PROVABLE DERIVATIVE-FREE INFERENCE WITH SCORE-BASED GENERATIVE PRIORS

Anonymous authors
Paper under double-blind review

ABSTRACT

A growing trend in solving inverse problems is to use pre-trained score-based generative models (SGMs) as plug-and-play priors. This paradigm retains the generative power of SGMs while allowing adaptation to different forward models without requiring re-training. In parallel, derivative-free posterior sampling algorithms have gained increasing attention for solving inverse problems where the derivative, pseudo-inverse, or full knowledge of the forward model is unavailable or impractical to compute. Despite their success, these methods lack principled foundations and provide no convergence guarantees to the true posterior distribution or to its ε-accurate approximation. We propose *zeroth-order annealed plug-and-play Monte Carlo (ZO-APMC)*, the first principled derivative-free framework for solving general inverse problems that requires only forward-model evaluations and a pre-trained SGM prior. We derive complexity bounds for obtaining samples with ε-relative Fisher information under a non-log-concave likelihood distribution and, under a Poincaré inequality assumption, ε-accuracy in total variation distance, and we establish weak convergence of ZO-APMC to the target posterior. We verify our theory with numerical experiments and demonstrate its performance on both linear and nonlinear inverse problems.

1 INTRODUCTION

The use of pre-trained score-based generative models (SGMs) (Song et al., 2020; Ho et al., 2020) as *plug-ang-play* priors for tackling inverse problems has become increasingly prominent, showing strong effectiveness across diverse domains such as image restoration (Wang et al., 2022; Rout et al., 2023), medical imaging (Song et al., 2021; Sun et al., 2024), and image and music generation (Rout et al., 2024). A primary advantage of this framework is its flexibility. It can be applied to various inverse problems without re-training, while preserving the expressive capacity of SGMs to capture complex, high-dimensional priors. In a parallel direction, recently, derivative-free posterior sampling methods with SGM priors (Tang et al., 2024; Huang et al., 2024; Zheng et al., 2024) has attracted growing interest to solve inverse problems where the privileged knowledge of the forward model such as its derivative (Song et al., 2023b; Chung et al., 2022), pseudo-inverse (Song et al., 2023a), or its parametrization (Chung et al., 2023) is unavailable or computationally prohibitive. For example, in many scientific applications (Oliver et al., 2008; Iglesias et al., 2013; Evensen & Van Leeuwen, 1996), forward models are systems of partial differential equations whose derivatives or pseudo-inverse are typically inaccessible or undefined. Despite their empirical success in reconstructing images, they lack theoretical convergence guarantees to the target posterior distribution or to ε-accurate approximation. In fact, even among posterior sampling methods with gradient access of a forward-model, rigorous guarantees are rare; when provided, they typically assume a linear forward operator, which is an assumption often violated in practice (Daras et al., 2024). Appendix **??**, Table 3 provides a conceptual comparison of prior work, highlighting the gap our method fills.

The goal of this work is to develop a theoretically grounded method for solving inverse problems that uses only black-box access to the forward model together with a pre-trained SGM prior. We position this as an important step toward posterior sampling in black-box settings that offers an algorithm with formal convergence guarantees and a solid foundation for future advances. A key challenge in this direction is that, although existing posterior sampling solvers employ principled formulations, they often rely on heuristic, intuition-driven approximations of the forward model's

score function (Iglesias et al., 2013; Huang et al., 2024; Tang et al., 2024), which makes rigorous convergence analysis difficult.

To tackle this issue, we develop a *zeroth-order (ZO) Markov chain Monte Carlo (MCMC)* sampling algorithm (Iglesias et al., 2013; Huang et al., 2024; Tang et al., 2024), in which the forward-model score is approximated from noisy function evaluations. In developing this approach, we face two key challenges. First, Langevin methods are known to exhibit mode collapse and slow convergence when sampling high-dimensional multimodal distributions; motivated by annealed importance sampling (Neal, 2001; Sun et al., 2024), we incorporate a weighted annealing scheduler to enhance exploration. Second, conventional zeroth-order methods are computationally expensive in high dimensions: accurately approximating the forward-model score requires large batch sizes due to the "curse of dimensionality". Additionally, the intricacy of posterior sampling renders simple randomized zeroth-order estimators impractical given their high variance and, thereby, large per-iteration cost. To make our approach practical for high-resolution image reconstruction, we adopt a PAGE-inspired variance-reduction strategy from the optimization literature (Li et al., 2021), which reduces estimator variance and maintains a *fixed per-iteration cost* without sacrificing accuracy. Our main contributions are the following:

- We propose zeroth-order annealed plug-and-play Monte Carlo (ZO-APMC), a completely derivative-free method that utilizes a pre-trained SGM prior in a plug-and-play fashion to tackle general inverse problems. ZO-APMC needs only black-box access to the forward model and works with different forward models without re-training.
- For general non-log-concave likelihood distributions, we establish that the averaged ZO-APMC algorithm exhibits weak convergence to the target distribution under decaying hyperparameters, and we provide non-asymptotic convergence guarantees showing that it attains ε-relative Fisher information after $O(1/\varepsilon^4)$ iterations with fixed per-iteration cost, and, assuming the target distribution satisfies the Poincaré inequality, ε-accuracy after $O(1/\varepsilon^4)$ iterations in total variation distance as well.
- We substantiate our theoretical findings through comprehensive numerical and statistical evaluations, and further demonstrate that our method achieves performance comparable to state-of-the-art gradient-free baselines across diverse inverse problems, including MRI reconstruction, black hole imaging, and Navier-Stokes equation.

2 BACKGROUND

Problem setting. We consider a general inverse problem modeled as

$$\boldsymbol{y} = \boldsymbol{A}(\boldsymbol{x}) + \xi, \quad \boldsymbol{x} \in \mathbb{R}^d, \quad \boldsymbol{y}, \xi \in \mathbb{R}^m, \tag{1}$$

where the objective is to recover the unknown signal $\boldsymbol{x}$ from noisy measurements $\boldsymbol{y}$. The forward operator $\boldsymbol{A} : \mathbb{R}^d \to \mathbb{R}^m$ characterizes the response of the imaging system, while $\xi \in \mathbb{R}^m$ denotes the measurement noise, typically modeled as Gaussian or Laplacian distribution. Recovering $\boldsymbol{x}$ from $\boldsymbol{y}$ amounts to inverting (1), which can be viewed as estimating either the most probable $\boldsymbol{x}$ or its full posterior distribution $\pi(\boldsymbol{x}|\boldsymbol{y})$ from Bayesian perspective. This inference problem is commonly expressed as $\pi(\boldsymbol{x}|\boldsymbol{y}) \propto \ell(\boldsymbol{y}|\boldsymbol{x})p(\boldsymbol{x})$, where $p(\boldsymbol{x})$ denotes the prior distribution over source signal, implemented here via a pre-trained SGM, and $\ell(\boldsymbol{y}|\boldsymbol{x})$ is the likelihood distribution defined by (1). In this work, we address the inverse problem using Bayesian inference based on MAP estimation, $\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{x}} \ell(\boldsymbol{y}|\boldsymbol{x})p(\boldsymbol{x})$. However, because we have only black-box access to $\boldsymbol{A}(\cdot)$, we can generate samples from $\ell(\boldsymbol{y}|\boldsymbol{x})$ but lack its explicit functional form.

Score-based generative models (SGMs). SGMs have emerged as a powerful deep learning (DL) framework for sampling from complex, high-dimensional distributions. At their core, they learn the perturbed score function $\nabla \log p_\sigma(\boldsymbol{x})$, where $p_\sigma(\boldsymbol{x}) = \int_{\mathbb{R}^d} p(\boldsymbol{z})\phi_\sigma(\boldsymbol{z} - \boldsymbol{x})\, d\boldsymbol{z}$ and ϕ_σ is the probability density function of $\mathcal{N}(0, \sigma^2 I)$. This score is learned using the *score matching* technique (Hyvärinen & Dayan, 2005; Vincent, 2011) and estimated via *Tweedie's* formula (Efron, 2011). The resulting score estimates are then integrated into MCMC samplers to perform iterative draws for unconditional image generation (Song & Ermon, 2020; 2019). In particular, sampling proceeds via a discretization of the Langevin diffusion process

$$d\mathbf{x}_t = \nabla \log p(\mathbf{x}_t)\, dt + \sqrt{2}\, d\boldsymbol{B}_t, \tag{2}$$

where $\{\boldsymbol{B}_t\}_{t\geq 0}$ denotes an n-dimensional Brownian motion and the learned score function $\mathcal{S}_\theta(\boldsymbol{x}_t, \sigma)$ approximates $\nabla \log p(\mathbf{x}_t)$ for sufficiently small $0 < \sigma \ll 1$. For posterior sampling, applying Bayes' rule and substituting the prior term with its estimate provided by $\mathcal{S}_\theta(\boldsymbol{x}, \sigma)$ yields

$$d\boldsymbol{x}_t = [\nabla \log \ell(\boldsymbol{y}|\boldsymbol{x}_t) - \mathcal{S}_\theta(\boldsymbol{x}_t, \sigma)]\, dt + \sqrt{2} d\boldsymbol{B}_t \tag{3}$$

as used in (Sun et al., 2024). Similarly, diffusion models (DMs), a class of SGMs, can be used to draw samples from the posterior $\pi(\boldsymbol{x}|\boldsymbol{y})$ by reversing a diffusion process from $\pi(\boldsymbol{x}|\boldsymbol{y})$ to a simple distribution (Yang et al., 2023) and approximating the time-dependent score $\nabla \log \ell_t(\boldsymbol{y}|\boldsymbol{x}_t)$ (Chung et al., 2022). One limitation of these approaches is their reliance on full acess to gradient $\nabla \log \ell(\boldsymbol{y}|\boldsymbol{x})$, which is not possible for many scientific inverse problems (Knape & De Valpine, 2012; Zheng et al., 2025). Another limitation is that, despite their empirical success, these methods rely on heuristic approximations of $\nabla \log \ell_t(\boldsymbol{y}|\boldsymbol{x}_t)$, which hinders rigorous theoretical analysis.

Derivative-free diffusion guidance for inverse problems. Recent studies increasingly explore derivative-free strategies for guiding SGMs in inverse problems. Three DM-based approaches have been proposed to date: Stochastic Control Guidance (SCG) (Huang et al., 2024) and Diffusion Policy Gradient (DPG) (Tang et al., 2024), which cast diffusion guidance in a stochastic control framework and steer the sampling process via an estimated value function, and Ensemble Kalman Guidance (EnKG) (Zheng et al., 2024), which uses statistical linearization to guide the diffusion process without explicit gradients. Although these methods have shown encouraging empirical results, they face a fundamental trade-off between broad applicability to highly nonlinear, black-box systems and the availability of rigorous convergence guarantees. In fact, even among gradient-based posterior sampling algorithms with SGM prior, only a few offer formal convergence results (Sun et al., 2024). This tension motivates our proposed approach, which seeks to combine the practical scope of derivative-free guidance with strong theoretical foundations.

Zeroth-order sampling. A zeroth-order (ZO) gradient estimator of a function f can be obtained using a forward finite difference along a random direction (Nesterov & Spokoiny, 2017):

$$\widetilde{\nabla} f(\boldsymbol{x}) = \frac{f(\boldsymbol{x} + \mu \boldsymbol{u}) - f(\boldsymbol{x})}{\mu}\, \boldsymbol{u}, \quad \boldsymbol{u} \sim \mathcal{N}(0, I), \tag{4}$$

where $\mu > 0$ is a small smoothing parameter. In our formulation, f represents the negative log-likelihood (or potential) function, with $\ell(\boldsymbol{y}|\boldsymbol{x}) \propto e^{-f(\boldsymbol{x})}$. By discretizing (3) and replacing the negative log-likelihood with its ZO estimator from (4), we obtain a zeroth-order Langevin sampling algorithm with SGM prior. (Roy et al., 2022) establish convergence guarantees for generating ε-approximate samples in Wasserstein distance under convex and smooth f; however, their analysis is purely theoretical, considers only settings without a prior, and assumes log-concave forward model distribution, an assumption typically violated in inverse problems. (He et al., 2024) establish asymptotic KL convergence but neither demonstrate the method on real-world problems nor consider posterior sampling.

More recently, (Sun et al., 2024) proposed annealed plug-and-play Monte Carlo (APMC), the closest work to ours in the literature, and derived an upper bound on the Fisher information, albeit under the assumption of access to the forward model's score function. In contrast, we prove convergence to stationary point in ε-relative Fisher information, to the total variation distance (assuming that the potential function of the forward model satisfies Poincaré inequality), and weak convergence to the posterior distribution $\pi(\boldsymbol{x}|\boldsymbol{y})$. Further discussion of the related work is available in Appendix **??**.

3 METHOD

To develop our ZO-APMC method, we first provide an interpretation of annealed Langevin dynamics and intuition behind the variance-reduction mechanism for zeroth-order estimate. Then, we present our algorithm with its convergence guarantees.

Annealed Langevin dynamics. As discussed in Section 2, given a SGM prior $\mathcal{S}_\theta(\boldsymbol{x}, \sigma) \approx \nabla \log p(\boldsymbol{x})$, we can discretize the Langevin diffusion in (3) and get the update rule as

$$\boldsymbol{x}_{k+1} \coloneqq \boldsymbol{x}_k + \gamma \left(\tilde{\nabla} f(\boldsymbol{x}_k) - \mathcal{S}_\theta(\boldsymbol{x}_k, \sigma) \right) + \sqrt{2\gamma} \boldsymbol{Z}_k, \tag{5}$$

where $\boldsymbol{Z}_k \sim \mathcal{N}(0, I)$ and $f(\boldsymbol{x}_k) \coloneqq -\log \ell(\boldsymbol{y}|\boldsymbol{x}_k)$ is the negative log-likelihood. Recall also that since we assume black-box setting, we replace $\nabla f(\boldsymbol{x})$ with its ZO estimate $\tilde{\nabla} f(\boldsymbol{x}) =$

Algorithm 1 ZO-APMC

Input: initial point $\boldsymbol{x}_0$, stepsize γ, minibatch size b, $b' < b$, probability $p \in (0, 1]$, and annealing parameters $\alpha_0 > 0$, $\sigma_0 > 0$.

1: $g_0 = \frac{1}{b}\sum_{i\in I} \tilde{\nabla} f_\mu(\boldsymbol{x}_0, \boldsymbol{u}_i)$ // I denotes random minibatch samples with $|I| = b$.
2: **for** $k = 0, 1, \ldots, N-1$ **do**
3: $\quad \boldsymbol{Z}_k \leftarrow \mathcal{N}(0, I)$
4: $\quad \sigma_k, \alpha_k \leftarrow \text{WeightedAnnealing}(\sigma_0, \alpha_0, k)$
5: $\quad \mathcal{G}_k(\boldsymbol{x}_k) \leftarrow \boldsymbol{g}_k - \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_k, \sigma_k)$
6: $\quad \boldsymbol{x}_{k+1} \leftarrow \boldsymbol{x}_k - \gamma \mathcal{G}_k(\boldsymbol{x}_k) + \sqrt{2\gamma}\boldsymbol{Z}_k$
7: $\quad \boldsymbol{g}_{k+1} = \begin{cases} \frac{1}{b}\sum_{i\in I} \tilde{\nabla} f_\mu(\boldsymbol{x}_{k+1}, \boldsymbol{u}_i), & \text{with prob. } p \\ \boldsymbol{g}_k + \frac{1}{b'}\sum_{i\in I'} \tilde{\nabla} f_\mu(\boldsymbol{x}_{k+1}, \boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_k, \boldsymbol{u}_i), & \text{with prob. } 1-p \end{cases}$
8: **end for**

Output: $\boldsymbol{x}_N$

$(1/b\mu)\sum_{i=0}^{N-1}(f(\boldsymbol{x} + \mu\boldsymbol{u}_i) - f(\boldsymbol{x}))\boldsymbol{u}_i$ where b is the batch size. In practice, Langeving algorithms often experience slow convergence and mode collapse when sampling from high-dimensional, multimodal distributions. Inspired by annealed importance sampling (Neal, 2001; Sun et al., 2024), we consider a sequence of posterior distributions for each step

$$\pi_{\sigma_k}^{(\alpha_k)}(\boldsymbol{x}|\boldsymbol{y}) \propto \ell(\boldsymbol{y}|\boldsymbol{x}) p_{\sigma_k}^{\alpha_k}(\boldsymbol{x}), \quad p_{\sigma_k}(\boldsymbol{x}) = \int_{\mathbb{R}^d} p(\boldsymbol{z})\phi_{\sigma_k}(\boldsymbol{z} - \boldsymbol{x})\, d\boldsymbol{z}, \tag{6}$$

where $\alpha_0 > \alpha_1 > \ldots > \alpha_K = \cdots = \alpha_{N-1} = 1$, $\sigma_0 > \sigma_1 > \ldots > \sigma_K = \cdots = \sigma_{N-1} \approx 0$, and ϕ_{σ_k} is the probability density function of $\mathcal{N}(0, \sigma_k^2 I)$. $\{\alpha_k\}_k^{N-1}$ and $\{\sigma_k\}_{k=0}^{N-1}$ are generally initialized with large values in practice and they decay to one and almost zero, respectively. Initially, the weighted posterior, is dominated by a smoothed prior, enabling rapid escape from gradient plateaus where $\nabla \log \pi(\boldsymbol{x}) \approx 0$. As iterations proceed, the likelihood influence grows and the smoothed posterior $\pi_{\sigma_k}^{(\alpha_k)}$ sharpens toward the true posterior distribution π. This annealing accelerates burn-in by first flattening and then gradually restoring distributional complexity. This process is illustrated with Figure 5 in Appendix A inspired by (Sun et al., 2024). With the annealing parameters, we can write the new update rule as

$$\boldsymbol{x}_{k+1} \coloneqq \boldsymbol{x}_k + \gamma\left(\tilde{\nabla} f(\boldsymbol{x}_k) - \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_k, \sigma_k)\right) + \sqrt{2\gamma}\boldsymbol{Z}_k. \tag{7}$$

3.1 Proposed Method: ZO-APMC

We now introduce our ZO-APMC algorithm and give an intuitive explanation of its variance-reduction mechanism. ZO gradient estimates (finite-difference, random direction, or coordinate sampling) are well-known for their high variance, especially in high-dimensional settings (Nesterov & Spokoiny, 2017). While despite this high variance, convergence guarantees are typically attainable in optimization problems (Nesterov & Spokoiny, 2017; Lan, 2020), our analysis unveils that the high variance prevents achieving strong convergence guarantees in the posterior sampling setting without requiring *growing batch sizes*, thereby rendering the resultant methods prohibitive in terms of their memory requirements. To resolve this critical issue, we replace the ZO estimate in (6) with a variance-reduced estimate g_k given in line 7 of ZO-APMC inspired by (Li et al., 2021) where $|I| = b$, $|I'| = b'$ denote the large and small batch sizes, respectively, and $\boldsymbol{u}_i \sim \mathcal{N}(0, I)$. In practice, we set $b' < b$, which significantly lowers the average number of function evaluations per iteration thanks to the reduced variance. With the variance-reduction mechanism inplace, we present the pseudocode for ZO-APMC in Algorithm 1. Moreover, while we provide the convergence results in Section 3.3, we now present an upper bound on the estimation error illustrating the variance-reduction mechanism for ZO posterior sampling.

