390†	5.7	—	—
	PDM	0.884	**1.00**	5.280	**10.0**	—	—
	CPS	0.880	**1.00**	5.339	9.9	—	—
	SafeDiffuser	0.806	0.97	5.122†	2.6	—	—
	DiRecT	0.744	0.83	4.574†	1.4	—	—
	ProxDG	**0.917**	**1.00**	5.205	6.9	—	—
Planar4Link	ALM	0.115	0.80	3.664†	13.5	0.0303†	14.02†
	DPS	0.462	0.85	5.735†	15.2	0.0065†	2.37†
	MPD	0.332	0.69	5.071†	31.2	0.0033†	2.09†
	PDM	0.615	0.97	6.169†	8.7	0.0054†	2.25†
	CPS	0.544	0.93	6.052†	24.0	0.0051†	2.78†
	SafeDiffuser	0.425	0.81	5.920†	15.1	0.0036†	2.28†
	DiRecT	0.690	0.99	7.180†	39.5	0.0035†	1.65†
	ProxDG	**0.813**	**1.00**	7.017	**64.0**	0.0034	3.94
Spheres3D	ALM	0.604	**1.00**	5.129	46.4	0.0086	0.31
	DPS	0.506	0.99	6.939†	48.1	0.0036†	1.46†
	MPD	0.247	0.94	6.369†	27.1	0.0103†	3.49†
	PDM	0.781	**1.00**	6.178	72.8	0.0047	2.51
	CPS	0.833	**1.00**	6.674	**78.8**	0.0059	3.54
	SafeDiffuser	—	—	—	—	—	—
	DiRecT	0.724	0.99	5.767†	68.1	0.0032†	0.33†
	ProxDG	**0.881**	**1.00**	5.562	72.4	0.0080	0.86
Warehouse	ALM	0.791	**1.00**	5.275	52.3	0.0154	1.28
	DPS	0.902	**1.00**	5.694	59.6	0.0066	2.62
	MPD	0.717	**1.00**	6.095	68.5	0.0068	3.53
	PDM	0.970	**1.00**	5.632	70.7	0.0066	3.13
	CPS	0.946	**1.00**	5.499	70.9	0.0092	5.15
	SafeDiffuser	—	—	—	—	—	—
	DiRecT	**0.996**	**1.00**	5.606	75.9	0.0053	1.03
	ProxDG	0.943	**1.00**	5.566	**86.4**	0.0074	4.21

Sampling time. Table 19 reports wall-clock sampling time on Planar4Link and Warehouse. ProxDG is the fastest method on Warehouse and costs about as much as DPS on Planar4Link. Running the ALM solver directly on prior samples is the slowest on both.

Table 19: Wall-clock time to sample one problem's batch of 100 trajectories on a single NVIDIA A100 GPU (80 GB).

Env	DPS	PDM	ALM	ProxDG
Planar4Link (s)	10.56	7.92	15.90	11.34
Warehouse (s)	12.88	8.31	31.47	7.99

I.5.5 Ablations

Scaling solver strength. We allow every guidance method the same number N of objective gradients per denoising step. DPS and MPD spend this budget on N guide gradient steps. ProxDG spends it on the Newton steps of its augmented Lagrangian solver, so $K_{\text{outer}} \cdot K_{\text{inner}} = N$. We sweep $N \in 1, 5, 10, 25$ on three environments. Tables 20, 21, and 22 report FV, the fraction of generated trajectories that are valid. Each entry aggregates 10 test sets with 10 start–goal pairs each, for 100 problems per entry. We tune hyperparameters at $N = 25$ for this ablation instead of reusing the main-results configuration, so its values differ from Table 18.

Table 20: NarrowPassageDense2D: FV vs. inner solver budget N.

Method	N=1	N=5	N=10	N=25
DPS	0.0700	0.0089	0.0010	0.0004
MPD	0.0452	0.2823	0.3194	0.2099
ProxDG	**0.2541**	**0.8367**	**0.8808**	**0.9329**

Table 21: Simple2D: FV vs. inner solver budget N.

Method	N=1	N=5	N=10	N=25
DPS	0.0660	0.0800	0.0212	0.0253
MPD	**0.2851**	0.3615	0.6885	0.7366
ProxDG	0.2321	**0.8922**	**0.8329**	**0.9311**

Table 22: Planar2Link: FV vs. inner solver budget N.

Method	N=1	N=5	N=10	N=25
DPS	0.7760	0.5669	0.1610	0.1691
MPD	0.5873	0.7265	0.8004	0.3086
ProxDG	**0.6116**	**0.9094**	**0.9344**	**0.9197**