LATENT-TIME STEERING FOR BAYESIAN OPTIMAL EXPERIMENTAL DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Bayesian optimal experimental design (BOED) selects experiments to maximally reduce uncertainty about unknown physical parameters, but the expected information gain (EIG) objective commonly used in BOED is often computationally expensive and difficult to estimate under complex priors. In this work, we propose a steering-cost surrogate for EIG that leverages pretrained diffusion-based generative priors over parameters and estimates EIG as a latent-time *steering cost* induced by likelihood guidance. Specifically, we combine the MINDE score–KL identity with a diffusion posterior sampling (DPS) plug-in approximation, yielding a tractable integral of squared likelihood guidance. The same pretrained prior supports posterior inference and design scoring without training a separate design policy. Additionally, we characterize the approximation error introduced by the proposed estimator. We validate our method on a range of design problems, including source localization, constant elasticity of substitution, and Darcy flow parameter estimation. For Darcy source-position design on a 64×64 grid, our method reduces mean relative L_2 reconstruction error by 24.7% and 24.9% after the first and second measurements, respectively, compared with ALINE at the same resolution over 20 test cases.

1 INTRODUCTION

The goal of Bayesian optimal experimental design (BOED) is to choose particular experiments that will provide maximal information about an unknown physical system. In many scientific applications—e.g., sensor placement, PDE-constrained inverse problems, and laboratory automation—measurements are costly, and design decisions must be made sequentially under tight budgets. The canonical BOED objective is the expected information gain (EIG), which selects designs that, in expectation, maximally reduce posterior model uncertainty. However, in modern simulator-based settings, the EIG is often computationally prohibitive (Kleinegesse & Gutmann, 2020; Foster et al., 2021; Ivanova et al., 2021; Huang et al., 2025). Evaluating a single candidate design may require many simulator calls, and accurate estimation typically involves nested Monte Carlo or density-ratio estimation under non-Gaussian, high-dimensional posteriors. This computational burden is a central bottleneck that limits the use of BOED in complex physical system inference problems.

Pretrained generative models suggest a different route to EIG estimation. Generative models (e.g. diffusion(Song et al., 2021a;b; Ho et al., 2020)) have been shown recently to be effective *priors* for complicated parameter distributions, enabling high-quality sampling and inference in settings where classical priors are inadequate. Importantly, these generative priors support *likelihood-guided* sampling procedures, e.g. Diffusion Posterior Sampling (DPS)-style guidance (Chung et al., 2023), where the posterior is approximated by steering an unconditional sampler using gradients of the likelihood. This perspective suggests an opportunity. Namely, if posterior sampling is implemented by guidance along latent time, then the *amount of steering* may itself quantify how informative a measurement is. The same pretrained diffusion prior supports both posterior inference and candidate scoring throughout sequential design, without training a separate design policy. This is particularly useful when a suitable prior has already been learned from expensive simulations and high-dimensional posterior inference is already required.

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

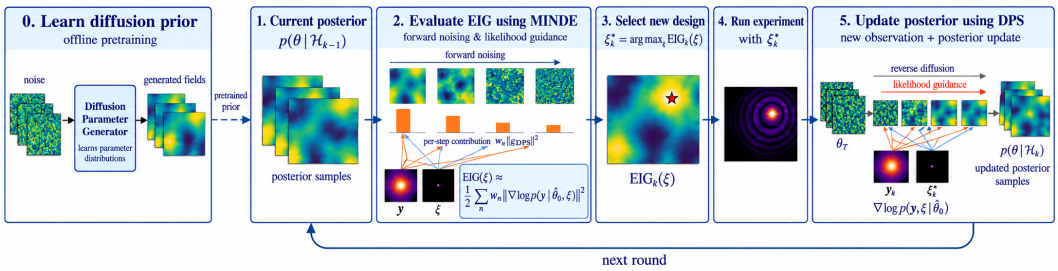

Figure 1: **Overview of sequential Bayesian experimental design with diffusion models. (0)** An unconditional diffusion model is trained offline to represent the prior over parameter fields. At round k, **(1)** samples from the current posterior $p(\theta \mid \mathcal{H}_{k-1})$ are used to **(2)** estimate the expected information gain $\mathrm{EIG}_k(\xi)$ for candidate designs using our MINDE-based diffusion-path estimator. **(3)** The design $\xi_k^\star = \arg\max_\xi \mathrm{EIG}_k(\xi)$ is selected, and **(4)** an experiment at $\xi_k^\star$ yields observation y_k. **(5)** Using all acquired observations, likelihood-guided DPS approximately samples the updated posterior $p(\theta \mid \mathcal{H}_k)$ for the next design round.

We propose a new EIG estimator based on the *steering cost* of likelihood-guided sampling under a generative prior. Our starting point is an unconditional diffusion model learned offline over the unknown parameters. Specifically, we connect this intuition to diffusion-path identities underlying Mutual Information Neural Diffusion Estimation (MINDE), which express KL divergence in terms of an integral of squared score differences along diffusion time. We combine this identity with likelihood-guided diffusion sampling, in the spirit of Diffusion Posterior Sampling (DPS) and related methods, and construct the required score differences using likelihood-score approximations available from the simulator. This workflow yields a practical estimator in which the information provided by an experiment is measured by the cumulative control energy required to steer prior trajectories toward the corresponding posterior. The overall workflow, from offline prior learning to EIG evaluation and posterior updating, is summarized schematically in Figure 1.

Recent work has also explored generative models for BOED from complementary perspectives. CoDiff (Iollo et al., 2025) uses pooled posterior sampling for gradient-based design optimization. DiffBED uses diffusion models to address optimization over high-dimensional design spaces (Saravanan et al., 2026). Our approach instead uses a diffusion prior over unknown parameters and constructs an information-gain surrogate from the likelihood guidance used for posterior inference. Specifically, we combine likelihood-guided posterior sampling with diffusion-path identities to estimate EIG through the steering energy induced by an observation. This formulation is particularly motivated by high-dimensional parameter inference, where explicitly estimating posterior densities or density ratios for each candidate design can be prohibitively expensive. We provide a detailed comparison with these approaches and other related methods in Appendix A.