Assumption 1 *We assume that the log-likelihood* $\log \ell(\boldsymbol{x}|\boldsymbol{y})$ *is Lipschitz continuous with constant* L_{f_2}*, namely, for any* $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n$, $\|\log \ell(\boldsymbol{y}|\boldsymbol{x}_1) - \log \ell(\boldsymbol{y}|\boldsymbol{x}_2)\| \leq L_{f_2}\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|$.

Remark 1. Assumption 1 is not satisfied under Gaussian noise, where the log-likelihood involves the squared residual $\|A(\boldsymbol{x}) - \boldsymbol{y}\|_2^2$ and is not Lipschitz. By contrast, with Laplacian noise the log-likelihood becomes $\log \ell(\boldsymbol{y}|\boldsymbol{x}) \propto \|A(\boldsymbol{x}) - \boldsymbol{y}\|_1$, whose absolute-value residual can admit a global

Lipschitz bound. In addition, Laplacian noise modeling is widely used in sparse modeling and imaging to robustly handle heavy-tailed errors and outliers (Chan & Esedoglu, 2005; Boyd & Vandenberghe, 2004; Rousseeuw & Hubert, 2011).

Proposition 1 *Under Assumption 1, let $\{\boldsymbol{x}_k\}_{k=0}^N$ denote the iterates produced by ZO-APMC for $N > 0$ steps. Define the estimation error of the forward-model score as $\boldsymbol{e}_k := \boldsymbol{g}_k - \nabla f_\mu(\boldsymbol{x}_k)$. Then, for each step k, the error variance satisfies*

$$\mathbb{E}\big[\|\boldsymbol{e}_k\|^2\big] \;\le\; (1-p^2)\,\mathbb{E}\big[\|\boldsymbol{e}_{k-1}\|^2\big] + \frac{4d(1-p)L_{f_2}^2}{b'\mu^2}\,\mathbb{E}\big[\|\boldsymbol{x}_k - \boldsymbol{x}_{k-1}\|^2\big] + \frac{p\,\mathbb{E}\big[\|\widetilde{\nabla} f_\mu(\boldsymbol{x}_k, \boldsymbol{u}_i)\|^2\big]}{b}.$$

The detailed derivation using the law of total covariance can be found in Appendix A.1.2. Proposition 1 provides an intuition that, thanks to the variance-reduction mechanism, larger batch sizes and smaller step sizes tend to make the bound dominated by the previous-step error variance with a contraction factor, which in turn yields progressively reduced estimator variance across iterations. For clarity, note that our ZO gradient estimator is biased ($\nabla f_\mu(\boldsymbol{x}) \neq \nabla f(\boldsymbol{x})$), but the bias vanishes as $\mu \to 0$. However, this raises the upper bound on the error variance and violates the contraction property. With an appropriate choice of the step size γ, this can be solved. We make this trade-off explicit in the main results, which follow after the discussion of our convergence criteria.

3.2 Optimization View of Langevin Diffusion

Consider the minimization of the *Kullback–Leibler (KL)* divergence over the Wasserstein space of probability distributions.

$$\hat{\nu} = \arg\min_{\nu} \mathrm{KL}(\nu\|\pi) \quad \text{where} \quad \mathrm{KL}(\nu\|\pi) = \int_{\mathbb{R}^d} \log \frac{\nu(\boldsymbol{x})}{\pi(\boldsymbol{x})} d\boldsymbol{x}, \tag{8}$$

where ν and π denote the estimate and desired posterior, respectively. Similar to the gradient concept in Euclidean space, we can write the Wasserstein gradient of $\mathrm{KL}(\nu\|\pi)$ as $\nabla_\nu \mathrm{KL}(\nu\|\pi) = \nabla \log(\nu(\boldsymbol{x})/\pi(\boldsymbol{x}))$ (Ambrosio et al., 2008) and its expected square norm gives us the *relative Fisher information (FI)* $\mathrm{FI}(\nu\|\pi) = \int_{\mathbb{R}^n} \|\nabla \log \nu(x) - \nabla \log \pi(x)\|_2^2 \nu(x)\, dx$. If ν_t evolves under Langevin diffusion in (3), then $\frac{d}{dt}\mathrm{KL}(\nu_t\|\pi) = -\mathrm{FI}(\nu_t\|\pi)$ (Ambrosio et al., 2008; Villani, 2009), showing that Langevin diffusion is a gradient flow in probability space. From an optimization viewpoint, $\mathrm{FI}(\nu_t\|\pi)$ serves as the analogue of the squared ℓ_2 gradient norm in $\mathbb{R}^d$ (Balasubramanian et al., 2022). Leveraging this analogy, we analyze the convergence of $\mathrm{FI}(\nu_t\|\pi)$ under a "linear interpolation" of the distributions generated by ZO-APMC, which in turn implies the stationarity of the discrete updates.

3.3 General Convergence Results

In this section, we state our main theoretical results establishing the convergence of ZO-APMC. We first state our assumptions.

Assumption 2 *The log-prior $\log p(\boldsymbol{x})$ is differentiable and $\nabla \log p(\boldsymbol{x})$ is L_{f_1}-Lipschitz, i.e. $\|\nabla \log p(\boldsymbol{x}_1) - \nabla \log p(\boldsymbol{x}_2)\| \le L_{f_1}\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|$ for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n$.*

Assumption 3 *Let $p_{\sigma_k}(\boldsymbol{x}) = \int_{\mathbb{R}^d} p(\boldsymbol{z})\phi_{\sigma_k}(\boldsymbol{x} - \boldsymbol{z})d\boldsymbol{z}$ denote the smoothed prior, where ϕ_{σ_k} is the probability density function of $\mathcal{N}(0, \sigma_k^2 I)$. We assume that for any $\sigma_k > 0$ and $\boldsymbol{x} \in \mathbb{R}^d$, $\|\nabla \log p_{\sigma_k}(\boldsymbol{x}) - \nabla \log p(\boldsymbol{x})\| \le \sigma_k C$.*

Assumption 4 *We assume that the log-likelihood $\log \ell(\boldsymbol{x}|\boldsymbol{y})$ is differentiable and has a Lipschitz continuous gradient with constant $L_{f_1} > 0$ for any $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^n$, that is, $\|\nabla \log \ell(\boldsymbol{y}|\boldsymbol{x}_1) - \nabla \log \ell(\boldsymbol{y}|\boldsymbol{x}_2)\| \le L_{f_1}\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|$.*

Assumption 5 *For any $\sigma_k > 0$ and all $\boldsymbol{x} \in \mathbb{R}^d$, the score network satisfies $\|\mathcal{S}_\theta(\boldsymbol{x}, \sigma_k) - \nabla \log p_{\sigma_k}(\boldsymbol{x})\| \le \varepsilon_{\sigma_k} < \infty$ and $\|\mathcal{S}_\theta(\boldsymbol{x}, \sigma_k)\| \le R_s$.*

Assumptions 2 and 4 correspond to standard conditions commonly adopted in the non-log-concave sampling literature (He & Zhang, 2025; Guo et al., 2024; Balasubramanian et al., 2022), and Assumption 3 captures the perturbation of the prior as in (Sun et al., 2024). Assumption 5 imposes

boundedness on both the score-network error and its output. The bounded error is a standard and essential requirement for the theoretical analysis of sampling with SGM priors (Sun et al., 2024; Lee et al., 2023; Wu et al., 2024; Lee et al., 2022; Chen et al., 2023), and the bounded output can be easily implemented in practice via simple clipping. Notably, unlike related works (Sun et al., 2024; Yang & Wibisono, 2022; Lee et al., 2022), we do not assume Lipschitz continuity of the SGM prior. Moreover, none of our assumptions restricts the likelihood distribution to be log-concave. Under these standard assumptions, we prove the convergence of ZO-APMC.

Theorem 1 *Let $\{\alpha_k\}_{k=0}^{N-1}$, $\{\sigma_k\}_{k=0}^{N-1}$ be decreasing annealing schedules with $\alpha_{K,\dots,N-1} = 1$, and let $\{\nu_t\}_{t\geq 0}$ denote the law of the continuous interpolation of $\{\boldsymbol{x}_k\}_{k=0}^{N}$ produced by ZO-APMC with $N > 0$ iterations under Assumptions 1–5. For any step size $\gamma \in \big(0, 1/(L_m\sqrt{85\phi(\mu)})\big]$, where $\phi(\mu) = 1 + 4(1-p)d/p\mu^2 b'$, the Fisher information satisfies*

$$\frac{1}{N\gamma}\int_0^{N\gamma} \mathrm{FI}(\nu_t\|\pi)\,dt \leq \frac{C_0}{N\gamma} + 8\gamma L_m^2 d\phi(\mu) + \frac{17d(d+2)L_{f_2}^2}{2b} + \frac{17\mu^2 L_{f_1}^2(d+3)^3}{8} + \bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2, \tag{9}$$

where $L_m = \max\{L_\pi, L_{f_2}\}$, $C_0 = 2\mathrm{KL}(\nu_0\|\pi) + 10\gamma\mathbb{E}\|g_0 - \nabla f_\mu(\boldsymbol{x}_0)\|^2/pL_m\sqrt{85\phi(\mu)}$, $\bar{\sigma}^2 = \frac{51C}{2N}\sum_{k=0}^{N-1}\sigma_k^2$, $\bar{\varepsilon}_\sigma^2 = \frac{51}{2N}\sum_{k=0}^{N-1}\varepsilon_{\sigma_k}^2$, and $\bar{\alpha}_{-1}^2 = \frac{51}{2N}\sum_{k=0}^{N-1}(\alpha_k - 1)^2 R_s^2$. Furthermore, suppose the parameters and schedules are chosen as $\gamma = \sqrt{C_0}/(2L_m\sqrt{Nd\phi(\mu)})$, $b = \lceil 51d(d+2)L_{f_2}^2/\varepsilon\rceil$, $p = 1/b$, $\mu = 2\sqrt{\varepsilon}/\sqrt{51}L_{f_1}(d+3)^{3/2}$ with annealing schedules $\sigma_n = O(n^{-\beta})$, $\alpha_n = 1 + O(n^{-\beta})$ for $\beta > 1/2$, and score-network error satisfying $\varepsilon_{\sigma_n} = O(n^{-\beta})$ for $\beta > 1/2$. Then an ε-approximate solution to (9) can be obtained with $N = O\big(d^7 L_m^6/\varepsilon^4\big)$ forward-model evaluations, using a fixed evaluation budget $pb = O(1)$ per iteration on average.

Theorem 1 (proof provided in Appendix A.1.3) shows ZO-APMC achieves ε-approximate solution in the Fisher information sense with $N = O\big(d^7 L_m^6/\varepsilon^4\big)$ forward-model evaluations and using a fixed evaluation budget $pb = O(1)$ per iteration.

Remark 2. In practice, the annealing schedules $\{\alpha_n\}_{k=0}^{N-1}$ and $\{\sigma_n\}_{k=0}^{N-1}$ are typically implemented using geometric decay (Sun et al., 2024; Song & Ermon, 2019), which decreases more rapidly than the polynomial rates selected for our analysis. Moreover, the condition $\varepsilon_{\sigma_n} = O(n^{-\beta})$ with $\beta > \frac{1}{2}$ characterizes the decay of the SGM generalization complexity across each noise level at step n. Recent studies on the generalization of SGMs report similar rates as the one used in our analysis (Fu & Lee, 2025; Zhang et al., 2024; Oko et al., 2023).

Leveraging the results of Theorem 1, we show that if the target posterior π further satisfies the Poincaré inequality, ZO-AMPC enjoys stronger sampling guarantees in total variation distance.

Assumption 6 *For every smooth, compactly supported function $f : \mathbb{R}^d \to \mathbb{R}$, the posterior distribution $\pi(\boldsymbol{x}|\boldsymbol{y})$ satisfies the Poincaré inequality $\mathrm{Var}_\pi(f) \leq C_{\mathrm{PI}}\mathrm{FI}(\nu\|\pi)$.*

Corollary 1 *Let $\{\nu_t\}_{t\geq 0}$ denote the law of the continuous interpolation $\{\boldsymbol{x}_k\}_{k=0}^{N}$ of ZO-APMC, and let the Assumptions Assumptions 1–6 hold. Then, if we choose $\gamma = \sqrt{C_0 C_{\mathrm{PI}}}/2L_m\sqrt{Nd\phi(\mu)}$, we have*

$$\|\bar{\nu}_{N\gamma} - \pi\|_{\mathrm{TV}} \leq 16L_m\sqrt{\frac{C_0 C_{\mathrm{PI}} d\phi(\mu)}{N}} + \frac{34d(d+2)C_{\mathrm{PI}}L_{f_2}^2}{b} + \frac{17}{2}\mu^2 C_{\mathrm{PI}}L_{f_1}^2(d+3)^3 + 4C_{\mathrm{PI}}(\bar{\sigma}^2 + \varepsilon_\sigma^2 + \bar{\alpha}_{-1}^2) \tag{10}$$

where $\bar{\nu}_{N\gamma} := (N\gamma)^{-1}\int_0^{N\gamma}\nu_t dt$. If we choose $b = \lceil 204d(d+2)C_{\mathrm{PI}}L_{f_2}^2/\varepsilon\rceil$, $p = 1/b$, $\mu = \sqrt{\varepsilon}/L_{f_1}(d+3)^{3/2}\sqrt{51C_{\mathrm{PI}}}$ with annealing schedules $\sigma_n = O(n^{-\beta})$, $\alpha_n = O(n^{-\beta})$ for $\beta > 1/2$ and assuming $\varepsilon_{\sigma_n} = O(n^{-\beta})$ for $\beta > 1/2$, an ε-approximate solution to (10) requires $N = O(d^7 L_m^6 C_{\mathrm{PI}}^3/\varepsilon^4)$ forward model evaluations, using a fixed ($pb = O(1)$) evaluation per iteration.

Theorem 1 (proof provided in Appendix A.1.4) shows ZO-APMC achieves ε-approximate solution in the stronger total variation sense with $N = O\big(d^7 L_m^6 C_{\mathrm{PI}}^3/\varepsilon^4\big)$ forward-model evaluations and using a fixed evaluation budget $pb = O(1)$ per iteration.

Remark 3. To generate a sample from $\bar{\nu}_{N\gamma}$, one may proceed as follows. First, draw a time $t \in [0, N\gamma]$ uniformly at random and determine the largest integer k such that $k\gamma \leq t$. Then perform a linear interpolation between the interval $[k\gamma, t]$ to produce $\boldsymbol{x}_t$ according to the update rule in line 7 of ZO-APMC. The resulting $\boldsymbol{x}_t$ is sample from $\bar{\nu}_{N\gamma}$.

Lastly, we establish the asymptotic convergence of the averaged ZO-APMC algorithm, with appropriately decaying hyperparameters, to the target posterior distribution under mild conditions.