Our main contributions in this paper are:

- **A diffusion-based framework for sequential BOED.** We develop a sequential design framework that reuses a pretrained diffusion prior for posterior sampling and candidate evaluation.
- **Steering-cost surrogate for expected information gain.** We derive a tractable EIG surrogate by connecting the per-observation KL term in BOED to a *latent-time steering cost* induced by likelihood guidance in diffusion posterior sampling and characterize the approximation error in terms of DPS guidance error.
- **High-dimensional sequential design.** We demonstrate sequential Darcy design with a fixed diffusion prior, achieving approximately 25% lower mean relative L_2 error than ALINE (64×64) in source-position design.

2 Problem Formulation

2.1 BOED

We consider sequential Bayesian optimal experimental design (BOED) applied to the solution of inverse problems in which we wish to infer unknown parameters $\boldsymbol{\theta} \in \mathbb{R}^D$ of a (differentiable) simulator $\mathcal{F}$ that models noisy measurements from a physical system. At BOED step $k \in \{1, \ldots, K\}$, a

design $\boldsymbol{\xi}_k \in \Xi$ (e.g., sensor locations or source terms) is selected, and an observation is drawn from $p(\mathbf{y}_k \mid \boldsymbol{\theta}, \boldsymbol{\xi}_k)$. For additive Gaussian observations, this specializes to

$$\mathbf{y}_k = \mathcal{F}(\boldsymbol{\theta}, \boldsymbol{\xi}_k) + \mathbf{n}_k, \qquad \mathbf{n}_k \sim \mathcal{N}(0, \boldsymbol{\Sigma}_y). \tag{1}$$

Let the measurement history up to step k be $H_k := \{(\mathbf{y}_1, \boldsymbol{\xi}_1), \ldots, (\mathbf{y}_k, \boldsymbol{\xi}_k)\}$. The current belief before selecting $\boldsymbol{\xi}_k$ is the posterior $p(\boldsymbol{\theta} \mid H_{k-1})$. A commonly adopted utility is the one-step EIG,

$$\mathrm{IG}_k = D_{\mathrm{KL}}\Big(p(\boldsymbol{\theta} \mid H_k) \,\Big\|\, p(\boldsymbol{\theta} \mid H_{k-1})\Big), \qquad \mathrm{EIG}_k(\boldsymbol{\xi}) = \mathbb{E}_{\mathbf{y} \sim p(\cdot \mid H_{k-1}, \boldsymbol{\xi})}[\mathrm{IG}_k]\,. \tag{2}$$

The next design is then selected as $\boldsymbol{\xi}_k \in \arg\max_{\boldsymbol{\xi} \in \Xi} \mathrm{EIG}_k(\boldsymbol{\xi})$.

Estimating EIG is challenging for complex simulators $\mathcal{F}$ as it requires comparing $p(\boldsymbol{\theta} \mid H_{k-1})$ to the updated posterior $p(\boldsymbol{\theta} \mid H_{k-1}, \mathbf{y}, \boldsymbol{\xi})$, which is often non-Gaussian and expensive to sample from or normalize. These challenges make EIG evaluation computationally demanding in high-dimensional, simulator-based BOED settings (Dong et al., 2025; Shen et al., 2025).

2.2 Efficient Steering-Cost Estimation

To make EIG practical for high-dimensional BOED problems, we need an accurate yet computationally efficient estimator of the KL term in (2). We estimate this quantity by measuring how much the current belief must be *steered* to obtain the updated belief. Our aim is to avoid evaluating normalization constants or fitting a new conditional density model after each observation, while reducing the cost of repeated inference for candidate designs. These requirements are especially stringent in sequential BOED where candidate designs must be scored many times, while the belief itself changes after every acquired observation. We here propose to integrate two complementary tools. MINDE (Franzese et al., 2024) converts the endpoint KL in EIG into a latent-time integral of squared score differences. DPS (Chung et al., 2023) then provides a denoiser-based likelihood-score plug-in that makes this integral computable for simulator-based models. Their integration yields a reusable steering-cost surrogate for scoring candidate designs throughout the sequential BOED loop.

Diffusion prior. We assume access to a pretrained *unconditional* diffusion model representing a prior distribution $p_{\mathrm{prior}}(\boldsymbol{\theta})$ over parameters $\boldsymbol{\theta}$. We use the VE-SDE (Song et al., 2021b) representation compatible with the EDM parameterization:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_0 + \sigma(t)\boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \qquad t \in [0, T]. \tag{3}$$

Assuming that $\sigma(0) = 0$ and $\sigma(\cdot)$ is nondecreasing and absolutely continuous, (3) has the same time-t marginals as the driftless SDE

$$\mathrm{d}\boldsymbol{\theta}_t = a(t)\,\mathrm{d}W_t, \qquad a(t)^2 = \frac{\mathrm{d}}{\mathrm{d}t}\big[\sigma(t)^2\big] = 2\sigma(t)\dot{\sigma}(t), \tag{4}$$

where W_t is standard Brownian motion.

MINDE. MINDE (Franzese et al., 2024) provides the identity underlying our steering-cost estimator. For two initial distributions p_0 and q_0, we express their endpoint KL divergence as a latent-time integral of squared score differences along the forward diffusion path.