Theorem 2 *Let $\{\nu_t\}_{t\geq 0}$ denote the law of the continuous interpolation $\{\boldsymbol{x}_k\}_{k=0}^N$ produced by ZO-APMC, and suppose Assumptions 1–5 hold for π; if ZO-APMC is initialized at a measure ν_0 with $\mathrm{KL}(\nu_0\|\pi) < \infty$ and uses $\gamma_n = \sqrt{b'}/(nL_m\sqrt{680d})$, $b_n = \lceil n^{1/2} \rceil$, $p_n = n^{-1/2}$, $\mu_n = n^{-1/8}$, and $\sigma_n = O(n^{-\beta})$ and if $\bar{\varepsilon}_\sigma^2 = O(n^{-\beta})$ for $\beta > 0$, then $\bar{\nu}_{\tau_n} \to \pi$ weakly, where $\tau_n = \sum_{k=1}^n \gamma_k$ and $\bar{\nu}_{\tau_n} = \tau_n^{-1} \int_0^{\tau_n} \nu_t \, dt$.*

We present the complete proof in Appendix A.1.5. To the best of our knowledge, this is the first work to establish weak convergence of ZO Langevin MC and ZO posterior sampling algorithms for non-log-concave distributions. This result follows directly from Theorem 1 together with the fact that $\mathrm{FI}(\mu\|\pi) = 0$ implies $\mu = \pi$. We also emphasize the key role of our proposed estimator (7) in establishing this result with *a fixed evaluation budget $pb = O(1)$ per iteration*.

4 Experiments

Baselines. Our primary focus is on gradient-free methods, which assume only black-box access to the forward model. We therefore benchmark against three gradient-free baselines: SCG (Huang et al., 2024), DPG (Tang et al., 2024), and EnKG (Zheng et al., 2024). We also include the Forward-GSG and Central-GSG baselines, introduced by (Zheng et al., 2024). These methods resemble Diffusion Posterior Sampling (DPS) (Chung et al., 2022) but approximate the forward-model gradient using Tweedie's formula together with forward and central ZO estimates of the forward score function. For completeness, we also evaluate gradient-based methods in settings where the forward-model gradient is available. Specifically, we compare our algorithm with DPS (Chung et al., 2022), PnPDM (Wu et al., 2024), and APMC (Sun et al., 2024), which is an annealed Langevin MC posterior sampling algorithm with gradient access and the closest approach to ours.

4.1 Toy Experiments

Numerical Validation. We test our theory that ZO-APMC converges in FI with fixed per-iteration cost on a synthetic bimodal 2D Gaussian-mixture prior with random $\boldsymbol{A}$ with $\xi \sim \mathcal{N}(0, I)$. Using the analytical score with added Gaussian noise $\varepsilon_{k^*} = 2.5$ to mimic SGM error, we generate 1000 samples with ZO-APMC from 20 random initializations and report the mean FI relative to the analytical posterior. Fig. 1a shows that with $b = 10$, $b' = 5$, ZO-APMC converges near zero for $p \in \{1, 0.75, 0.5\}$, matching gradient-based APMC but becoming unstable at $p = 0.3$ due to fixed b. Fig. 1b shows that increasing b while keeping $pb = 10$ restores stability and achieves convergence ($FI \leq 0.01$), confirming our theoretical results.

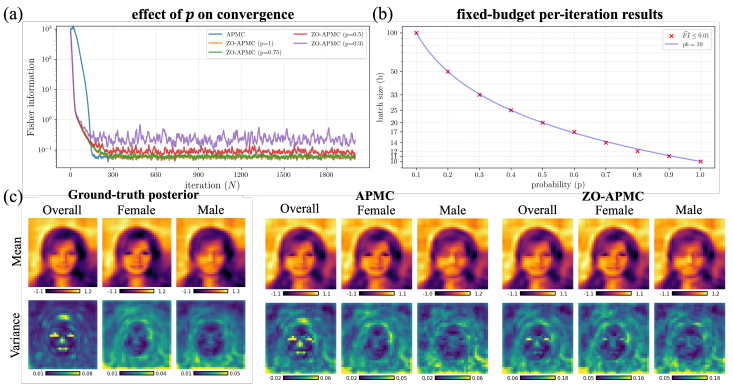

Figure 1: (a) Convergence of ZO-APMC with $b = 10$, $b' = 5$, and $\varepsilon_{k^*} = 2.5$ for various p, alongside APMC convergence with gradient access. (b) Convergence results for fixed-budget per-iteration. Each red "x" marks cases where the FI falls below 0.01 after 2000 iterations for the corresponding p and b. (c) Comparison of sample statistics obtained by ZO-APMC and APMC versus the ground-truth posterior.

Table 1: Quantitative comparison with baselines. The best values of each metric for black-box and gradient-access settings are highlighted in **bold** and underline, respectively.

	PSNR (dB)↑	SSIM↑	NRMSE↓	SD↓	MSE↓
PnPDM	30.81	0.946	3.76e-2	2.16e-2	8.46e-4
DPS	34.38	0.965	2.54e-2	2.06e-2	4.07e-4
APMC	36.55	0.973	1.99e-2	2.0e-2	2.55e-4
Forward-GSG	27.8	0.918	5.42e-2	3.26e-2	19.1e-4
Central-GSG	27.78	0.917	5.43e-2	3.27e-2	19.2e-4
SCG	7.1	0.711	7.67	1.38	0.21
DPG	32.17	0.953	5.4e-2	**2.69e-2**	6.5e-4
EnKG	31.32	0.934	5.72e-2	2.92e-2	6.72e-4
ZO-APMC (ours)	**35.29**	**0.966**	**2.28e-2**	2.99e-2	**3.29e-4**

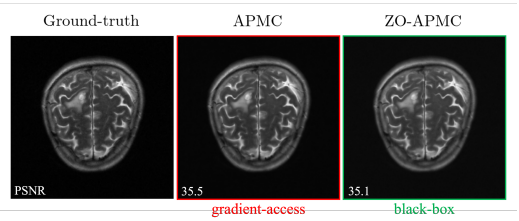

Figure 2: Visual comparison of pathological brain MRI with corresponding PSNR values obtained using APMC with gradient access and ZO-APMC in the black-box setting.

Statistical Validation. We assess ZO-APMC's ability to sample from multiple modes of a multimodal distribution under black-box conditions. Using the same setup as our previous validation with random $\boldsymbol{A} \in \mathbb{R}^{115\times1024}$, we construct a two-mode Gaussian mixture prior from CelebA (Liu et al., 2018) images normalized to $[0, 1]$, where the "male" and "female" attributes define the modes. To ensure clear separation between the modes, we shift them by $+1$ and -1, respectively. A shallow SGM is then trained by customizing the U-Net from (Nichol & Dhariwal, 2021) for this data. As shown in Fig. 1c, ZO-APMC with $p=0.5$, $b=50$, $b'=5$, accurately recovers the posterior statistics of both modes, comparable to APMC with gradient access, though with slightly higher variance due to ZO estimation, which can be mitigated by increasing b. For extended results and further details of validations, see Appendix A.2.

4.2 Magnetic Resonance Imaging (MRI)

Image inverse problems (i.e., MRI recon.) are widely used benchmarks. Although we focus on more challenging black-box forward models, we also evaluate our method on the linear MRI recon. problem for completeness and demonstrate the capability of our variance-reduction mechanism on high-resolution data.

Problem Setting We consider the radial subsampling mask with acceleration factor of $4\times$. For evaluation, we use the SGM prior from Sun et al. (2024), which was pre-trained on the FastMRI brain dataset (Zbontar et al., 1811), and evaluate all algorithms on a separate test set provided in that work to ensure a consistent comparison. We randomly select 40 images at a resolution of 256×256 pixels and generate 20 reconstructions per algorithm. For each method, we report the mean image-quality metrics along with the average per-pixel standard deviation (SD). In this experiment, we use $p = 0.2$, $b = 10^4$, and $b' = 10^3$ to run ZO-APMC.

Results Table 1 shows that ZO-APMC consistently achieve higher reconstruction quality than other black-box baselines in all image quality metrics and closely matches the APMC with gradient access. Fig. 2 further demonstrates that both ZO-APMC and APMC yield visually indistinguishable pathological brain MRI reconstructions, with ZO-APMC accurately capturing fine details without gradient information. Our method yields slightly higher standard deviation than DPG but this can be alleviated by increasing p, albeit at increased computational cost.

4.3 Black-hole Imaging

Problem Setting The black-hole interferometric imaging system reconstructs images of black holes from "visibility" measurements collected by Earth-based telescope arrays. We adopt the SGM prior (pre-trained on the GRMHD dataset at 64×64 resolution), the highly non-linear forward model, and the 100-sample test set, as provided by the InverseBench benchmark (Zheng et al., 2025; Wong et al., 2022). For each method, we generate five samples and report their mean results. Since the resolution of the images are low, we use $p = 1$ with $b = 1024$. Evaluation is based on the chi-square errors of the closure phases (χ^2_{cph}) and closure amplitudes (χ^2_{camp}), which quantify how well the reconstructions fit the measurements. Because the black-hole imaging system captures only low spatial frequencies, we follow Akiyama et al. (2019) and compute PSNR for both the original and blurred reconstructions at the system's intrinsic resolution.

Results Fig. 3 shows two examples of black-hole reconstructions of our ZO-APMC method and other gradient-free baselines against the ground truth. ZO-APMC yields black-hole reconstructions

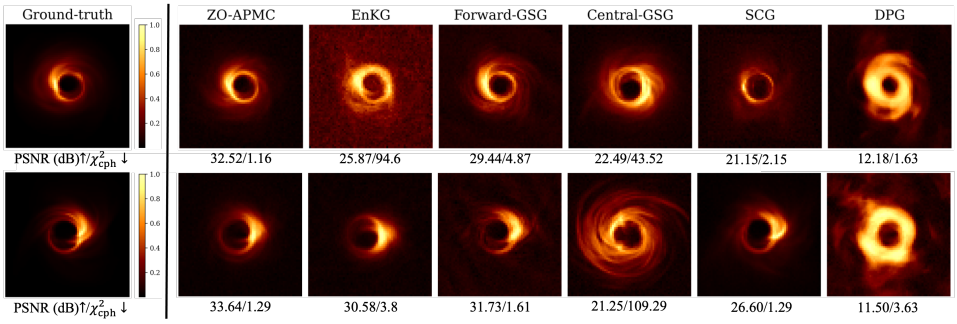

Figure 3: Visualization of samples generated for the black-hole imaging inverse problem. Reconstructions of two examples by gradient-free methods are shown in the top and bottom rows.

Table 2: Quantitative evaluation of reconstructed black-hole images (SD: sample standard deviation).

	PSNR↑	Blurred PSNR↑	χ^2_{cph} ↓	χ^2_{camp} ↓	SD↓
PnPDM	<u>26.48</u>	32.31	<u>11.48</u>	23.54	4.5e-2
DPS	25.61	<u>30.84</u>	12.39	17.72	<u>4.32e-2</u>
APMC	26.23	31.32	11.78	19.23	4.34e-2
Forward-GSG	26.21	31.47	6.77	14.06	2.99e-2
Central-GSG	21.63	23.73	80.31	78.5	4.5e-2
SCG	22.21	25.51	23.72	14.23	1.7e-2
DPG	12.33	14.02	8.17	30.44	**1.6e-2**
EnKG	22.86	27.69	64.37	33.44	0.925
ZO-APMC (Ours)	**26.71**	**32.86**	**5.42**	**11.23**	3.02e-2

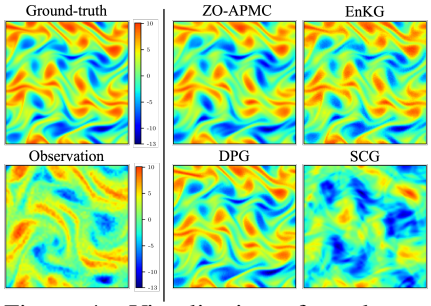

Figure 4: Visualization of results on Navier-Stokes inverse problem.

with visual characteristics most closely match the ground truth among other baselines. Table 2 shows the quantitative comparison. ZO-APMC outperforms all baselines across metrics except SD, which can be mitigated by increasing batch size b at additional cost.

4.4 Navier-Stokes Equation

Problem Setting The Navier–Stokes equation is a standard fluid-dynamics benchmark (Iglesias et al., 2013), widely used from ocean dynamics to climate modeling, where atmospheric observations calibrate initial conditions for numerical forecasts. Computing forward-model gradients via auto-differentiation is impractical because it requires differentiating through a PDE solver. We evaluate the gradient-free methods on 10 test samples from InverseBench using the SGM prior provided by the benchmark, generating five reconstructions per method and reporting the mean performance. We report the NRMSE (relative ℓ_2 error) to evaluate the accuracy of reconstructions with sample SD. For additional details on the experiments, and ablation studies, please refer to Appendix **??**.

Results. Fig. 4 demonstrates that ZO-APMC produces solutions that qualitatively preserve key flow features, comparable to EnKG and DPG, while SCG fails. Moreover, EnKG yields noticeably noisier reconstructions than ZO-APMC. Additional quantitative results and more representative cases, showing our algorithm's performance comparable to the baselines, are provided in Appendix **??**.

5 Conclusion

We proposed ZO-APMC, the first provable derivative-free framework for posterior sampling with a pre-trained SGM prior. It provides non-asymptotic complexity guarantees for reaching an ε-relative Fisher information stationary point and provably converges to the target posterior under mild assumptions using only forward-model evaluations. Toy experiments confirm that our variance-reduction scheme with fixed per-iteration cost ensures convergence in Fisher information across batch sizes, while the annealing mechanism enables accurate sampling from multimodal distributions. On both linear and highly non-linear inverse problems, ZO-APMC matches the performance of state-of-the-art gradient-free methods. The main limitations are higher runtime than gradient-based methods due to Langevin diffusion and the absence of manifold projection as in Chung et al. (2022). Future work includes extending our theoretical analysis with Riemannian zeroth-order derivative estimation (Li et al., 2023) and incorporating faster sampling methods (Yin et al., 2024; Song et al., 2023c).

6 Reproducibility Statement

The full design of the toy experiments is detailed in Appendix A.2. For image-based experiments, we employed publicly available datasets. For all inverse problems other than brain MRI, we adopted the forward models from the reference implementations provided by the InverseBench benchmark Zheng et al. (2025), while for brain MRI experiments we followed the implementation of Sun et al. (2024). Furthermore, we use the original implementations of all baseline methods and include our code as supplementary material with the submission.

References

Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First m87 event horizon telescope results. iv. imaging the central supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L4, 2019.

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2nd edition, 2008.

Krishna Balasubramanian, Sinho Chewi, Murat A Erdogdu, Adil Salim, and Shunshi Zhang. Towards a theory of non-log-concave sampling: first-order stationarity guarantees for langevin monte carlo. In *Conference on Learning Theory*, pp. 2896–2923. PMLR, 2022.

Charles A Bouman and Gregery T Buzzard. Generative plug and play: Posterior sampling for inverse problems. In *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1–7. IEEE, 2023.

Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Andrew A Chael, Michael D Johnson, Katherine L Bouman, Lindy L Blackburn, Kazunori Akiyama, and Ramesh Narayan. Interferometric imaging directly with closure phases and closure amplitudes. *The Astrophysical Journal*, 857(1):23, 2018.

Tony F Chan and Selim Esedoglu. Aspects of total variation regularized l 1 function approximation. *SIAM Journal on Applied Mathematics*, 65(5):1817–1837, 2005.

Hongrui Chen, Holden Lee, and Jianfeng Lu. Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pp. 4735–4763. PMLR, 2023.

Sinho Chewi, Murat A Erdogdu, Mufan Li, Ruoqi Shen, and Matthew S Zhang. Analysis of langevin monte carlo from poincare to log-sobolev. *Foundations of Computational Mathematics*, pp. 1–51, 2024.

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

Hyungjin Chung, Jeongsol Kim, Sehui Kim, and Jong Chul Ye. Parallel diffusion models of operator and image for blind inverse problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6059–6069, 2023.

Florentin Coeurdoux, Nicolas Dobigeon, and Pierre Chainais. Plug-and-play split gibbs sampler: embedding deep generative priors in bayesian inference. *IEEE Transactions on Image Processing*, 33:3496–3507, 2024.

Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G. Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems, 2024. URL https://arxiv.org/abs/2410.00083.

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

Geir Evensen and Peter Jan Van Leeuwen. Assimilation of geosat altimeter data for the agulhas current using the ensemble kalman filter with a quasigeostrophic model. *Monthly weather review*, 124(1):85–96, 1996.

Berthy T Feng, Jamie Smith, Michael Rubinstein, Huiwen Chang, Katherine L Bouman, and William T Freeman. Score-based diffusion models as principled priors for inverse imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10520–10531, 2023.

Guoji Fu and Wee Sun Lee. Approximation and generalization abilities of score-based neural network generative models for sub-gaussian distributions. *arXiv preprint arXiv:2505.10880*, 2025.

Arnaud Guillin, Christian Léonard, Liming Wu, and Nian Yao. Transportation-information inequalities for markov processes. *Probability theory and related fields*, 144(3):669–695, 2009.

Wei Guo, Molei Tao, and Yongxin Chen. Provable benefit of annealed langevin monte carlo for non-log-concave sampling. *arXiv preprint arXiv:2407.16936*, 2024.

Ye He, Kevin Rojas, and Molei Tao. Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion. *Advances in Neural Information Processing Systems*, 37:71122–71161, 2024.

Yinnian He and Weiwei Sun. Stability and convergence of the crank–nicolson/adams–bashforth scheme for the time-dependent navier–stokes equations. *SIAM Journal on Numerical Analysis*, 45(2):837–869, 2007.

Yuchen He and Chihao Zhang. On the query complexity of sampling from non-log-concave distributions. *arXiv preprint arXiv:2502.06200*, 2025.

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

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

Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.

Marco A Iglesias, Kody JH Law, and Andrew M Stuart. Ensemble kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.

Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust compressed sensing mri with deep generative priors. *Advances in neural information processing systems*, 34:14938–14954, 2021.

Bahjat Kawar, Gregory Vaksman, and Michael Elad. Snips: Solving noisy inverse problems stochastically. *Advances in Neural Information Processing Systems*, 34:21757–21769, 2021.

Jonas Knape and Perry De Valpine. Fitting complex population models by combining particle filters with markov chain monte carlo. *Ecology*, 93(2):256–263, 2012.

Guanghui Lan. *First-order and stochastic optimization methods for machine learning*, volume 1. Springer, 2020.

Rémi Laumont, Valentin De Bortoli, Andrés Almansa, Julie Delon, Alain Durmus, and Marcelo Pereyra. Bayesian imaging using plug & play priors: when langevin meets tweedie. *SIAM Journal on Imaging Sciences*, 15(2):701–737, 2022.

Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence for score-based generative modeling with polynomial complexity. *Advances in Neural Information Processing Systems*, 35:22870–22882, 2022.

Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pp. 946–985. PMLR, 2023.

Jiaxiang Li, Krishnakumar Balasubramanian, and Shiqian Ma. Stochastic zeroth-order riemannian derivative estimation and optimization. *Mathematics of Operations Research*, 48(2):1183–1211, 2023.

Zhize Li, Hongyan Bao, Xiangliang Zhang, and Peter Richtárik. Page: A simple and optimal probabilistic gradient estimator for nonconvex optimization. In *International conference on machine learning*, pp. 6286–6295. PMLR, 2021.

Jiaming Liu, Rushil Anirudh, Jayaraman J Thiagarajan, Stewart He, K Aditya Mohan, Ulugbek S Kamilov, and Hyojin Kim. Dolce: A model-based probabilistic diffusion framework for limited-angle ct reconstruction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10498–10508, 2023.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August*, 15(2018):11, 2018.

Radford M Neal. Annealed importance sampling. *Statistics and computing*, 11(2):125–139, 2001.

Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.

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

Kazusato Oko, Shunta Akiyama, and Taiji Suzuki. Diffusion models are minimax optimal distribution estimators. In *International Conference on Machine Learning*, pp. 26517–26582. PMLR, 2023.

Dean S Oliver, Albert C Reynolds, and Ning Liu. *Inverse theory for petroleum reservoir characterization and history matching*. 2008.

Peter J Rousseeuw and Mia Hubert. Robust statistics for outlier detection. *Wiley interdisciplinary reviews: Data mining and knowledge discovery*, 1(1):73–79, 2011.

Litu Rout, Negin Raoof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai. Solving linear inverse problems provably via posterior sampling with latent diffusion models. *Advances in Neural Information Processing Systems*, 36:49960–49990, 2023.

Litu Rout, Yujia Chen, Nataniel Ruiz, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu. Rb-modulation: Training-free personalization of diffusion models using stochastic optimal control. *arXiv preprint arXiv:2405.17401*, 2024.

Abhishek Roy, Lingqing Shen, Krishnakumar Balasubramanian, and Saeed Ghadimi. Stochastic zeroth-order discretizations of langevin diffusions for bayesian inference. *Bernoulli*, 28(3):1810–1834, 2022.

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

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

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.

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

Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023c.

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

Yu Sun, Zihui Wu, Yifan Chen, Berthy T Feng, and Katherine L Bouman. Provable probabilistic imaging using score-based generative priors. *IEEE Transactions on Computational Imaging*, 2024.

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

Cédric Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009.

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

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

George N Wong, Ben S Prather, Vedant Dhruv, Benjamin R Ryan, Monika Mościbrodzka, Chi-kwan Chan, Abhishek V Joshi, Ricardo Yarza, Angelo Ricarte, Hotaka Shiokawa, et al. Patoka: Simulating electromagnetic observables of black hole accretion. *The Astrophysical Journal Supplement Series*, 259(2):64, 2022.

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

Kaylee Yingxi Yang and Andre Wibisono. Convergence of the inexact langevin algorithm and score-based generative models in kl divergence. *arXiv preprint arXiv:2211.01512*, 2022.

Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM computing surveys*, 56(4):1–39, 2023.

Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6613–6623, 2024.

J Zbontar, F Knoll, A Sriram, T Murrell, Z Huang, MJ Muckley, A Defazio, R Stern, P Johnson, M Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. arxiv 2018. *arXiv preprint arXiv:1811.08839*, 1811.

Kaihong Zhang, Caitlyn H Yin, Feng Liang, and Jingbo Liu. Minimax optimality of score-based diffusion models: Beyond the density lower bound assumptions. *arXiv preprint arXiv:2402.15602*, 2024.

Hongkai Zheng, Wenda Chu, Austin Wang, Nikola Kovachki, Ricardo Baptista, and Yisong Yue. Ensemble kalman diffusion guidance: A derivative-free method for inverse problems. *arXiv preprint arXiv:2409.20175*, 2024.

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

A Appendix

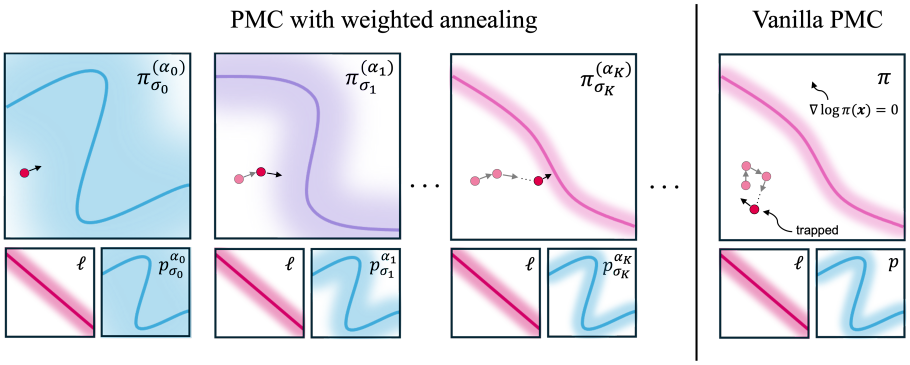

Figure 5: Illustrating how weighted annealing improves the convergence of PMC algorithms by introducing weighted posteriors $\{\pi_{\sigma_k}^{(\alpha_k)}\}$. The solid lines and shaded regions indicate the mean and probability density of the distribution, respectively, while the unshaded area corresponds to $\nabla \log p(\boldsymbol{x}) = 0$. Weighted annealing helps the vanilla PMC algorithm escape plateaus in $\nabla \log p(\boldsymbol{x})$ by gradually reducing the prior's smoothing parameter σ_k and its weight w.r.t. the likelihood ℓ.

A.1 Proofs

Notation. Throughout the proof, we work within the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω denotes the sample space, $\mathcal{F}$ the σ-algebra, and $\mathbb{P}$ the probability measure. For a random variable $X : \Omega \to \mathbb{R}^n$, we write its expectation as

$$\mathbb{E}[X] = \int_\Omega \zeta(\omega)\, \mathbb{P}(d\omega).$$

The posterior distribution of interest is of the form

$$\pi(\boldsymbol{x}|\boldsymbol{y}) \propto \ell(\boldsymbol{y}|\boldsymbol{x}) p(\boldsymbol{x}),$$

where we define $f(\boldsymbol{x}) = -\log \ell(\boldsymbol{y}|\boldsymbol{x})$ and $h(\boldsymbol{x}) = -\log p(\boldsymbol{x})$. Moreover, the gradient of the perturbed log-prior is denoted by $\nabla h_{\sigma_k}(\boldsymbol{x}) := -\nabla \log p_{\sigma_k}(\boldsymbol{x})$. For simplicity, we omit the explicit dependence on $\boldsymbol{y}$. Recall that

$$-\nabla \log \pi(\boldsymbol{x}) = \nabla f(\boldsymbol{x}) + \nabla h(\boldsymbol{x}). \tag{11}$$

We denote the zeroth-order approximation of the forward model gradient as follows

$$\tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}) := \frac{f(\boldsymbol{x}_{k\gamma} + \mu \boldsymbol{u}) - f(\boldsymbol{x}_{k\gamma})}{\mu} \boldsymbol{u}, \tag{12}$$

where $\boldsymbol{u} \sim \mathcal{N}(0, I)$ and $\mu > 0$. The expectation of the zeroth-order approximation is denoted as $\nabla f_\mu(\boldsymbol{x}_{k\gamma}) := \mathbb{E}_{\boldsymbol{u}}\left[\tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u})\right]$. For notational convenience, we also define

$$\Delta_k := \mathbb{E}\left[\|\boldsymbol{x}_{(k+1)\gamma} - \boldsymbol{x}_{k\gamma}\|^2\right], \tag{13}$$

as the expected squared ℓ_2-distance between consecutive iterates. For convenience, we recall the definition of the Kullback–Leibler (KL) divergence between two probability densities ν and π:

$$\mathrm{KL}(\nu \| \pi) = \int_{\mathbb{R}^n} \nu(x) \log \frac{\nu(x)}{\pi(x)}\, dx.$$

Table 3: A conceptual overview of posterior sampling approaches for probabilistic imaging. The *"Annealing"* column highlights distinctions among MCMC-based methods. The *"Black-box access"* column shows whether the corresponding method assumes black-box access and works when gradients of the forward model is unavailable. The $\boldsymbol{A}(\cdot)$ column shows the assumption on the type of forward model. This table extends that of Sun et al. (2024) by incorporating additional black-box posterior sampling algorithms.

Category	Reference	Generative prior	Model agnostic	$\boldsymbol{A}(\cdot)$	Convergence guarantees	Annealing	Black-box access
Variational Bayesian	(Sun & Bouman, 2021)	✗	✗	General	✗	–	✗
	(Feng et al., 2023)	✓	✗	General	✗	–	✗
DM-based	(Song et al., 2021)	✓	✓	Linear	✗	–	✗
	(Chung et al., 2022)	✓	✓	General	✗	–	✗
	(Liu et al., 2023)	✓	✗	Linear	✗	–	✗
	(Tang et al., 2024)	✓	✗	General	✗	–	✓
	(Huang et al., 2024)	✓	✗	General	✗	–	✓
	(Zheng et al., 2024)	✓	✗	General	✗	–	✓
MCMC-based	(Jalal et al., 2021)	✓	✓	Linear	✓[1]	✓	✗
	(Kawar et al., 2021)	✓	✓	Linear	✗	✓	✗
	(Laumont et al., 2022)	✗	✓	General	✓	✗	✗
	(Coeurdoux et al., 2024)	✓	✓	Linear	✗	✗	✗
	(Bouman & Buzzard, 2023)	✗	✓	Linear	✓[2]	✗	✗
	(Sun et al., 2024)	✓	✓	General	✓	✓	✗
MCMC-based	**Ours**	✓	✓	General	✓	✓	✓

[1]Requires $\boldsymbol{A}(\cdot)$ to be a Gaussian random matrix. [2]Guarantees on asymptotic convergence.

The Fisher information (FI) is given by

$$\mathrm{FI}(\nu\|\pi) = \int_{\mathbb{R}^n} \left\|\nabla \log \tfrac{\nu(x)}{\pi(x)}\right\|_2^2 \nu(x)\, dx = \int_{\mathbb{R}^n} \|\nabla \log \nu(x) - \nabla \log \pi(x)\|_2^2 \nu(x)\, dx.$$

The Total Variation (TV) distance between two probability measures μ and ν on a measurable space $(\mathcal{X}, \mathcal{F})$ is given b by

$$\|\mu - \nu\|_{\mathrm{TV}} \;\coloneqq\; \sup_{A\in\mathcal{F}} \big|\mu(A) - \nu(A)\big| \;=\; \frac{1}{2}\int_{\mathcal{X}} \big|d\mu - d\nu\big|.$$

Unless otherwise stated, $\|\cdot\|^2$ denotes the squared ℓ_2-norm, i.e. $\|\cdot\|_2^2$.

A.1.1 LEMMAS

We begin by reviewing the key lemmas from the zeroth-order optimization and non-log-concave sampling literature. The following section summarizes the fundamental properties of zeroth-order approximations that will be used in our analysis.

Lemma 1 ((Lan, 2020)) *Suppose that $f(\boldsymbol{x}) \in C_L^{1,1}$, and let $f_\mu(\boldsymbol{x}) := \mathbb{E}_{\boldsymbol{u}}[f(\boldsymbol{x}+\mu\boldsymbol{u})]$. Then the following statements hold:*

(a) $f_\mu \in C_\mu^{1,1}(\mathbb{R}^d)$, *where* $L_\mu \le L$,

(b) $\|\nabla f_\mu(\boldsymbol{x}) - \nabla f(\boldsymbol{x})\| \le \frac{1}{2}\mu L(d+3)^{\frac{3}{2}}$,

(c) $\mathbb{E}_{\boldsymbol{u}}[\|\tilde{\nabla} f_\mu(\boldsymbol{x},\boldsymbol{u})\|^2] \le \frac{1}{2}\mu^2 L^2(d+6)^3 + 2(d+4)\|\nabla f(\boldsymbol{x})\|_2^2$,

where $\tilde{\nabla} f_\mu(\boldsymbol{x},\boldsymbol{u}) := \frac{f(\boldsymbol{x}+\mu\boldsymbol{u})-f(\boldsymbol{x})}{\mu}\boldsymbol{u}$ *for* $\boldsymbol{u} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_d)$ *and any* $\boldsymbol{x} \in \mathbb{R}^d$, $\mu > 0$.

The following lemma concerns the density evolution of an interpolated diffusion process.

Lemma 2 ((Balasubramanian et al., 2022)) *Consider the stochastic process defined by*

$$\boldsymbol{x}_t := \boldsymbol{x}_0 - t g_0 + \sqrt{2}\boldsymbol{B}_t, \quad \text{with } g_0 = g(\boldsymbol{x}_0), \quad \boldsymbol{x}_0 \sim \nu_0 \tag{14}$$

where g_0 is integrable and $\{B_t\}_{t\ge 0}$ is a standard Brownian motion in $\mathbb{R}^d$ independent of $(\boldsymbol{x}_0, g_0)$. Then, writing ν_t for the probability density of $\boldsymbol{x}_t$, we have

$$\frac{d}{dt}\mathrm{KL}(\nu_t\|\pi) \le -\frac{3}{4}\mathrm{FI}(\nu_t\|\pi) + \mathbb{E}\left[\|\nabla f(\boldsymbol{x}_t) - g_0\|^2\right], \tag{15}$$

where we recall that $\pi \propto e^{-f}$, and the expectation in the last term is with respect to $x_0 \sim \nu_0$ and $x_t \sim \nu_t$.

We also used the following lemma to bound the Fisher information, which is taken from (Chewi et al., 2024).

Lemma 3 ((Chewi et al., 2024)) *Assume that $\nabla \log \pi(\boldsymbol{x})$ is L_π-Lipschitz. For any probability measure ν, it holds that*

$$\mathbb{E}_\mu\left[\|\nabla \log \pi(\boldsymbol{x})\|^2\right] \le \mathrm{FI}\left(\nu\|\pi\right) + 2dL_\pi. \tag{16}$$

We use the following lemma to derive an upper bound on Total Variation (TV) distance.