Proposition 1 (MINDE Score–KL identity)**.** *(Eq. (10) in Franzese et al., 2024) Let $T < \infty$ and $K_{0,t}(\theta_0, \cdot)$ be the Markov transition kernel of the forward diffusion, i.e. $K_{0,t}(\theta_0, A) = \mathbb{P}(\Theta_t \in A \mid \Theta_0 = \theta_0)$. Given two initial distributions p_0 and q_0, define the time-t marginals*

$$p_t := p_0 K_{0,t}, \qquad q_t := q_0 K_{0,t}.$$

Assume p_t and q_t admit densities and that the scores $\nabla_\theta \log p_t(\theta)$ and $\nabla_\theta \log q_t(\theta)$ exist and are square-integrable. Then

$$D_{\mathrm{KL}}(q_0 \| p_0) = D_{\mathrm{KL}}(q_T \| p_T) + \frac{1}{2}\int_0^T a(t)^2\, \mathbb{E}_{\boldsymbol{\theta} \sim q_t}\Big[\|\nabla_{\boldsymbol{\theta}} \log q_t(\boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}} \log p_t(\boldsymbol{\theta})\|_2^2\Big]\,\mathrm{d}t. \tag{5}$$

The forward semigroup contracts KL divergence. If the noising process additionally satisfies $D_{\mathrm{KL}}(q_T \| p_T) \to 0$ as $T \to \infty$, the terminal residual becomes negligible at sufficiently large noise. Thus, for sufficiently large terminal noise/time,

$$D_{\mathrm{KL}}(q_0 \| p_0) \;\approx\; \int_0^T \frac{a(t)^2}{2}\, \mathbb{E}_{q_t}\Big[\big\|\nabla \log q_t - \nabla \log p_t\big\|_2^2\Big]\, dt,$$

where $a(t)$ is the diffusion coefficient defined in (4).

For BOED, we will specialize p_0 to the current belief and q_0 to the belief updated with a prospective observation. MINDE therefore gives the desired pathwise representation of the per-observation information gain. For a candidate design $\boldsymbol{\xi}$ and prospective observation $\mathbf{y}$, we specialize the MINDE identity by setting $p_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid H_{k-1})$, $q_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid H_{k-1}, \mathbf{y}, \boldsymbol{\xi})$.

Let p_t and q_t denote their corresponding forward-noised marginals. By Bayes' rule,

$$\nabla_{\boldsymbol{\theta}_t} \log q_t(\boldsymbol{\theta}_t) - \nabla_{\boldsymbol{\theta}_t} \log p_t(\boldsymbol{\theta}_t) = \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, H_{k-1}, \boldsymbol{\xi}).$$

Thus, the squared score difference in MINDE is the squared likelihood-induced steering term for the prospective observation.

At each step, we assume $p(\mathbf{y}_i \mid \boldsymbol{\theta}, H_{i-1}, \boldsymbol{\xi}_i) = p(\mathbf{y}_i \mid \boldsymbol{\theta}, \boldsymbol{\xi}_i)$. The posterior given H_k satisfies

$$p(\boldsymbol{\theta} \mid H_k) \propto p_{\mathrm{prior}}(\boldsymbol{\theta}) \prod_{i=1}^{k} p(\mathbf{y}_i \mid \boldsymbol{\theta}, \boldsymbol{\xi}_i).$$

This factorization expresses the full-history log-likelihood as a sum of individual log-likelihood terms. We use these terms to guide the latent-time sampler associated with the pretrained unconditional diffusion prior, obtaining approximate posterior samples without training a new conditional generative model for each H_k.

DPS Plug-in for the Steering Cost. The MINDE specialization identifies the steering cost, but the noised likelihood score $\nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})$ is generally intractable. We use DPS (Chung et al., 2023) to approximate this score by evaluating the likelihood at a denoised prediction and backpropagating to the latent state. For notational simplicity, we suppress the dependence on H_{k-1} in the DPS expressions below.

At reverse step t with latent $\boldsymbol{\theta}_t$, the diffusion model provides a denoised (clean) prediction

$$\hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t) \;\approx\; \boldsymbol{\theta}_0, \tag{6}$$

obtained from the score/denoiser network (by Tweedie's formula). DPS then evaluates the likelihood at the denoised prediction and backpropagates the resulting gradient to the current latent state, yielding the plug-in approximation

$$\nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}) \;\approx\; \nabla_{\boldsymbol{\theta}_t} \log p\Big(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t), \boldsymbol{\xi}\Big). \tag{7}$$

Operationally, each reverse transition first performs an unconditional update to obtain a proposal $\boldsymbol{\theta}'_{t-1}$ and then applies a likelihood-guidance correction,

$$\boldsymbol{\theta}_{t-1} \leftarrow \boldsymbol{\theta}'_{t-1} + \zeta_t \nabla_{\boldsymbol{\theta}_t} \log p\Big(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t), \boldsymbol{\xi}\Big), \tag{8}$$

where ζ_t is the guidance step size. For sequential BOED, this plug-in is applied to the sum of likelihood terms in H_k. When scoring a candidate, the prospective term $p(\mathbf{y} \mid \boldsymbol{\theta}, \boldsymbol{\xi})$ is appended to the guidance for H_{k-1}. Repeating this procedure over $t = T, \ldots, 1$ yields samples whose terminal denoised output $\hat{\boldsymbol{\theta}}_0$ approximates draws from the desired posterior $p(\boldsymbol{\theta} \mid H_k)$.

3 METHOD

Building on the complex-prior setting and the guidance-based posterior sampling framework in Section 2, we now derive our new BOED objective based on a *latent-time steering cost* accumulated along likelihood-guided diffusion trajectories. We, therefore, arrive at a practical MINDE-based surrogate for the per-step EIG when the relevant distributions are complex.

3.1 MINDE FOR INFORMATION GAIN

At BOED step k, given the current history H_{k-1} and a candidate design $\boldsymbol{\xi}$, we must draw posterior samples and evaluate the step-k EIG defined in (2). When the prior and/or posterior is complex, direct posterior sampling and EIG estimation can be computationally demanding. We therefore estimate the EIG using the MINDE latent-time identity. Specifically, we derive Corollary 2 based on Proposition 1, to express the information gain as an expected steering cost computed along the corresponding guided posterior-sampling trajectory.

Corollary 2. *Fix history H_{k-1}, experiment $\boldsymbol{\xi}$, and an observation $\mathbf{y}$. Let the prior be $p_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid H_{k-1})$ and the posterior be $q_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta} \mid H_k) \propto p_0(\boldsymbol{\theta})\, p(\mathbf{y} \mid \boldsymbol{\theta}, \boldsymbol{\xi})$. Let q_t and p_t be the corresponding forward-noised marginals under the same kernel. Then, using Bayes' rule at latent time t,*

$$\nabla_{\boldsymbol{\theta}_t} \log q_t(\boldsymbol{\theta}_t) = \quad \nabla_{\boldsymbol{\theta}_t} \log p_t(\boldsymbol{\theta}_t) + \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}). \tag{9}$$

Substituting this expression into (5) *gives*

$$D_{\mathrm{KL}}\Big(p(\boldsymbol{\theta} \mid H_{k-1}, \boldsymbol{\xi}, \mathbf{y}) \,\Big\|\, p(\boldsymbol{\theta} \mid H_{k-1})\Big) \approx \frac{1}{2} \int_0^T a(t)^2\, \mathbb{E}_{\boldsymbol{\theta}_t \sim q_t}\Big[\big\|\nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})\big\|_2^2\Big]\, \mathrm{d}t. \tag{10}$$

Taking $\mathbb{E}_{\mathbf{y} \sim p(\cdot \mid H_{k-1}, \boldsymbol{\xi})}[\,\cdot\,]$ gives the corresponding approximation to $\mathrm{EIG}_k(\boldsymbol{\xi})$.

Corollary 2 shows that the per-outcome information gain can be approximated by a latent-time integral of a squared *noised-likelihood score*, and $\mathrm{EIG}_k(\boldsymbol{\xi})$ is obtained by averaging this quantity over the predictive distribution $p(\mathbf{y} \mid H_{k-1}, \boldsymbol{\xi})$. To turn this identity into a practical estimator for BOED, we must (i) express the integrand in terms of quantities available from our guided diffusion sampler, and (ii) discretize the latent-time integral in a way that can be evaluated efficiently along sampled trajectories. Using the DPS plug-in defined above, we next construct a Monte Carlo estimator by averaging over sampled latent states and discretizing the latent-time integral. The detailed derivation of Corollary 2 is provided in Appendix B.

3.2 MONTE CARLO ESTIMATION FOR CANDIDATE DESIGNS

For each candidate design $\boldsymbol{\xi}$, we estimate the MINDE–DPS surrogate in (10) by averaging the steering cost over prospective observations and posterior samples. The computation consists of three steps: generating clean posterior samples $\boldsymbol{\theta}_0$, constructing forward-noised samples $\boldsymbol{\theta}_t$, and accumulating their squared likelihood guidance.

We first draw M observations $\mathbf{y}^{(m)} \sim p(\cdot \mid H_{k-1}, \boldsymbol{\xi})$. For each observation, we use DPS with N_{diff} reverse-diffusion steps to generate S approximate samples $\boldsymbol{\theta}_0^{(m,s)}$ from $p(\boldsymbol{\theta} \mid H_{k-1}, \mathbf{y}^{(m)}, \boldsymbol{\xi})$. For the construction below, we also retain the intermediate states of these DPS trajectories.

Hybrid forward noising. The steering-cost integral requires samples at multiple noise levels. For Darcy source-position design, we construct these samples by combining direct noising of clean posterior samples with perturbations of saved DPS states. Let $t_0 < \cdots < t_N$ denote the subsampled levels in increasing noise order, write $\sigma_n = \sigma(t_n)$, and let $\mathcal{I} \subseteq \{1, \ldots, N\}$ index the retained intervals.

For each clean posterior sample $\boldsymbol{\theta}_0^{(m,s)}$, direct noising generates L_{dir} samples at each retained noise level:

$$\boldsymbol{\theta}_{t_n}^{(m,s,\ell,\mathrm{dir})} = \boldsymbol{\theta}_0^{(m,s)} + \sigma_n \boldsymbol{\epsilon}_{n,\ell}^{(m,s,\mathrm{dir})}, \qquad \ell = 1, \ldots, L_{\mathrm{dir}}.$$

For the second component, we generate L_{traj} perturbations of the saved DPS state $\boldsymbol{\theta}_{t_{n-1}}^{(m,s,\mathrm{DPS})}$ at the preceding noise level:

$$\boldsymbol{\theta}_{t_n}^{(m,s,\ell,\mathrm{traj})} = \boldsymbol{\theta}_{t_{n-1}}^{(m,s,\mathrm{DPS})} + \sqrt{\sigma_n^2 - \sigma_{n-1}^2}\, \boldsymbol{\epsilon}_{n,\ell}^{(m,s,\mathrm{traj})}, \qquad \ell = 1, \ldots, L_{\mathrm{traj}}.$$

All $\boldsymbol{\epsilon}$ are independent standard Gaussian draws. The two components reuse the clean outputs and saved intermediate states of the same DPS trajectories. The resulting latent samples do not constitute additional independent posterior draws.

Steering-cost estimation. At each retained noise level, we pool the $L = L_{\mathrm{dir}} + L_{\mathrm{traj}}$ latent samples and evaluate the DPS likelihood guidance $\boldsymbol{g}_{\mathrm{DPS}}(\boldsymbol{\theta}_t, t; \mathbf{y}, \boldsymbol{\xi}) := \nabla_{\boldsymbol{\theta}_t} \log p\Big(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t), \boldsymbol{\xi}\Big)$.

Averaging the squared guidance over observations and samples, and applying quadrature in noise variance, gives

$$\widehat{\mathrm{EIG}}_k(\boldsymbol{\xi}) = \frac{1}{2MSL} \sum_{m=1}^{M} \sum_{s=1}^{S} \sum_{n \in \mathcal{I}} (\sigma_n^2 - \sigma_{n-1}^2) \sum_{\ell=1}^{L} \left\| \boldsymbol{g}_{\mathrm{DPS}}\Big(\boldsymbol{\theta}_{t_n}^{(m,s,\ell)}, t_n; \mathbf{y}^{(m)}, \boldsymbol{\xi}\Big) \right\|_2^2, \tag{11}$$

where ℓ indexes the pooled states obtained by direct noising and by perturbing saved DPS states. We average their squared guidance contributions.