Lemma 4 ((Guillin et al., 2009)) *If π satisfies a Poincaré inequality, i.e. for every smooth, compactly supported $f : \mathbb{R}^d \to \mathbb{R}$,*

$$\mathrm{Var}_\pi(f) \;\le\; C_{\mathrm{PI}}\,\mathbb{E}_\pi\big[\|\nabla f\|^2\big],$$

then for any probability measure μ,

$$\|\mu - \pi\|_{\mathrm{TV}}^2 \;\le\; 4C_{\mathrm{PI}}\,\mathrm{FI}(\mu\|\pi).$$

A.1.2 Proof of Proposition 1

For simplicity, let our PAGE estimator be defined as

$$\boldsymbol{g}_k \coloneqq \begin{cases} \dfrac{1}{b}\displaystyle\sum_{i\in I}\tilde{\nabla} f_\mu(\boldsymbol{x}_k,\boldsymbol{u}_i), & \boldsymbol{B}_k = \text{ref}, \\ \boldsymbol{g}_{k-1} + \dfrac{1}{b'}\displaystyle\sum_{i\in I'}\Big(\tilde{\nabla} f_\mu(\boldsymbol{x}_k,\boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k-1},\boldsymbol{u}_i)\Big), & \boldsymbol{B}_k = \text{corr}, \end{cases} \tag{17}$$

where "ref" and "corr" denote the "refresh" and "correction" branches of the estimate and $\boldsymbol{B}_k \in \{\text{corr}, \text{ref}\}$ is a random variable such that $\mathbb{P}(\boldsymbol{B}_k = \text{ref}) = p$ and $\mathbb{P}(\boldsymbol{B}_k = \text{corr}) = 1-p$. Additionally, define the mini-batch estimators as

$$\tilde{\boldsymbol{v}}_b(\boldsymbol{x}_k) = \frac{1}{b}\sum_{i\in I}\tilde{\nabla} f_\mu(\boldsymbol{x}_k,\boldsymbol{u}_i) \quad \text{and} \quad \boldsymbol{\delta}_k = \frac{1}{b'}\sum_{i\in I'}\Big(\tilde{\nabla} f_\mu(\boldsymbol{x}_k,\boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k-1},\boldsymbol{u}_i)\Big), \tag{18}$$

where $|I| = b$ and $|I'| = b'$. Then, $\boldsymbol{g}_k$ can be written as

$$\boldsymbol{g}_k \coloneqq \begin{cases} \tilde{\boldsymbol{v}}_b(\boldsymbol{x}_k), & \boldsymbol{B}_k = \text{ref}, \\ \boldsymbol{g}_{k-1} + \boldsymbol{\delta}_k, & \boldsymbol{B}_k = \text{corr}. \end{cases} \tag{19}$$

Let $\mathcal{F}_k \coloneqq \sigma(\boldsymbol{x}_0, \boldsymbol{B}_1, \boldsymbol{Z}_1, \boldsymbol{u}_1, \ldots, \boldsymbol{x}_k, \boldsymbol{B}_k, \boldsymbol{Z}_k, \boldsymbol{u}_k)$ be the sigma-algebra generated by all the random variables revealed up to the end of iteration k. From (6), recall that $\boldsymbol{Z}_k$ is due to the discretization of the Langevin diffusion. Then, conditioning on the history $\mathcal{F}_{k-1}$ and on $\boldsymbol{x}_k$, we have

$$\mathbb{E}[\tilde{\boldsymbol{v}}_b(\boldsymbol{x}_k)|\mathcal{F}_{k-1}, \boldsymbol{x}_k] = \nabla f_\mu(\boldsymbol{x}_k), \quad \mathbb{E}[\boldsymbol{\delta}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k] = \nabla f_\mu(\boldsymbol{x}_k) - \nabla f_\mu(\boldsymbol{x}_{k-1}). \tag{20}$$

Then, we have $\mathbb{E}[\boldsymbol{g}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k] = \nabla f_\mu(\boldsymbol{x}_k)$. Using this property inductively, we can obtain $\mathbb{E}[\boldsymbol{g}_k] = \mathbb{E}[\nabla f_\mu(\boldsymbol{x}_k)]$ after taking the expectation of both sides. Therefore, if we define the error as $\boldsymbol{e}_k \coloneqq \boldsymbol{g}_k - \nabla f_\mu(\boldsymbol{x}_k)$, then $\mathbb{E}[\boldsymbol{e}_k] = 0$, which implies that $\boldsymbol{g}_k$ is unbiased estimate of $\nabla f_\mu(\boldsymbol{x}_k)$. Let's consider the error propagation at each branch separately.

For the correction branch, assume $\boldsymbol{B}_k = \text{corr}$ is true. We can define zero-mean fluctuation at step k as

$$\tilde{\boldsymbol{\delta}}_k \coloneqq \boldsymbol{\delta}_k - \mathbb{E}[\boldsymbol{\delta}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k], \tag{21}$$

where $\mathbb{E}[\tilde{\boldsymbol{\delta}}_k|\mathcal{F}_{k-1}|\boldsymbol{x}_k] = 0$. Then,

$$\begin{aligned}\boldsymbol{e}_k = \boldsymbol{g}_k - \nabla f_\mu(\boldsymbol{x}_k) &= \underbrace{\boldsymbol{g}_{k-1} - \nabla f_\mu(\boldsymbol{x}_{k-1})}_{\boldsymbol{e}_{k-1}} + \boldsymbol{\delta}_k - (\nabla f_\mu(\boldsymbol{x}_k) - \nabla f_\mu(\boldsymbol{x}_{k-1})) \\ &= \boldsymbol{e}_{k-1} + \tilde{\boldsymbol{\delta}}_k. \end{aligned} \tag{22}$$

Note that $\boldsymbol{u}_i \sim \mathcal{N}(0, I)$ random vectors selected at step k are independent of the ones selected at step $k-1$. Therefore, $\tilde{\boldsymbol{\delta}}_k$ is conditionally independent of $\boldsymbol{e}_{k-1}$, so we can write

$$\text{Cov}(\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k, \boldsymbol{B}_k = \text{corr}) = \text{Cov}(\boldsymbol{e}_{k-1}|\mathcal{F}_{k-1}) + \text{Cov}(\tilde{\boldsymbol{\delta}}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k). \tag{23}$$

In the refresh branch, we have $\boldsymbol{B}_k = \text{ref}$. Then, the error term can be written as $\boldsymbol{e}_k = \tilde{\boldsymbol{v}}_b(\boldsymbol{x}_k) - \nabla f_\mu(\boldsymbol{x}_k)$ and its covariance is

$$\text{Cov}(\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k, \boldsymbol{B}_k = \text{corr}) = \Sigma_b(\boldsymbol{x}_k), \tag{24}$$

where $\Sigma_b(\boldsymbol{x}_k) \coloneqq \frac{1}{b}\text{Cov}(\tilde{\nabla} f_m u(\boldsymbol{x}_k, \boldsymbol{u}_i))$. Furthermore, using the definition of $\boldsymbol{e}_k$, we have

$$\mathbb{E}[\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k, \boldsymbol{B}_k = \text{corr}] = \boldsymbol{e}_{k-1} \quad \text{and} \quad \mathbb{E}[\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k, \boldsymbol{B}_k = \text{ref}] = 0. \tag{25}$$

Using the law of total variance, we can write the covariance matrix conditioned on the history $\mathcal{F}_{k-1}$ and $\boldsymbol{x}_k$ as

$$\text{Cov}(\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k) = \mathbb{E}[\text{Cov}(\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k, \boldsymbol{B}_k)] + \text{Cov}\left(\mathbb{E}[\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k, \boldsymbol{B}_k]\right). \tag{26}$$

If we plug (23) and (24) with the conditional means in (25), we get

$$\text{Cov}(\boldsymbol{e}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k) = (1-p)\left(\text{Cov}(\boldsymbol{e}_{k-1}|\mathcal{F}_{k-1}) + \text{Cov}(\tilde{\delta}_k|\mathcal{F}_{k-1}, \boldsymbol{x}_k)\right) + p\Sigma_b(\boldsymbol{x}_k) + p(1-p)\boldsymbol{e}_{k-1}\boldsymbol{e}_{k-1}^T. \tag{27}$$

Taking the expectation of both sides, we get

$$\text{Cov}(\boldsymbol{e}_k) = (1-p^2)\text{Cov}(\boldsymbol{e}_{k-1}) + (1-p)\mathbb{E}[\text{Cov}(\tilde{\boldsymbol{\delta}}_k)] + p\mathbb{E}[\Sigma_b(\boldsymbol{x}_k)]. \tag{28}$$

The factor $(1-p^2)$ is the contraction on the previous error covariance in expectation. Note that $\boldsymbol{u}_i$ are i.i.d. and recalling the definition of $\tilde{\delta}_k$, we have

$$\text{Cov}(\tilde{\delta}_k) \preceq \text{Cov}(\delta_k) \preceq \frac{4dL_{f_1}^2}{b'\mu^2}\|\boldsymbol{x}_k - \boldsymbol{x}_{k-1}\|^2 I, \tag{29}$$

where we use the Assumption 1 to get the second inequality. This shows that the correction-step noise is small when the iterate moves only a little between steps. Taking the trace of both sides in (28) and plugging (29), we get

$$\text{Cov}(\boldsymbol{e}_k) \preceq (1-p^2)\text{Cov}(\boldsymbol{e}_{k-1}) + \frac{4d(1-p)L_{f_2}^2}{b'\mu^2}\mathbb{E}[\|\boldsymbol{x}_k - \boldsymbol{x}_{k-1}\|^2]I + p\mathbb{E}[\Sigma_b(\boldsymbol{x}_k)]. \tag{30}$$

Equivalently, this can be written as

$$\mathbb{E}[\|\boldsymbol{e}_k\|^2] \leq (1-p^2)\mathbb{E}[\|\boldsymbol{e}_{k-1}\|^2] + \frac{4d(1-p)L_{f_2}^2}{b'\mu^2}\mathbb{E}[\|\boldsymbol{x}_k - \boldsymbol{x}_{k-1}\|^2] + \frac{p}{b}\sigma^2, \tag{31}$$

where $\sigma_g^2 \leq \sigma^2$ and $\sigma_g^2(\boldsymbol{x}) \coloneqq \text{Tr}\left(\text{Cov}(\tilde{\nabla} f_\mu(\boldsymbol{x}, \boldsymbol{u}_i))\right)$.

A.1.3 PROOF OF THEOREM 1

We can construct the following interpolation for ZO-APMC

$$\boldsymbol{x}_t = \boldsymbol{x}_{k\gamma} - (t - k\gamma)\mathcal{G}(\boldsymbol{x}_{k\gamma}) + \sqrt{2}(\boldsymbol{B}_t - \boldsymbol{B}_{k\gamma}) \quad \text{for} \quad t \in [k\gamma, (k+1)\gamma] \tag{32}$$

where $\mathcal{G}(\boldsymbol{x}_{k\gamma}) = \boldsymbol{g}_k - \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_{k\gamma})$, $\boldsymbol{g}_k$ is an estimate of the forward model gradient with zeroth-order approximation and variance-reduction mechanism, α_k and σ_k are annealing parameters. By Assumption 2 and 4 and triangle inequality, we know that the target posterior score function $\nabla \log \pi(\boldsymbol{x})$ is Lipschitz continuous with Lipschitz constant $L_\pi = L_p + L_{f_1}$. Furthermore, by Assumptions 3 and 5, the error between the prior score function and the score estimate scaled by annealing parameter can be bounded by

$$\|\nabla h(\boldsymbol{x}_{k\gamma}) + \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_{k\gamma})\| \leq \sigma_k C + \varepsilon_{\sigma_k} + (\alpha_k - 1)R_s, \tag{33}$$

where we recall $\nabla h(\boldsymbol{x}_{k\gamma}) = -\nabla \log p(\boldsymbol{x}_{k\gamma})$. Note that we add and subtract $\nabla h_{\sigma_k}(\boldsymbol{x}_{k\gamma})$ and use triangle inequality. Now we can provide the proof for Theorem 1. Combining Lemma 2 with the interpolation argument in (32), it follows that for every $t \in [k\gamma, (k+1)\gamma]$,

$$\frac{d}{dt}\mathrm{KL}(\nu_t \| \pi) \leq -\frac{3}{4}\mathrm{FI}(\nu_t \| \pi) + \mathbb{E}\left[\|\nabla \log \pi(\boldsymbol{x}_t) + \boldsymbol{g}_k - \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_{k\gamma}, \sigma_k)\|^2\right]. \tag{34}$$

Adding and subtracting the following values $\nabla f_\mu(\boldsymbol{x}_t)$, $\nabla f_\mu(\boldsymbol{x}_{k\gamma})$, $\nabla h(\boldsymbol{x}_{k\gamma})$, $\nabla h_{\sigma_k}(\boldsymbol{x}_{k\gamma})$ inside the expectation and using the convexity of ℓ_2 norm with the upper bound in (33), we get

$$\begin{aligned}\frac{d}{dt}\mathrm{KL}(\nu_t \| \pi) \leq &-\frac{3}{4}\mathrm{FI}(\nu_t \| \pi) + 4\mathbb{E}\left[\|\boldsymbol{g}_k - \nabla f_\mu(\boldsymbol{x}_{k\gamma})\|^2\right] + 4L_\pi^2 \mathbb{E}\left[\|\boldsymbol{x}_t - \boldsymbol{x}_{k\gamma}\|^2\right] \\ &+ \mu^2 L_{f_1}^2 (d+3)^3 + 4(\sigma_k C + \varepsilon_{\sigma_k} + (\alpha_k - 1)R_s)^2.\end{aligned} \tag{35}$$

Let $e_k^2 \coloneqq \mathbb{E}\left[\left\|\boldsymbol{g}_k - \nabla f_\mu(\mathbf{x}_{k\gamma})\right\|^2\right]$, which quantifies the squared error between the zeroth-order estimate g_k and the true score $\nabla f_\mu(\mathbf{x}_{k\gamma})$ of the μ-perturbed forward model. Here the expectation is taken with respect to both the randomness of the zeroth-order approximation and the measure $\mathcal{F}_k$ associated with the data $\boldsymbol{x}_{k\gamma}$. Note that the bias term due to the zeroth-order approximation appears as the fourth term of the previous inequality. Using the definition of $\boldsymbol{g}_k$, we can expand the error term as

$$\begin{aligned}e_{k+1}^2 = p\,\mathbb{E}&\left[\left\|\nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - \frac{1}{b}\sum_{i=1}^{b} \tilde{\nabla} f_\mu\big(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i\big)\right\|^2\right] \\ &+ (1-p)\,\mathbb{E}\left[\left\|\nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - g_k - \frac{1}{b'}\sum_{i=1}^{b'}\Big(\tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}_i)\Big)\right\|^2\right]\end{aligned} \tag{36}$$

where b', b denote the small and large batch sizes, respectively, and $\boldsymbol{u}_i \sim \mathcal{N}(0, I)$ in $\mathbb{R}^d$. We can upper bound the first expectation as

$$\mathbb{E}\left[\left\|\nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - \frac{1}{b}\sum_{i=1}^{b} \tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma,}, \boldsymbol{u}_i)\right\|^2\right] \leq \frac{1}{b}\mathbb{E}\left[\left\|\nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i)\right\|^2\right] \tag{37}$$

$$\leq \frac{1}{b}\mathbb{E}\left[\left\|\tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i)\right\|^2\right] \tag{38}$$

$$\leq \frac{L_{f_2}^2}{b}\mathbb{E}\left[\|\boldsymbol{u}_i\|^4\right] \tag{39}$$

$$= \frac{d(d+2)L_{f_2}^2}{b}. \tag{40}$$

In (37), we use the fact that the random variables $\boldsymbol{u}_i$ are i.i.d. In (38), we use the second-moment bound on the variance. Finally, in (39), we use the zeroth-order definition of $\tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}_i)$ with the

Assumption 1 and evaluate the expectation under $\boldsymbol{u}_i \sim \mathcal{N}(0, I)$ to get (40). Plugging this upper bound into (36), we get

$$e_{k+1}^2 \leq \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)\mathbb{E}\Bigg[\Big\|\nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - g_k - \frac{1}{b'}\sum_{i=1}^{b'}\Big(\tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}_i)\Big)\Big\|^2\Bigg] \tag{41}$$

$$= \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)\mathbb{E}\Bigg[\Big\|\nabla f_\mu(\boldsymbol{x}_{k\gamma}) - g_k + \nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - \nabla f_\mu(\boldsymbol{x}_{k\gamma}) - \frac{1}{b'}\sum_{i=1}^{b'}\Big(\tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}_i)\Big)\Big\|^2\Bigg] \tag{42}$$

$$= \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)e_k^2 + \frac{1}{b'}\mathbb{E}\Bigg[\Big\|\nabla f_\mu(\boldsymbol{x}_{(k+1)\gamma}) - \nabla f_\mu(\boldsymbol{x}_{k\gamma}) - \Big(\tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}_i)\Big)\Big\|^2\Bigg] \tag{43}$$

$$\leq \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)e_k^2 + \frac{1}{b'}\mathbb{E}\Bigg[\Big\|\tilde{\nabla} f_\mu(\boldsymbol{x}_{(k+1)\gamma}, \boldsymbol{u}_i) - \tilde{\nabla} f_\mu(\boldsymbol{x}_{k\gamma}, \boldsymbol{u}_i)\Big\|^2\Bigg] \tag{44}$$

$$= \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)e_k^2 + \frac{1}{\mu^2 b'}\mathbb{E}\Bigg[\Big\|\Big(f(\boldsymbol{x}_{(k+1)\gamma} + \mu\boldsymbol{u}_i) - f(\boldsymbol{x}_{k\gamma} + \mu\boldsymbol{u}_i)\Big) - \Big(f(\boldsymbol{x}_{(k+1)\gamma}) - f(\boldsymbol{x}_{k\gamma})\Big)\Big\|^2\Big\|\boldsymbol{u}_i\Big\|^2\Bigg] \tag{45}$$

$$\leq \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)e_k^2 + \frac{4L_{f_2}^2\Delta_k}{\mu^2 b'}\mathbb{E}\Big[\|\boldsymbol{u}_i\|^2\Big] \tag{46}$$

$$= \frac{pd(d+2)L_{f_2}^2}{b} + (1-p)e_k^2 + \frac{4dL_{f_2}^2\Delta_k}{\mu^2 b'} \tag{47}$$

where $\Delta_k \coloneqq \mathbb{E}\left[\|\boldsymbol{x}_{(k+1)\gamma} - \boldsymbol{x}_{k\gamma}\|^2\right]$. Note that we add and subtract $\nabla f_\mu(\boldsymbol{x}_{k\gamma})$ in (42). To get (43), we use the fact that random variables $\boldsymbol{u}_i \sim \mathcal{N}(0, I)$ are i.i.d., and calculate conditional expectation conditioned with respect to $\mathcal{F}_k$ and then use the definition of e_k^2. We use second-moment bound on variance in (44) and use the zeroth-order definition to get (45). Following that, we first apply Assumption 1 and then exploit the independence between $\boldsymbol{x}_{k\gamma}$, $\boldsymbol{x}_{(k+1)\gamma}$, and $\boldsymbol{u}_i$ to obtain (46). Dividing both sides by p and rearranging the terms, we get an upper bound on the error term

$$e_k^2 \leq \frac{d(d+2)L_{f_2}^2}{b} + \left(\frac{1-p}{p}\right)\frac{4dL_{f_2}^2}{\mu^2 b'}\Delta_k - \frac{1}{p}(e_{k+1}^2 - e_k^2) \tag{48}$$