The accumulated history and the prospective observation are both used for posterior sampling, while the steering-cost integrand includes only the guidance term associated with the prospective observation. Once the DPS trajectories have been generated, we construct the quadrature samples by adding noise to their clean outputs and saved intermediate states, without generating a new reverse trajectory for each perturbation.

3.3 Error analysis

We study how accurately the DPS steering cost approximates the step-k expected information gain for a candidate design $\boldsymbol{\xi}$. For an observation $\mathbf{y}$, let $\boldsymbol{g}_t^\star(\boldsymbol{\theta}_t) := \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})$ denote the exact noised-likelihood guidance, and let $\boldsymbol{g}_{\mathrm{DPS}}(\boldsymbol{\theta}_t)$ denote its learned DPS approximation. We write q_t for the forward-noised marginal of the exact posterior and $\widehat{q}_t$ for the corresponding marginal induced by DPS. The result below quantifies the resulting EIG approximation error.

Define the weighted DPS guidance discrepancy as

$$\varepsilon_{\mathrm{DPS}}(\boldsymbol{\xi}, T) := \left\{ \mathbb{E}_{\mathbf{y}} \left[\int_0^T a(t)^2 \, \mathbb{E}_{q_t} \left[\|\boldsymbol{g}_t^\star - \boldsymbol{g}_{\mathrm{DPS}}\|_2^2 \right] \mathrm{d}t \right] \right\}^{1/2}, \tag{12}$$

and let $\mathcal{R}_T(\boldsymbol{\xi}) := \mathbb{E}_{\mathbf{y}}[D_{\mathrm{KL}}(q_T \| p_T)]$ denote the terminal residual.

Theorem 3 (MINDE–DPS EIG error rate)**.** *Under the assumptions stated in Appendix E, for a fixed diffusion horizon T and trained model, the estimator based on M independent outer samples satisfies*

$$\left| \mathrm{EIG}_k(\boldsymbol{\xi}) - \widehat{\mathrm{EIG}}_{k,M}(\boldsymbol{\xi}) \right| = O_p\Big(\varepsilon_{\mathrm{DPS}}(\boldsymbol{\xi}, T) + \sqrt{\mathcal{R}_T(\boldsymbol{\xi})} + \mathcal{R}_T(\boldsymbol{\xi}) + M^{-1/2} \Big). \tag{13}$$

The bound applies to an idealized continuous-time estimator with direct forward noising under the stated assumptions. For fixed T and model, increasing M reduces Monte Carlo error at rate $O_p(M^{-1/2})$, but not guidance or terminal-truncation bias. See Appendix E for the proof.

3.4 Sequential BOED Implementation

After acquiring the history $H_k = \{(\mathbf{y}_i, \boldsymbol{\xi}_i)\}_{i=1}^k$, we reuse the same unconditional diffusion prior and condition the reverse trajectory on all acquired observations:

$$\boldsymbol{\theta}_{t-1} \leftarrow \boldsymbol{\theta}'_{t-1} + \zeta_t \sum_{i=1}^{k} \nabla_{\boldsymbol{\theta}_t} \log p\Big(\mathbf{y}_i \mid \hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t), \boldsymbol{\xi}_i\Big), \tag{14}$$

where $\boldsymbol{\theta}'_{t-1}$ is the unconditional reverse-diffusion proposal in (8). Thus, the retained samples at step k approximate the current belief $p(\boldsymbol{\theta} \mid H_k)$ using both the prior dynamics and the full measurement history.

At BOED step k, the EIG contains two nested expectations. We first approximate the outer predictive expectation in (2) by drawing $\boldsymbol{\theta}^{(m)} \sim p(\boldsymbol{\theta} \mid H_{k-1})$ and $\mathbf{y}_k^{(m)} \sim p(\cdot \mid \boldsymbol{\theta}^{(m)}, \boldsymbol{\xi})$. The samples $\boldsymbol{\theta}^{(m)}$ are generated using (14) with the accumulated history H_{k-1}.

For each predictive draw $\mathbf{y}_k^{(m)}$, MINDE in (10) requires the inner expectation with respect to the updated posterior $p(\boldsymbol{\theta}_t \mid H_{k-1}, \mathbf{y}_k^{(m)}, \boldsymbol{\xi})$. We generate these samples by appending the final likelihood

term associated with $(\mathbf{y}_k^{(m)}, \boldsymbol{\xi})$ to the accumulated guidance in (14), and estimate the resulting inner expectation using the Monte Carlo procedure in Section 3.2. Although the posterior trajectory uses the prior and all terms in H_{k-1}, the MINDE score difference retains only the likelihood guidance induced by the newly appended observation. We evaluate $\mathrm{EIG}_k(\boldsymbol{\xi})$ for each candidate, select the maximizer, append its acquired observation to the history, and repeat this procedure at step $k+1$.

4 Numerical Experiments

We evaluate whether the guided trajectories used for posterior inference also provide a useful signal for choosing the next experiment. We first validate the steering-cost surrogate on source localization, compare achieved information gain on source localization and Constant Elasticity of Substitution (CES), and then assess high-dimensional parameter-field reconstruction in Darcy flow. The design variables are sensor locations for source localization, basket pairs for CES, and observation patches or forcing-source locations for Darcy flow. Full problem formulations and implementation details are provided in Appendix F.

4.1 Performance on BOED Benchmarks

A key practical distinction of our approach is that source-localization design does not require training an additional design network: the EIG surrogate is obtained directly from guided posterior-sampling trajectories. We first examine whether these trajectories provide useful information for both inference and design. As sensors are placed sequentially, posterior samples become increasingly concentrated around the true source locations (Figure 2).

To assess the design information captured by the surrogate, we compare its candidate rankings with a sampling-based sequential Prior Contrastive Estimator (sPCE) reference (Foster et al., 2021). Across ten BOED steps, Spearman correlations range from 0.855 to 0.924, indicating substantial agreement over the candidate pool. This agreement is less consistent at the point of selection: the chosen design's sPCE score falls below the highest candidate score by 5.3%–33.7% of the within-step score range. The trajectories therefore provide a useful ranking signal, but do not always identify the same maximizer as the reference. Per-step metrics and additional visualizations are provided in Table 3 and Appendix G.1.

We next evaluate how this design signal translates into achieved information gain on source localization and CES (Table 1). On source localization, MINDE-DPS achieves an EIG lower bound of 6.92, improving on SG-BOED 5.55, while remaining below DAD 7.97 and ALINE 8.51. The same approach also improves on VPCE 9.40 and DAD 10.77 on CES, reaching 12.97 compared with 13.50 for ALINE. These results demonstrate the utility of the trajectory-based design signal while also showing that its agreement with reference rankings does not guarantee the strongest sequential design performance. In source localization, the forward model maps four unknown source coordinates to one scalar measurement, making the inverse problem underdetermined from a single observation. We discuss how this ill-posedness may affect likelihood-guided inference and design scoring in Section 5. The source-localization results for Random, SG-BOED, and DAD in Table 1 are taken from Ivanova et al. (2021), and the CES baseline results are taken from Huang et al. (2025). We

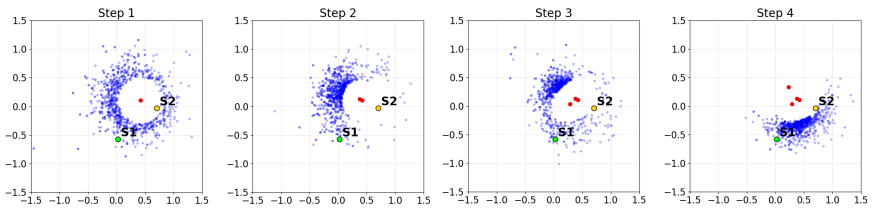

Figure 2: Posterior evolution under sequential sensor placement. Blue points show 200 posterior samples conditioned on the measurements available at each displayed step; red markers indicate the sensors installed so far, and green and yellow circles mark the true source locations. Samples become more concentrated around the true locations as measurements accumulate.

reevaluate ALINE under our source-localization setting. This ALINE result and all MINDE-DPS results are from our experiments.

Table 1: BOED performance on source localization and CES.

Source Localization		*Constant Elasticity of Substitution*	
Method	EIG(↑)	Method	EIG(↑)
Random	4.79 ± 0.04	Random	9.05 ± 0.26
SG-BOED	5.55 ± 0.03	VPCE	9.40 ± 0.27
DAD	7.97 ± 0.03	DAD	10.77 ± 0.15
ALINE	8.51 ± 0.07	ALINE	13.50 ± 0.15
MINDE-DPS	6.92 ± 0.06	MINDE-DPS	12.97 ± 0.19

4.2 High-Dimensional Experimental Design for Inverse Problems: Darcy Flow

We next consider a larger Darcy flow inverse problem. The unknown parameters form a spatially varying parameter field represented on a 64×64 grid. We consider two design problems: selecting observation patches and selecting source positions in the forcing channel. Both experiments aim to improve inference of the parameter field. We use a fixed diffusion model pretrained jointly on the forcing, parameter field, and solution fields following FunDPS. Our BOED procedure uses only the parameter field channel for posterior inference and design scoring. The details about the PDE, data generation, and diffusion-prior settings are given in Appendix F.3.

4.2.1 Sequential Observation-Patch Selection

Inspired by the patchwise acquisition setting of Iollo et al. (2025), we test whether the proposed criterion can guide sequential acquisition under a limited observation budget. Patchwise acquisition defines a finite set of candidate regions and allows reconstruction quality to be evaluated as the observation budget increases. We evaluate 400 instances sharing the same coefficient field but differing in source position. We partition the 64×64 domain into 25 non-overlapping patches arranged in a 5×5 grid. At each step, we recompute the scores of the remaining patches under the posterior conditioned on all previous observations, acquire the highest-scoring patch, and update the reconstruction using the accumulated measurements. The reconstructions and error maps in Figure 7 show improved recovery of the field's spatial structure as more patches are observed. At a fixed budget of three patches, EIG-guided selection also yields lower parameter reconstruction errors than random selection in the examples shown in Appendix G.2. This comparison illustrates the benefit of selecting informative patches under the same observation budget.

4.2.2 Sequential Source-Position Design

We next consider source-position design, motivated by the practical question of where to apply forcing to obtain informative observations for parameter field reconstruction. We evaluate 100 candidate source locations, computing their scores from saved DPS trajectories using the Monte Carlo procedure described in Appendix F.3. Figure 3 illustrates the two-stage source-position design process. Stage 1 scores candidate forcing locations and selects the highest-scoring location. After assimilating the first measurement, Stage 2 recomputes the scores under the updated posterior to assess the incremental information provided by each candidate. In the illustrated example, the high-scoring region shifts away from the previously selected location, consistent with seeking information that complements the first measurement. The reconstruction results show that these sequentially selected

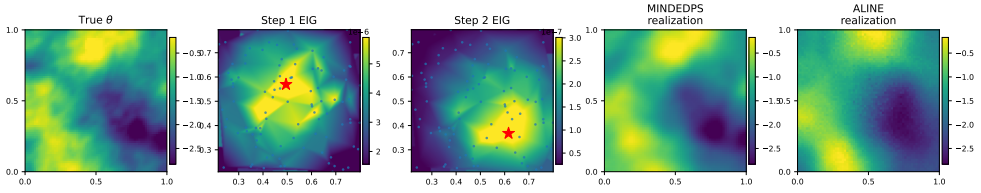

Figure 3: Two-stage sequential EIG for Darcy source-position design. Stage 1 scores candidate forcing locations. Stage 2 recomputes EIG after the first measurement. Red stars mark the location selected at each stage.