Plugging this upper bound into (35), we get

$$\frac{d}{dt}\mathrm{KL}(\nu_t\|\pi) \leq -\frac{3}{4}\mathrm{FI}(\nu_t\|\pi) + 4L_\pi^2\mathbb{E}\Big[\|\boldsymbol{x}_t - \boldsymbol{x}_{k\gamma}\|^2\Big] + \mu^2 L_{f_1}^2(d+3)^2 + \frac{4d(d+2)L_{f_2}^2}{b} + \left(\frac{1-p}{p}\right)\frac{16dL_{f_2}^2}{\mu^2 b'}\mathbb{E}\Big[\|\boldsymbol{x}_{(k+1)\gamma} - \boldsymbol{x}_{k\gamma}\|^2\Big] - \frac{4}{p}(e_{k+1}^2 - e_k^2) + 4(\sigma_k C + \varepsilon_{\sigma_k} + (\alpha_k - 1)R_s)^2 \tag{49}$$

By the interpolation argument in (32), we have

$$\mathbb{E}[\|\boldsymbol{x}_t - \boldsymbol{x}_{k\gamma}\|^2] = (t-k\gamma)^2\mathbb{E}\left[\|\mathcal{G}(x_{k\gamma})\|^2\right] + 2(t-k\gamma)d \tag{50}$$

$$\leq \gamma^2\mathbb{E}\left[\|\mathcal{G}(x_{k\gamma})\|^2\right] + 2\gamma d \tag{51}$$

$$= \mathbb{E}[\|\boldsymbol{x}_{(k+1)\gamma} - \boldsymbol{x}_{k\gamma}\|^2] \tag{52}$$

for $t \in [k\gamma, (k+1)\gamma]$ because dimensionality of vectors $d > 0$ and $\mathbb{E}\left[\|\mathcal{G}(x_{k\gamma})\|^2\right] \geq 0$. We use the bound on (52) in (49) and get

$$\frac{d}{dt}\mathrm{KL}(\nu_t\|\pi) \leq -\frac{3}{4}\mathrm{FI}(\nu_t\|\pi) + \mu^2 L_{f_1}^2(d+3)^2 + \frac{4d(d+2)L_{f_2}^2}{b} - \frac{4}{p}(e_{k+1}^2 - e_k^2)$$
$$+ 4L_m^2\left[1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right]\Delta_k + 4(\sigma_k C + \varepsilon_{\sigma_k} + (\alpha_k - 1)R_s)^2 \tag{53}$$

where $L_m \coloneqq \max\{L_{f_2}, L_\pi\}$. We use the interpolation argument given in (32) to put an upper bound on Δ_k as

$$\Delta_k = \gamma^2\mathbb{E}[\|\boldsymbol{g}_k - \alpha_k\mathcal{S}_\theta(\boldsymbol{x}_{k\gamma}, \sigma_k)\|^2] + 2\gamma d \tag{54}$$

$$\leq 5\gamma^2 e_k^2 + \frac{5\gamma^2\mu^2 L_{f_1}^2(d+3)^3}{4} + 5\gamma^2 L_\pi^2\Delta_k + 5\gamma^2\mathbb{E}\left[\|\nabla\log\pi(\boldsymbol{x}_t)\|^2\right]$$
$$+ 5\gamma^2(\sigma_k C + \varepsilon_{\sigma_k}2 + (\alpha_k - 1)R_s)^2 + 2\gamma d \tag{55}$$

where we add and subtract $\nabla f_\mu(\mathbf{x}_{k\gamma})$, $\nabla f(\mathbf{x}_{k\gamma})$, $\nabla f(\mathbf{x}_t)$, $\nabla h(\mathbf{x}_t)$, $\nabla h(\mathbf{x}_{k\gamma})$, and $\nabla h_{\sigma_k}(\mathbf{x}_{k\gamma})$ inside the expectation, and then apply the convexity of the ℓ_2 norm together with the part (b) of Assumption 5 to obtain (55). Using the bound on e_k^2 in (48), we get

$$\Delta_k \leq 5\gamma^2 L_m^2\left[1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right]\Delta_k + \frac{5\gamma^2 d(d+2)L_{f_2}^2}{b} + \frac{5\gamma^2\mu^2 L_{f_1}^2(d+3)^3}{4}$$
$$+ 5\gamma^2\mathbb{E}\left[\|\nabla\log\pi(\boldsymbol{x}_t)\|^2\right] - \frac{5\gamma^2}{p}\left(e_{k+1}^2 - e_k^2\right) + 5\gamma^2(\sigma_k C + \varepsilon_{\sigma_k}^2 + (\alpha_k - 1)R_s)^2 + 2\gamma d \tag{56}$$

Assume that $\gamma^2 \leq \left[85L_m^2\left(1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right)\right]^{-1}$, then we have

$$\frac{16}{17}\Delta_k \leq \frac{5\gamma^2 d(d+2)L_{f_2}^2}{b} + \frac{5\gamma^2\mu^2 L_{f_1}^2(d+3)^3}{4} + 5\gamma^2\mathbb{E}\left[\|\nabla\log\pi(\boldsymbol{x}_t)\|^2\right] - \frac{5\gamma^2}{p}\left(e_{k+1}^2 - e_k^2\right)$$
$$+ 5\gamma^2(\sigma_k C + \varepsilon_{\sigma_k}^2 + (\alpha_k - 1)R_s)^2 + 2\gamma d \tag{57}$$

Multiplying both sides by $\frac{17}{16}$, we get

$$\Delta_k \leq \frac{85\gamma^2 d(d+2)L_{f_2}^2}{16b} + \frac{85\gamma^2\mu^2 L_{f_1}^2(d+3)^3}{64} + \frac{85}{16}\gamma^2\mathbb{E}\left[\|\nabla\log\pi(\boldsymbol{x}_t)\|^2\right]$$
$$- \frac{85\gamma^2}{16p}\left(e_{k+1}^2 - e_k^2\right) + \frac{85}{16}\gamma^2(\sigma_k C + \varepsilon_{\sigma_k}^2 + (\alpha_k - 1)R_s)^2 + \frac{17}{8}\gamma d \tag{58}$$

We can use Lemma 3 to put an upper bound to the third term

$$\Delta_k \leq \frac{85\gamma^2 d(d+2)L_{f_2}^2}{16b} + \frac{85\gamma^2\mu^2 L_{f_1}^2(d+3)^3}{64} + \frac{85}{16}\gamma^2\mathrm{FI}(\nu_t\|\pi) + \frac{85}{8}\gamma^2 L_\pi d$$
$$- \frac{85\gamma^2}{16p}\left(e_{k+1}^2 - e_k^2\right) + \frac{85}{16}\gamma^2(\sigma_k C + \varepsilon_{\sigma_k}^2 + (\alpha_k - 1)R_s)^2 + \frac{17}{8}\gamma d. \tag{59}$$

We can combine the fourth and the last term by using $\gamma \leq \left[85L_m^2\left(1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right)\right]^{-\frac{1}{2}}$ and concavity of the square root. Note that

$$\gamma \leq \frac{1}{\sqrt{85L_m^2\left(1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right)}} \leq \frac{2}{L_m\sqrt{170} + \sqrt{\left(\frac{1-p}{p}\right)\frac{8d}{\mu^2 b'}}}. \tag{60}$$

We can use this and get

$$\frac{85}{8}\gamma^2 dL_\pi \leq \frac{85\gamma d}{4}\left(\frac{L_m}{L_m\sqrt{170}+\sqrt{\left(\frac{1-p}{p}\right)\frac{8d}{\mu^2 b'}}}\right) \leq \frac{85}{4\sqrt{170}}\gamma d. \tag{61}$$

Finally, we get the upper bound for Δ_k as

$$\begin{aligned}\Delta_k \leq{}& \frac{85\gamma^2 d(d+2)L_{f_2}^2}{16b} + \frac{85\gamma^2\mu^2 L_{f_1}^2(d+3)^3}{64} + \frac{85}{16}\gamma^2\mathrm{FI}(\nu_t\|\pi) + 4\gamma d \\ &- \frac{85\gamma^2}{16p}\left(e_{k+1}^2 - e_k^2\right) + \frac{85}{16}\gamma^2(\sigma_k C + \varepsilon_{\sigma_k}^2 + (\alpha_k-1)R_s)^2.\end{aligned} \tag{62}$$

We define $\phi(\mu) \coloneqq 1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}$ for convenience. Plugging (62) into (53), we get

$$\begin{aligned}\frac{d}{dt}\mathrm{KL}(\nu_t\|\pi) \leq{}& \left(-\frac{3}{4} + \frac{85\gamma^2 L_m^2\phi(\mu)^2}{4}\right)\mathrm{FI}(\nu_t\|\pi) + \left(1 + \frac{85\gamma^2 L_m^2\phi(\mu)}{16}\right)\frac{4d(d+2)L_{f_2}^2}{b} \\ &+ \left(1 + \frac{85\gamma^2 L_m^2\phi(\mu)}{16}\right)\mu^2 L_{f_1}^2(d+3)^3 - \frac{4}{p}\left(1 + \frac{85\gamma^2 L_m^2\phi(\mu)}{16}\right)\left(e_{k+1}^2 - e_k^2\right) \\ &+ 4\left(1 + \frac{85\gamma^2 L_m^2\phi(\mu)}{16}\right)(\sigma_k C + \varepsilon_{\sigma_k} + (\alpha_k-1)R_s)^2 + 4\gamma d L_m^2\phi(\mu) \qquad (63) \\ \leq{}& -\frac{1}{2}\mathrm{FI}(\nu_t\|\pi) + \frac{17d(d+2)L_{f_2}^2}{4b} + \frac{17}{16}\mu^2 L_{f_1}^2(d+3)^3 + 4\gamma d L_m^2\phi(\mu) \\ &- \frac{4}{p}\left(1 + \frac{85\gamma^2 L_m^2\phi(\mu)}{16}\right)\left(e_{k+1}^2 - e_k^2\right) + \frac{17}{4}(\sigma_k C + \varepsilon_{\sigma_k} + (\alpha_k-1)R_s)^2, \qquad (64)\end{aligned}$$

where we use the fact that $85\gamma^2 L_m^2\phi(\mu) \leq 1$ to get (64). Integrating both sides between $[k\gamma, (k+1)\gamma]$, we get

$$\begin{aligned}\mathrm{KL}(\nu_{(k+1)\gamma}\|\pi) - \mathrm{KL}(\nu_{k\gamma}\|\pi) \leq{}& -\frac{1}{2}\int_{k\gamma}^{(k+1)\gamma}\mathrm{FI}(\nu_t\|\pi)dt + \frac{17\gamma d(d+2)L_{f_2}^2}{4b} + \frac{17}{16}\gamma\mu^2 L_{f_1}^2(d+3)^3 \\ &+ 4\gamma^2 d L_m^2\phi(\mu) - \frac{\gamma}{p}\left(4 + \frac{85\gamma^2 L_m^2\phi(\mu)}{16}\right)\left(e_{k+1}^2 - e_k^2\right) \\ &+ \frac{51\gamma}{4}(\sigma_k^2 C^2 + \varepsilon_{\sigma_k}^2 + (\alpha_k-1)^2 R_s^2)\end{aligned} \tag{65}$$

Note that for the last term, we use Jensen's inequality. Let $\mathcal{L}_k \coloneqq \mathrm{KL}(\nu_{k\gamma}\|\pi) + \frac{\gamma}{p}\left[4 + \frac{85}{16}\gamma^2 L_m^2\phi(\mu)\right]e_k^2$, iterating for $k = 0, \ldots, N-1$, multiplying both sides by $\frac{2}{N\gamma}$, rearranging the terms and using the fact that $\mathcal{L}_k \geq 0$, we get

$$\begin{aligned}\frac{1}{N\gamma}\int_0^{N\gamma}\mathrm{FI}(\nu_t\|\pi)dt \leq{}& \frac{2\mathcal{L}_0}{N\gamma} + \frac{17d(d+2)L_{f_2}^2}{2b} + \frac{17}{8}\mu^2 L_{f_1}^2(d+3)^3 + 8\gamma L_m^2 d\phi(\mu) \\ &+ \bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2,\end{aligned} \tag{66}$$

where $\bar{\sigma}^2 \coloneqq \frac{51C}{2N}\sum_{k=0}^{N-1}\sigma_k^2$, $\bar{\varepsilon}_\sigma^2 \coloneqq \frac{51}{2N}\sum_{k=0}^{N-1}\varepsilon_{\sigma_k}^2$, and $\bar{\alpha}_{-1} \coloneqq \frac{51}{2N}\sum_{k=0}^{N-1}(\alpha_k-1)^2 R_s^2$. Since $\gamma^2 L_m^2\phi(\mu) \leq \frac{1}{85}$, we have

$$\mathcal{L}_0 = \mathrm{KL}(\nu_0\|\pi) + \frac{4\gamma}{p}\left(1 + \frac{85}{64}\gamma^2 L_m^2\phi(\mu)\right)e_0^2 \leq \mathrm{KL}(\nu_0\|\pi) + \frac{5\gamma e_0^2}{p}. \tag{67}$$

We can let $C_0 = 2\mathrm{KL}(\nu_0\|\pi) + \frac{10\gamma e_0^2}{p}$ be a constant and this completes the proof of the first statement. To find an ε-approximate solution, we choose the step size as

$$\gamma = \frac{1}{2L_m}\sqrt{\frac{C_0}{Nd\phi(\mu)}} \tag{68}$$

so that we get

$$\frac{1}{N\gamma}\int_0^{N\gamma}\mathrm{FI}(\nu_t\|\pi)dt \leq 6L_m\sqrt{\frac{C_0 d\phi(\mu)}{N}} + \frac{17d(d+2)L_{f_2}^2}{2b} + \frac{17}{8}\mu^2 L_{f_1}^2(d+3)^3 + \bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2. \tag{69}$$

We know that $\{\alpha\}_{k=0}^{N-1}$ and $\{\sigma\}_{k=0}^{N-1}$ are nonincreasing sequences such that

$$\alpha_0 > \alpha_1 > \ldots > \alpha_K = \cdots = \alpha_{N-1} = 1 \quad \text{and} \quad \sigma_0 > \sigma_1 > \ldots > \sigma_K = \cdots = \sigma_{N-1} \approx 0. \tag{70}$$

Thus, we can choose the annealing schedules as $\sigma_n = O(n^{-\beta})$, $\alpha_n = 1 + O(n^{-\beta})$ for $\beta > 1/2$ so that we have $\bar{\sigma}^2 \leq \frac{\varepsilon}{6}$ and $\bar{\alpha}_{-1}^2 \leq \frac{\varepsilon}{6}$. In addition, if the training error is $\varepsilon_{\sigma_n} = O(n^{-\beta})$ for $\beta > 1/2$, then

$$\bar{\varepsilon}_\sigma^2 = \frac{51}{2N}\sum_{k=0}^{N-1}\varepsilon_{\sigma_k}^2 = \frac{51S}{2N} = O\left(\frac{1}{N}\right), \tag{71}$$

where sum converges to a constant S. For the convergence of other terms, we can choose a sufficiently large batch size b and sufficiently small smoothing parameter μ for the zeroth-order estimate as

$$b = \left\lceil \frac{51d(d+2)L_{f_2}^2}{\varepsilon} \right\rceil \quad \text{and} \quad \mu^2 = \frac{4\varepsilon}{51L_{f_1}^2(d+3)^3} \tag{72}$$

where $\lceil . \rceil$ rounds the value to the larger closest integer. Using these values, we get

$$\frac{1}{N\gamma}\int_0^{N\gamma}\mathrm{FI}(\nu_t\|\pi)dt \leq 6L_m\sqrt{\frac{C_0 d\phi(\mu)}{N}} + \frac{5\varepsilon}{6}. \tag{73}$$

To make the per-iteration complexity constant, we can choose $p = \frac{1}{b}$ so that we have $pb = O(1)$. We can find the lower bound on N by using the following inequality

$$6L_m\sqrt{\frac{C_0 d\left(1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right)}{N}} \leq 6L_m\sqrt{\frac{C_0 d\left(1 + \frac{4d}{p\mu^2 b'}\right)}{N}} \leq \frac{\varepsilon}{6}. \tag{74}$$

Plugging the chosen values for p and μ, we get

$$N \geq \frac{C_1 dL_m^2}{\varepsilon^2} + \frac{C_2(d+2)^7 L_m^6}{\varepsilon^4}, \tag{75}$$

where $C_1 = 1296\mathrm{KL}(\nu_0\|\pi) + \frac{6480\gamma}{p}$ and $C_2 = 36^2 \times 51^2$ are numerical constants. That implies

$$N = O\left(\frac{d^7 L_m^6}{\varepsilon^4}\right) \tag{76}$$

number of iteration complexity, which is also equivalent to number forward model evaluations because per-iteration complexity is fixed $pb = O(1)$. That concludes the proof of Theorem 1. □

A.1.4 Proof of Corollary 1

Recall that from Theorem 1 proof, we have inequality (66) as

$$\begin{aligned}\frac{1}{N\gamma}\int_0^{N\gamma}\mathrm{FI}(\nu_t\|\pi)dt \leq \frac{C_0}{N\gamma} + \frac{17d(d+2)L_{f_2}^2}{2b} + \frac{17}{8}\mu^2 L_{f_1}^2(d+3)^3 + 8\gamma L_m^2 d\phi(\mu) \\ + \bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2,\end{aligned} \tag{77}$$

where $C_0 = 2\mathrm{KL}(\nu_0\|\pi) + \frac{10\gamma e_0^2}{p}$. By the convexity of the Fisher information, we have

$$\begin{aligned}\mathrm{FI}(\bar{\nu}\|\pi) &\leq \frac{1}{N\gamma}\int_0^{N\gamma}\mathrm{FI}(\nu_t\|\pi)dt \\ &\leq \frac{C_0}{N\gamma} + \frac{17d(d+2)L_{f_2}^2}{2b} + \frac{17}{8}\mu^2 L_{f_1}^2(d+3)^3 + 8\gamma L_m^2 d\phi(\mu) \\ &\quad + \bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2\end{aligned} \tag{78}$$

where $\bar{\nu} \coloneqq \frac{1}{N\gamma}\int_0^{N\gamma} \nu_t dt$. By Assumption 6, we can invoke Lemma 4 and get

$$\begin{aligned}\|\bar{\nu} - \pi\|_{\mathrm{TV}} &\le 4C_{\mathrm{PI}}\mathrm{FI}(\bar{\nu}\|\pi) \\ &\le \frac{4C_0C_{\mathrm{PI}}}{N\gamma} + \frac{34d(d+2)C_{\mathrm{PI}}L_{f_2}^2}{b} + \frac{17}{2}\mu^2 C_{\mathrm{PI}}L_{f_1}^2(d+3)^3 + 32\gamma C_{\mathrm{PI}}L_m^2 d\phi(\mu) \\ &\quad + 4C_{\mathrm{PI}}(\bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2). \end{aligned} \tag{79}$$

If we let

$$\gamma = \frac{1}{2L_m}\sqrt{\frac{C_0C_{\mathrm{PI}}}{Nd(1+\left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'})}} \quad b = \left\lceil \frac{204d(d+2)C_{\mathrm{PI}}L_{f_2}^2}{\varepsilon} \right\rceil, \quad \mu = \sqrt{\frac{\varepsilon}{51C_{\mathrm{PI}}L_{f_1}^2(d+3)^3}}$$
$$p = \frac{1}{b} \quad \text{and} \quad \bar{\sigma}^2 + \bar{\varepsilon}_\sigma^2 + \bar{\alpha}_{-1}^2 \le \varepsilon/8C_{\mathrm{PI}},$$

plugging these values in (79), we get

$$\|\bar{\nu} - \pi\|_{\mathrm{TV}} \le 24L_m\sqrt{\frac{C_pC_{\mathrm{PI}}d}{N}\left(1 + \left(\frac{1-p}{p}\right)\frac{4d}{\mu^2 b'}\right)} + \frac{5\varepsilon}{6}. \tag{80}$$

We can bound the first term with $\frac{\varepsilon}{6}$ if we choose the number of iteration as

$$N \ge \frac{C_1C_{\mathrm{PI}}L_m^2 d}{\varepsilon^2} + \frac{C_2C_{\mathrm{PI}}^3L_m^6(d+3)^7}{\varepsilon^4}, \tag{81}$$

where $C_1 = 36\times 16^2 C_0$ and $C_2 = 36\times 16^2\times 24\times 34\times 51C_0$ are numerical constants. This yields a forward model complexity of

$$N = O\left(\frac{d^7L_m^6C_{\mathrm{PI}}^3}{\varepsilon^4}\right),$$

with a fixed per-iteration cost $pb = O(1)$ to achieve $\|\bar{\nu} - \pi\|_{\mathrm{TV}} \le \varepsilon$. That proves Corollary 1. □

A.1.5 Proof of Theorem 2

Given step–dependent parameters $\gamma_n, b_n, p_n, \mu_n$, define

$$\tau_n \coloneqq \sum_{k=1}^{n}\gamma_k, \qquad \bar{\nu}_{\tau_n} \coloneqq \frac{1}{\tau_n}\int_0^{\tau_n}\nu_t\,dt,$$

where ν_t denotes the law of the process $\boldsymbol{x}_t$ specified by

$$\boldsymbol{x}_t \coloneqq \boldsymbol{x}_{\tau_{n-1}} - (t-\tau_{n-1})\,\mathcal{G}(\boldsymbol{x}_{\tau_{n-1}}) + \sqrt{2}\,(\boldsymbol{B}_t - \boldsymbol{B}_{\tau_{n-1}}), \qquad t\in[\tau_{n-1},\tau_n]. \tag{82}$$

Here, γ_n denotes the step size used at iteration n, while τ_n denotes the cumulative time elapsed up to iteration n. We note that the steps up to (65) in the proof of Theorem 1 hold for $t\in[\tau_{n-1},\tau_n]$ with step-dependent parameters. Then, we can write the one-step recursion of ZO-APMC as

$$\begin{aligned}\mathrm{KL}(\nu_{\tau_n}\|\pi) - \mathrm{KL}(\nu_{\tau_{n-1}}\|\pi) &\le -\frac{1}{2}\int_{\tau_{n-1}}^{\tau_n}\mathrm{FI}(\nu_t\|\pi)dt + \frac{17\gamma_n d(d+2)L_{f_2}^2}{4b} + \frac{17}{16}\gamma_n\mu_n^2L_{f_1}^2(d+3)^3 \\ &\quad + 4\gamma_n^2 dL_m^2\phi(\mu) - \frac{\gamma_n}{p_n}\left(4 + \frac{85\gamma_n^2L_m^2\phi_n(\mu_n)}{16}\right)(e_n^2 - e_{n-1}^2) \\ &\quad + \frac{51\gamma}{4}(\sigma_n^2C^2 + \varepsilon_{\sigma_n}^2 + (\alpha_n-1)^2R_s^2) \end{aligned} \tag{83}$$

for $t\in[\tau_{n-1},\tau_n]$. We choose the parameters as follows

$$\gamma_n = \frac{1}{n}\sqrt{\frac{b'}{680L_m^2 d}}, \quad b = \lceil n^{\frac{1}{2}}\rceil, \quad p = n^{-\frac{1}{2}}, \quad \mu_n = n^{-\frac{1}{8}}. \tag{84}$$

We emphasize that the chosen step size for γ_n, namely $\gamma_n \in \big(0,\, 1/\big(85\, L_m^2 \phi_n(\mu_n)\big)\big]$, satisfies this condition for all $n \in \mathbb{N}^+$. Plugging these selected values into (83), we get

$$\mathrm{KL}(\nu_{\tau_n}\|\pi) - \mathrm{KL}(\nu_{\tau_{n-1}}\|\pi) \leq -\frac{1}{2}\int_{\tau_{n-1}}^{\tau_n} \mathrm{FI}(\nu_t\|\pi)dt + \frac{A_1}{n^{\frac{3}{2}}} + \frac{A_2}{n^{\frac{5}{4}}} + \frac{A_3}{n^2}\left(1 + \frac{4dn^{\frac{3}{4}}}{b'}\right) - c_n\left(e_n^2 - e_{n-1}^2\right) + \frac{51\gamma_n}{4}(\sigma_n^2 C^2 + \varepsilon_{\sigma_n}^2 + (\alpha_n - 1)^2 R_s^2), \quad (85)$$

where

$$A_1 = \frac{17d(d+2)L_{f_2}^2}{4}\sqrt{\frac{b'}{680L_m^2 d}}, \quad A_2 = \frac{17L_{f_1}^2(d+3)^3}{16}\sqrt{\frac{b'}{680L_m^2 d}}, \quad A_3 = 4dL_m^2\sqrt{\frac{b'}{680L_m^2 d}}, \quad (86)$$

and all of them are constant. Moreover, we define

$$c_n \coloneqq \frac{\gamma_n}{p_n}\left(4 + \frac{85\gamma_n^2 L_m^2 \phi_n(\mu_n)}{16}\right)\left(e_n^2 - e_{n-1}^2\right), \quad (87)$$

for $n \in \mathbb{N}^+$. Iterating the bound in (85), we obtain

$$\mathrm{KL}(\nu_{\tau_n}\|\pi) \leq \mathrm{KL}(\nu_0\|\pi) - \frac{1}{2}\int_0^{\tau_n} \mathrm{FI}(\nu_t\|\pi)dt + A_1S_1 + A_2S_2 + A_3S_3 - \sum_{k=1}^{n} c_k\left(e_k^2 - e_{k-1}^2\right) + \frac{51}{4}\sum_{k=1}^{n}\gamma_k(\sigma_k^2 C^2 + \varepsilon_{\sigma_k}^2 + (\alpha_k - 1)^2 R_s^2), \quad (88)$$

where we have

$$S_1 = \sum_{k=1}^{n} k^{-3/2} \leq \sum_{k=1}^{\infty} k^{-3/2} < \infty, \quad S_2 = \sum_{k=1}^{n} k^{-5/4} \leq \sum_{k=1}^{\infty} k^{-5/4} < \infty.$$

$$S_3 = \sum_{k=1}^{n}\Big(k^{-2} + \frac{4dk^{-5/4}}{b'}\Big) \leq \sum_{k=1}^{\infty}\Big(k^{-2} + \frac{4dk^{-5/4}}{b'}\Big) < \infty.$$

Thus, A_1S_1, A_2S_2, and A_3S_3 are bounded constants and are independent of n. Furthermore, if we assume that c_n is nonnegative and nonincreasing sequence (i.e. $0 \leq c_{n+1} \leq c_n$), we can bound the summation of difference of the error terms in (88) as follows

$$-\sum_{k=1}^{n} c_k\left(e_k^2 - e_{k-1}^2\right) = c_1e_0^2 + \sum_{k=1}^{n_1}(c_{k+1} - c_k)e_k^2 - c_ne_n^2 \leq c_1e_0. \quad (89)$$

Thus, we need to prove that c_n is nonnegative and nonincreasing sequence. Plugging the variables in (84) into the definition of c_n, we get

$$c_n = \frac{1}{n^{1/2}}\sqrt{\frac{b'}{680L_m^2 d}}\left[4 + \frac{85}{12n^2}\frac{b'}{680d}\right] + \frac{1}{n^{9/4}}\left(1 - \frac{1}{n^{1/2}}\right)\sqrt{\frac{b'}{680L_m^2 d}}\frac{85b'}{12 \times 680d}. \quad (90)$$

Therefore, $c_n \geq 0$ for any $n \in \mathbb{N}^+$ and we have $c_n = O(n^{-1/2})$ so it's a nonnincreasing sequence. Using this upper bound on the differences of error terms in (85), we get

$$\mathrm{KL}(\nu_{\tau_n}\|\pi) \leq \mathrm{KL}(\nu_0\|\pi) - \frac{1}{2}\int_0^{\tau_n} \mathrm{FI}(\nu_t\|\pi)dt + A_1S_1 + A_2S_2 + A_3S_3 + c_1e_0^2 + \frac{51}{4}\sum_{k=1}^{n}\gamma_k(\sigma_k^2 C^2 + \varepsilon_{\sigma_k}^2 + (\alpha_k - 1)^2 R_s^2) \quad (91)$$

We need to upper bound the last term related to the annealing parameters and the score network. Note that $\{\alpha_k\}_{k=0}^{N-1}$ is a decreasing sequence with $\alpha_{K,\ldots,N-1} = 1$ for some K. Hence, we have

$$\frac{51}{4}\sum_{k=1}^{n}\gamma_k(\alpha_k - 1)^2 R_s^2 \leq \frac{51}{4}\sum_{k=1}^{\infty}\gamma_k(\alpha_k - 1)^2 R_s^2 = \frac{51}{4}\sum_{k=1}^{K}\gamma_k(\alpha_k - 1)^2 R_s^2 = C_\alpha, \quad (92)$$

where C_α is bounded constant. Additionally, we can choose a schedule for σ_n such that $\sigma_n^2 = O(n^{-\beta})$ for $\beta > 0$. For example, in practice, one often uses an exponentially decaying schedule $\sigma_n = \sigma_0\, \xi^{\,n}$ (Sun et al., 2024; Song & Ermon, 2019), where ξ denotes the decay rate. Then, we get

$$\frac{51}{4}\sum_{k=1}^{n}\gamma_n\sigma_k^2C^2 \le \frac{51}{4}\sum_{k=1}^{\infty}\gamma_n\sigma_k^2C^2 = C_\sigma < \infty. \tag{93}$$

Lastly, if we assume that $\varepsilon_{\sigma_n}^2 = O(n^{-\beta})$ for some $\beta > 0$, then

$$\frac{51}{4}\sum_{k=1}^{n}\gamma_n\varepsilon_{\sigma_k}^2 \;\le\; \frac{51}{4}\sum_{k=1}^{\infty}\gamma_n\varepsilon_{\sigma_k}^2 \;=\; C_{\varepsilon_\sigma} \;<\; \infty. \tag{94}$$

This assumption implies that, as the noise level decreases, the test error of the score network forms a monotonically decreasing sequence of order $O(n^{-\beta})$ with $\beta > 0$. Such a decay pattern is commonly observed in diffusion models, where networks trained across noise scales exhibit progressively lower errors at finer (less noisy) levels (Song et al., 2020; Ho et al., 2020). Combining (92), (93), and (94) in (91), we get

$$\begin{aligned}\mathrm{KL}(\nu_{\tau_n}\|\pi) \le \mathrm{KL}(\nu_0\|\pi) - \frac{1}{2}\int_0^{\tau_n}\mathrm{FI}(\nu_t\|\pi)dt + A_1S_1 + A_2S_2 + A_3S_3 + c_1e_0^2 \\ + C_\alpha + C_{\varepsilon_\sigma} + C_\sigma\end{aligned} \tag{95}$$

Rearranging the terms, using the convexity of Fisher information and dividing both sides by τ_n, we get

$$\begin{aligned}\mathrm{FI}(\bar{\nu}_{\tau_n}\|\pi) &\le \frac{1}{\tau_n}\int_0^{\tau_n}\mathrm{FI}(\nu_t\|\pi)\,dt \\ &\le \frac{2\,\mathrm{KL}(\nu_0\|\pi)}{\tau_n} + \frac{2}{\tau_n}\Big(A_1S_1 + A_2S_2 + A_3S_3 + c_1e_0^2 + C_\alpha + C_{\varepsilon_\sigma} + C_\sigma\Big)\end{aligned} \tag{96}$$

where $A_1S_1 + A_2S_2 + A_3S_3 + c_1e_0^2 + C_\alpha + C_{\varepsilon_\sigma} + C_\sigma < \infty$. Alternatively, if $t \in [\tau_n, \tau_{n+1}]$, integrating (64) between τ_n and t and dropping the negative integral over the Fisher information give us

$$\begin{aligned}\mathrm{KL}(\nu_t\|\pi) &\le \mathrm{KL}(\tau_n\|\pi) + \frac{17(t-\tau_n)d(d+2)L_{f_2}^2}{4b} + \frac{17}{16}(t-\tau_n)\mu^2L_{f_1}^2(d+3)^3 \\ &\quad + 4(t-\tau_n)\gamma dL_m^2\phi(\mu) - \frac{4(t-\tau_n)}{p}\left(1 + \frac{85\gamma^2L_m^2\phi(\mu)}{16}\right)e_{n+1}^2 \\ &\quad + \frac{17(t-\tau_n)}{4}(\sigma_kC + \varepsilon_{\sigma_k} + (\alpha_k-1)R_s)^2 \qquad (97) \\ &\le \mathrm{KL}(\nu_0\|\pi) + 2A_1S_1 + 2A_2S_2 + 2A_3S_3 + 2C_\alpha + 2C_{\varepsilon_\sigma} + 2C_\sigma + c_1e_0^2. \qquad (98)\end{aligned}$$

where, in the second inequality, for positive terms, we use the fact that $\frac{1}{n+1} \le \sum_{k=1}^{n}\frac{1}{k^\beta}$ for $\beta > 1$. With (98), we show that $\{\mathrm{KL}(\nu_t\|\pi)\}_{t\ge0}$ is bounded. By the convexity of the KL divergence, this implies that the sequence $\{\mathrm{KL}(\bar{\mu}_{\tau_n}\|\pi)\}_{n\in\mathbb{N}}$ is uniformly bounded as well. Since the sublevel sets of $\mathrm{KL}(\cdot\|\pi)$ are weakly compact, $(\bar{\nu}_{\tau_n})_{n\in\mathbb{N}}$ is tight. To establish that $\bar{\nu}_{\tau_n} \rightharpoonup \pi$ weakly, it suffices to verify that every cluster point of $(\bar{\nu}_{\tau_n})_{n\in\mathbb{N}}$ equal to π.

Take a subsequence $(\bar{\nu}_{\tau_n})_{n\in\mathbb{N}}$ converging to some limit $\bar{\nu}$. Sending $n \to \infty$ in (96) and using $\tau_n \to \infty$ gives $\mathrm{FI}(\bar{\nu}_{\tau_n}\|\pi) \to 0$, so the same holds along the subsequence. By the weak lower semicontinuity of the Fisher information along the subsequence, we have $\mathrm{FI}(\bar{\nu}\|\pi) = 0$. Writing $\psi := \frac{d\bar{\nu}}{d\pi}$, this means $\sqrt{\psi} \in \mathrm{dom}\,\mathcal{E}$ and $\mathcal{E}(\sqrt{\psi}) = 0$. Since $\nabla \log \pi$ is Lipschitz by Assumption 1, π has a continuous and strictly positive density on $\mathbb{R}^d$, so $\mathcal{E}(\sqrt{\psi}) = 0$, which implies that ψ must be a constant π-a.e., hence $\bar{\nu} = \pi$. □

A.2 Extended Experimental Results

A.2.1 Numerical Validation

In this section, we give more details on our numerical validation experiments. Similar to the related works (Sun et al., 2024; Song & Ermon, 2020), we run ZO-APMC with an exponential annealing

schedule:

$$\sigma_k := \max\{\sigma_0 r^k, \sigma_{\min}\}, \quad \alpha_k = \max\{\alpha_0 \sigma_k^2, 1\}, \tag{99}$$

where r is the decay factor and k is the step number. We always choose $\alpha_0 \leq 1/\sigma_{\min}^2$ so that $\{\alpha_k\}$ converges to 1. For our numerical validation results, we set $\sigma_0 = 10$, $\alpha_0 = 10$, $r = 0.975$, $\sigma_{\min} = 0$ and $\gamma = 0.1$. For ZO estimator, we choose the smoothing parameter as $\mu = 10^{-4}$. We run ZO-APMC with 1000 sample points initialized with uniform distribution $\mathrm{U}[-50, 50]^2$ on $[-50, 50]^2$ grid for $N = 2000$ iterations. At each step, we use a Gaussian mixture model (GMM) to fit a distribution to the samples at intermediate steps, which allows us to compute the probability of an arbitrary value on $[-50, 50]^2$ grid. Then, for each intermediate distribution, we calculate the empirical Fisher information relative to target posterior whose analytical posterior can be calculated. We discretize the grid to 1000×1000 unit areas in $[-50, 50]^2$ and calculate the Fisher information for each unit area. The total sum over the grid gives us the approximate relative Fisher information. We perform additional experimental results for showing the effect of p on the convergence with mean and standard deviation values in Fig. 6.

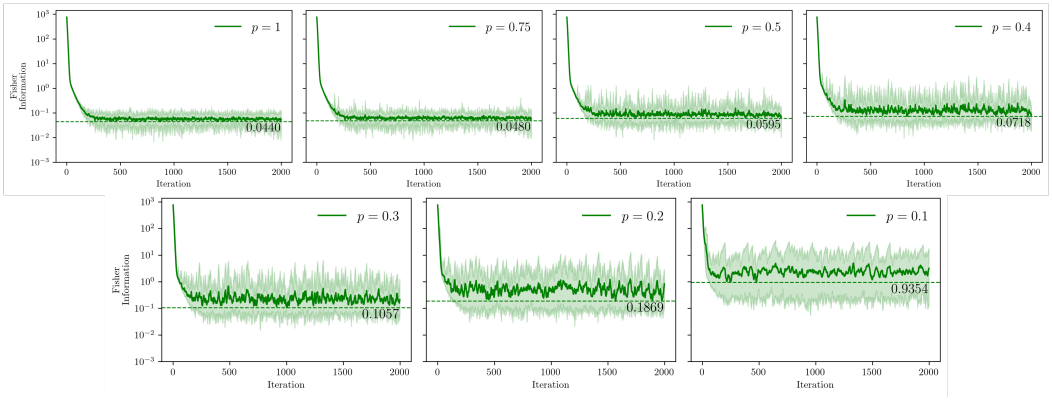

Figure 6: Effect of p on the convergence of ZO-APMC to the true posterior distribution in terms of relative Fisher information. The solid lines show the mean values and shaded areas show the standard deviations.

At each iteration, ZO-APMC performs a zero-order estimate with a large batch size $b = 10$ with probability p, while with probability $1-p$ it uses a smaller batch size $b' = 1$, whose gradient estimate is aggregated with the previous step's update.

A.2.2 Statistical Validation

The SGM used in this experiment is U-Net taken from (Nichol & Dhariwal, 2021) with some of its layers removed to process the 32×32 images, which are taken from CelebA (Liu et al., 2018) dataset. Each image is normalized to $[-1, 1]$ and downscaled to 32×32 pixels for simplicity. The forward operator is generated as random Gaussian matrix and for each test image, we inject a Gaussian noise with variance 0.01 as a measurement noise. We construct a bimodal distribution by selecting male and female images from the CelebA dataset and fitting a Gaussian mixture model (GMM) to the combined data. To ensure adequate separation, the two modes are shifted by $+1$ and -1. The SGM prior is then trained on samples drawn from this synthetic multimodal distribution. Because the synthetic Gaussian images lack the structural richness of natural images, the score network's results on this dataset should not be taken as representative of its performance on real-world data. For comparison, we compute the target modes and posterior statistics using the statistics derived from male and female images in the CelebA dataset.

A.3 Additional Details for Inverse Problem Experiments

A.3.1 MRI Experiment Details

We evaluate the reconstruction quality of the samples generated by ZO-APMC and other baselines methods by using peak signal to noise (PSNR) ratio, structural similarity index measure (SSIM),

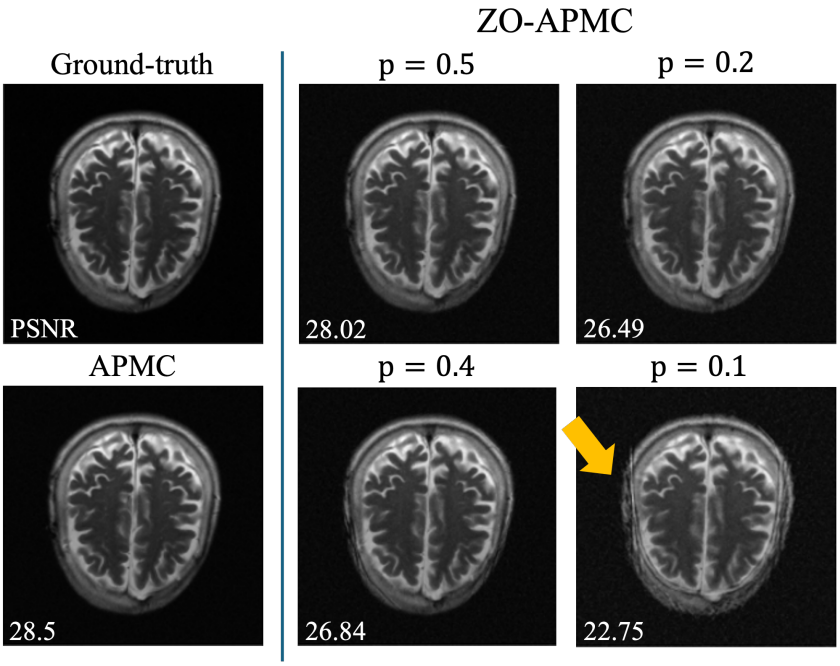

Figure 7: Comparison of the ground-truth brain MRI with APMC and ZO-APMC reconstructions for various probabilities $p \in \{0.1, 0.2, 0.4, 0.5\}$, using a large batch size of $b = 10^4$ and a small batch size of $b' = 10^3$. PSNR values for each reconstruction are displayed in the lower-left corner of the corresponding image.

normalized root mean square (NRMSE), and standard deviation (SD). Given an estimate $\hat{\boldsymbol{x}} \in \mathbb{R}^d$ and the ground truth $\boldsymbol{x}_{\text{GT}} \in \mathbb{R}^d$, we define the error metrics as

$$\text{MSE}(\hat{\boldsymbol{x}}, \boldsymbol{x}_{\text{GT}}) = \frac{1}{d}\|\hat{\boldsymbol{x}} - \boldsymbol{x}_{\text{GT}}\|_2^2, \quad \text{NRMSE}(\hat{\boldsymbol{x}}, \boldsymbol{x}_{\text{GT}}) = \frac{\|\hat{\boldsymbol{x}} - \boldsymbol{x}_{\text{GT}}\|_2}{\|\boldsymbol{x}_{\text{GT}}\|_2},$$

$$\text{PSNR}(\hat{\boldsymbol{x}}, \boldsymbol{x}_{\text{GT}}) = 10 \log_{10}\left(\frac{\max(\boldsymbol{x}_{\text{GT}})^2}{\text{MSE}(\hat{\boldsymbol{x}}, \boldsymbol{x}_{\text{GT}})}\right).$$

where d is the number of elements in $\boldsymbol{x}$, and $\max$ denotes the maximum possible value of the signal (e.g., 1 for normalized data or 255 for 8-bit images).

Ablation Study. Among the inverse problems considered in this work, MRI reconstruction involves the largest image size (256×256), which necessitates a larger batch size in our ZO estimator to accurately compute the forward-model gradient. To identify the optimal value of p, we subsample examples from the validation set of FastMRI (Zbontar et al., 1811) and evaluate reconstruction quality across different values, $p \in \{0.1, 0.2, 0.4, 0.5\}$, as illustrated in Fig. 7. As indicated by the orange arrow, reducing p excessively while keeping b fixed produces visible artifacts in the generated samples. ZO-APMC maintains reasonable reconstruction quality down to $p = 0.2$. Even when using a smaller batch size of $b' = 10^3$, an order of magnitude lower than b, in about half of the iterations on average, ZO-APMC maintains high reconstruction quality that is very close both visually and quantitatively to the reconstruction of APMC. As opposed to APMC, ZO-APMC achieves this without any gradient information and uses only forward model function evaluations. Because the performance gain beyond $p = 0.2$ is not significant and the gap between $p = 0.5$ and $p = 0.2$ can be further reduced by averaging multiple parallel outputs, we set $p = 0.2$ for our brain MRI inverse problem experiments.

Moreover, ZO estimators are widely recognized in the literature for exhibiting high variance in high-dimensional settings, as they rely on first-order approximations of the function along random directions. To evaluate our proposed variance-reduction mechanism, we compare the reconstructions of our method with DPS and APMC, which do not assume black-box setting and have access to

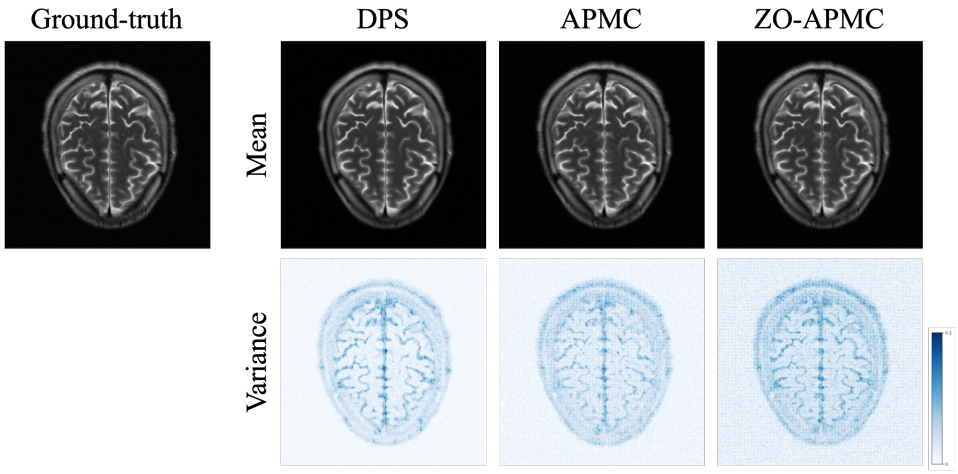

Figure 8: Comparison of the ground-truth brain MRI with reconstructions from ZO-APMC and the gradient-based approaches DPS and APMC. Each method generates 20 samples from the same measurements; the first row shows the mean reconstructions and the second row shows the corresponding variance maps. Owing to its variance-reduction mechanism, ZO-APMC produces variance maps comparable to those of the gradient-based algorithms despite relying on noisy evaluations of the forward model.

gradients of the forward model. Results in Fig. 8 show that although our proposed method ZO-APMC assumes black-box setting and uses noisy forward model evaluations to approximate the gradient of the forward model, it has similar variance compared to DPS and APMC, which assumes access to the gradients, thanks to our proposed variance-reduction mechanism.

A.3.2 Black-hole Imaging Experiment Details

In black-hole imaging, very long baseline interferometry (VLBI) uses an array of ground-based telescopes. Each telescope pair (a, b) at time t produces a complex visibility $V_t^{a,b}$. To mitigate atmospheric and thermal phase errors, visibilities are combined into noise-robust *closure* measurements (Chael et al., 2018): closure phases $\boldsymbol{y}^{\text{cph}}_{t,(a,b,c)}$ and log-closure amplitudes $\boldsymbol{y}^{\text{camp}}_{t,(a,b,c,d)}$. Following Sun & Bouman (2021); Zheng et al. (2024), we use the following likelihood model:

$$\ell(\boldsymbol{y} \mid \boldsymbol{x}) = \sum_t \frac{\left\|\mathcal{A}_t^{\text{cph}}(\boldsymbol{x}) - \boldsymbol{y}_t^{\text{cph}}\right\|_2^2}{2\beta_{\text{cph}}^2} + \sum_t \frac{\left\|\mathcal{A}_t^{\text{camp}}(\boldsymbol{x}) - \boldsymbol{y}_t^{\text{camp}}\right\|_2^2}{2\beta_{\text{camp}}^2} + \frac{\rho}{2}\left\|\sum_i x_i - y^{\text{flux}}\right\|_2^2. \quad (100)$$

Here, $\mathcal{A}_t^{\text{cph}}$ and $\mathcal{A}_t^{\text{camp}}$ map an image $\boldsymbol{x}$ to predicted closure phases and log-closure amplitudes, respectively; β_{cph} and β_{camp} are instrument-specific noise scales. The first two sums act as chi-squared penalties for the closure measurements, while the final term enforces the total-flux constraint with weight ρ and target flux y^{flux}. For our experiments, we adopted the dataset, pre-trained SGM prior, forward model implementation, and baseline methods provided by Zheng et al. (2025). For EnKG, we adopt the hyperparameters recommended by Zheng et al. (2024), and for the baseline methods we use the hyperparameters provided by Zheng et al. (2025).

A.3.3 Navier-Stokes Equation Experiment Details

In our experiments, we study the two-dimensional Navier–Stokes equations for a viscous, incompressible fluid in vorticity form on a torus. Let $u \in C([0,T]; H^r_{\text{per}}((0,2\pi)^2, \mathbb{R}^2))$ for any $r > 0$ denote the velocity field, and let $w = \nabla \times u$ be the vorticity. The initial vorticity is $w_0 \in L^2_{\text{per}}((0,2\pi)^2; \mathbb{R})$, the viscosity coefficient is $\nu \in \mathbb{R}_+$, and the forcing term is $f \in L^2_{\text{per}}((0,2\pi)^2; \mathbb{R})$. The solution operator $\mathcal{G}$ maps the initial vorticity to the vorticity at time T, i.e. $\mathcal{G} : w_0 \mapsto w_T$. In

our experiments, we implement $\mathcal{G}$ using a pseudo-spectral solver following He & Sun (2007):

$$\partial_t w(x,t) + u(x,t) \cdot \nabla w(x,t) = \nu \Delta w(x,t) + f(x), \qquad x \in (0,2\pi)^2,\ t \in (0,T], \tag{101}$$

$$\nabla \cdot u(x,t) = 0, \qquad x \in (0,2\pi)^2,\ t \in (0,T], \tag{102}$$

$$w(x,0) = w_0(x), \qquad x \in (0,2\pi)^2. \tag{103}$$

The task is to infer the initial vorticity field from noisy and sparsely observed vorticity data at time $T = 1$. Since Eq. (21) admits no closed-form solution, the corresponding derivative of the solution operator is also unavailable. Furthermore, the computation of accurate numerical derivatives via automatic differentiation through the solve is challenging since the extensive computation graph can span thousands of discrete time steps.

We follow the approach in Zheng et al. (2024; 2025) and first solve the equation up to time $T = 5$ starting from random Gaussian initial conditions, which are highly nontrivial due to the nonlinearity of the Navier–Stokes equations. We use the SGM-prior, which was pre-trained over $20{,}000$ vorticity fields, and use the test set consisting of 10 samples from InverseBench. For EnKG, we use the hyperparameters recommended by Zheng et al. (2024), and for the baseline methods we adopt the hyperparameters provided by Zheng et al. (2025). Quantitative results are presented in Fig. 4. Our method demonstrates a performance comparable to most black-box posterior samplers, while distinctively providing rigorous theoretical guarantees of convergence to the target posterior guarantees, which is not established for the baseline methods.

Table 4: Quantitative results for the Navier–Stokes equation benchmark. For each case, the best-performing method is shown in **bold**. Baseline results are taken from Zheng et al. (2024).

	NRMSE ($\sigma_{\text{noise}} = 0$)↓	NRMSE ($\sigma_{\text{noise}} = 1$)↓	NRMSE ($\sigma_{\text{noise}} = 2$)↓
Forward-GSG	1.687	1.612	1.454
Central-GSG	2.203	2.117	1.746
SCG	0.908	0.928	0.966
DPG	0.325	0.408	0.466
EnKG	**0.120**	**0.191**	**0.294**
ZO-APMC (Ours)	0.459	0.463	0.472