observations support improved parameter recovery, with lower mean reconstruction error at Stage 2 than at Stage 1 (Table 2). Additional examples are provided in Figure 6. For our comparison with ALINE, both methods use the same dataset of $1{,}000$ parameter fields and the same candidate pool. MINDE-DPS achieves lower mean reconstruction error than both ALINE configurations at both stages (Table 2). Across 20 test cases, its mean relative L_2 error is 0.229 at Stage 1 and 0.214 at Stage 2, compared with 0.304 and 0.285 for ALINE at 64×64 resolution. These correspond to relative reductions of 24.7% and 24.9%, respectively. All three methods improve after the second measurement, with MINDE-DPS retaining the lowest mean error. These results support the use of a fixed diffusion prior for sequential design and high-dimensional field reconstruction.

Table 2: Reconstruction performance on the Darcy problem. Relative L_2 errors are reported as mean $\pm$ standard deviation over 20 test cases.

Method	Stage 1	Stage 2
ALINE (8×8)	0.328 ± 0.115	0.314 ± 0.106
ALINE (64×64)	0.304 ± 0.121	0.285 ± 0.111
MINDE-DPS	$\mathbf{0.229 \pm 0.166}$	$\mathbf{0.214 \pm 0.127}$

5 Discussion

We reframe BOED with pretrained generative priors by viewing *posterior conditioning as latent-time steering*. The resulting steering-cost utility is computed along the guided diffusion process, providing a practical BOED objective for complex, high-dimensional posteriors.

When guidance is reliable. When the observation model is well specified and likelihood guidance is reliable, the same pretrained prior can support both high-dimensional posterior reconstruction and design scoring. Our Darcy flow results illustrate this possibility.

Understanding the performance gap. MINDE-DPS combines an unconditional prior with likelihood guidance at deployment, so errors in posterior sampling and the DPS likelihood-score approximation can affect design evaluation. In the small-scale benchmarks, however, ALINE performs better. Its transformer-based model learns directly from simulated parameter–observation relationships and acquisition histories, jointly learning inference and acquisition for the specific task (Huang et al., 2025). In source localization, the forward model maps four unknown source coordinates to one scalar measurement, making the inverse problem underdetermined from a single observation. The gap is particularly pronounced here, where ill-posedness and multimodality can make likelihood-guided inference more challenging—and where substantial agreement in candidate rankings does not necessarily translate into equally effective sequential designs.

Practical relevance. The practical value of MINDE-DPS lies in reusing a generative prior for both posterior inference and experimental design without training a separate design policy. This is particularly relevant when a suitable pretrained prior is available and high-dimensional posterior inference is already required. Darcy flow illustrates this setting. ALINE's policy training involves observation-dependent rollouts, so longer rollouts require more sequential network evaluations and more computation during backpropagation. MINDE-DPS uses its fixed prior at inference time, allowing candidate and Monte Carlo evaluations within each design stage to be batched or parallelized. Reusing a prior learned from existing simulations is also attractive when forward-model solves make additional training-data generation costly. We therefore view MINDE-DPS as a complementary approach for such settings, rather than a uniformly more accurate or computationally cheaper alternative to amortized design methods.

Limitations and future directions. A limitation is potential bias and instability in DPS guidance. The intractable noised-likelihood score $\nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})$ is replaced by a denoiser-based surrogate. If the score model is undertrained or likelihood gradients are poorly scaled, the guidance can become overly large, leading to inflated steering costs and degraded posterior samples. The sensitivity results in Table 4 show that guidance strength and the number of diffusion steps affect both reconstruction error and the estimated steering cost. Our error analysis in Section 3.3 separates guidance approximation error from terminal truncation and Monte Carlo sampling error. Increasing the sampling budget can reduce Monte Carlo variability, but does not by itself correct biased likelihood guidance. Several extensions could improve robustness and efficiency. A distributional estimate of the steering-cost contribution at each latent time could drive adaptive discretization, concentrating evaluations where the contribution is large. Faster generative priors (e.g., flow matching), combined with stochastic posterior samplers (Kim et al., 2025), could reduce the number of latent steps required for sampling and cost accumulation.

AI USE STATEMENT

In this work, we used generative AI tools for assisting with the implementation of our computational methods, translating text, designing and providing feedback on experiments, interpreting results, and reviewing author-written proof drafts for missing assumptions and logical gaps. We have not used generative AI tools for providing critical ingredients for proving mathematical claims, and tasks involving qualitative and thematic data analysis are not applicable to this work. Additionally, we used generative AI tools for creating and modifying scientific figures, suggesting experimental parameters, writing and revising software code, drafting parts of the manuscript, and editing the manuscript to improve clarity and readability. We have reviewed all AI-assisted work. We reviewed the AI-assisted manuscript revisions and verified and tested the AI-assisted code for correctness. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work studies Bayesian optimal experimental design using synthetic benchmarks and simulated physical systems. The experiments do not involve human participants or personal data. We have not identified specific ethical concerns associated with the methods and experiments presented in this work.

REPRODUCIBILITY STATEMENT

The proposed method and its assumptions are described in Section 3, with supporting derivations and proofs provided in Appendices B and E. Experimental setups, data-generation procedures, and implementation details are described in Appendix F and Appendix D.2. We will make our source code, including scripts for data generation, training, and evaluation, together with the experiment configurations, publicly available to support reproduction of our results.

REFERENCES

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

Jiayuan Dong, Christian Jacobsen, Mehdi Khalloufi, Maryam Akram, Wanjiao Liu, Karthik Duraisamy, and Xun Huan. Variational bayesian optimal experimental design with normalizing flows. *Computer Methods in Applied Mechanics and Engineering*, 433:117457, 2025.

Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential Bayesian experimental design. In *International conference on machine learning*, pp. 3384–3395. PMLR, 2021.

Giulio Franzese, Mustapha Bounoua, and Pietro Michiardi. MINDE: Mutual information neural diffusion estimation. In *International Conference on Learning Representations (ICLR)*, 2024. arXiv:2310.09031v2.

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

Daolang Huang, Xinyi Wen, Ayush Bharti, Samuel Kaski, and Luigi Acerbi. Aline: Joint amortization for bayesian inference and active data acquisition. *arXiv preprint arXiv:2506.07259*, 2025.

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

Jacopo Iollo, Christophe Heinkelé, Pierre Alliez, and Florence Forbes. Bayesian experimental design via contrastive diffusions. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=h8yg0hT96f.

Desi R Ivanova, Adam Foster, Steven Kleinegesse, Michael U Gutmann, and Thomas Rainforth. Implicit deep adaptive design: Policy-based experimental design without likelihoods. *Advances in Neural Information Processing Systems*, 34:25785–25798, 2021.

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

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

Steven Kleinegesse and Michael U Gutmann. Bayesian experimental design for implicit models by mutual information neural estimation. In *International conference on machine learning*, pp. 5316–5326. PMLR, 2020.

Adhithya Saravanan, Rik Knowles, Gavin Kerrigan, and Tom Rainforth. Diffbed: Scaling bayesian experimental design to high-dimensions. In *International Conference on Learning Representations*, volume 2026, pp. 131048–131073, 2026.

Wanggang Shen, Jiayuan Dong, and Xun Huan. Variational sequential optimal experimental design using reinforcement learning. *Computer Methods in Applied Mechanics and Engineering*, 444: 118068, 2025.

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

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

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

A Related Work

Diffusion and score-based generative models. Diffusion models and score-based generative models learn complex data distributions by reversing a forward noising process, yielding either a denoiser or a score network that parameterizes $\nabla_x \log p_t(x)$ at intermediate noise levels. Canonical formulations include Denoising Diffusion Probabilistic Models (DDPMs) (Ho et al., 2020) and the Stochastic Differential Equation (SDE) based score framework (Song et al., 2021b), which also admits a probability-flow Ordinary Differential Equation (ODE) interpretation for deterministic sampling. Because these models provide high-quality unconditional samples and tractable gradient information through the learned score/denoiser, they have been widely adopted as expressive learned priors for downstream inference and decision-making tasks.

Inverse problems and PDE-constrained optimization with diffusion models. Substantial literature uses pretrained diffusion priors as plug-and-play posteriors for inverse problems by injecting measurement information into the reverse-time sampler. For linear inverse problems, Denoising Diffusion Restoration Model (DDRM) provides an efficient unsupervised posterior sampling procedure built on a pretrained DDPM (Kawar et al., 2022). DPS extends this idea to general noisy (including nonlinear) inverse problems by combining diffusion sampling with likelihood-gradient guidance and a denoiser-based plug-in for intermediate likelihood terms (Chung et al., 2023). Recent work pushes

these ideas toward scientific and Partial Differential Equation (PDE) settings. Namely, the DiffusionPDE trains a diffusion model over (coefficient, solution) pairs and guides sampling with sparse observations and PDE residuals to solve forward and inverse PDE problems (Huang et al., 2024). FunDPS develops a function-space diffusion framework with guidance and theoretical support via extensions of Tweedie-type identities to infinite-dimensional settings (Yao et al., 2025).

A closely related line of work is Bayesian experimental design with *Contrastive Diffusions* (CoDiff) (Iollo et al., 2025), which derives a tractable estimator for gradient-based design optimization. A key device is a *pooled posterior* proposal distribution, defined by pooling of posteriors induced by likelihood function. In addition, CoDiff adopts a filtering perspective for conditional diffusion sampling by conditioning on a *noised observation* at intermediate time t to obtain tractable likelihood-score terms in linear measurements. Our work uses a complementary conditional mechanism based on a DPS-style denoiser plug-in. Another related line of work DiffBED (Saravanan et al., 2026) is complementary. It addresses high-dimensional *design* spaces by learning a diffusion prior over feasible designs and guiding its reverse process with an EIG, thereby generating realistic and informative designs such as high-resolution images. Our diffusion prior instead models the unknown *parameter* $\boldsymbol{\theta}$. We use DPS likelihood guidance to sample from the corresponding approximate posterior, then use the MINDE score–KL identity to express EIG as a latent-time steering cost along that guided trajectory. Thus, DiffBED regularizes optimization through a prior on $\boldsymbol{\xi}$, whereas we use a prior on $\boldsymbol{\theta}$ to make posterior sampling and EIG evaluation practical for high-dimensional, non-Gaussian parameter beliefs. The novelty is this steering-cost reformulation of EIG, rather than the separate use of MINDE or DPS: a single pretrained diffusion network is reused for posterior sampling and EIG scoring throughout the BOED loop, without fitting a separate posterior estimator after each observation.

B Derivation

Here we provide the detailed derivation of Corollary 2.

Proof. Fix the history H_{k-1}, the candidate design $\boldsymbol{\xi}$, and an observation $\mathbf{y}$. Throughout this proof, H_{k-1} is fixed, so we suppress it in the notation:

$$p_0(\boldsymbol{\theta}) := p(\boldsymbol{\theta} \mid H_{k-1}), \quad p(\mathbf{y} \mid \boldsymbol{\xi}) := p(\mathbf{y} \mid H_{k-1}, \boldsymbol{\xi}), \quad p_t(\boldsymbol{\theta}_t) := p_t(\boldsymbol{\theta}_t \mid H_{k-1}),$$

and we write $p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})$ for the *induced* time-t likelihood defined below (which still depends on p_0).

Step 1 (Bayes at the clean stage). Define

$$p_0(\boldsymbol{\theta}_0) = p(\boldsymbol{\theta}_0 \mid H_{k-1}), \qquad q_0(\boldsymbol{\theta}_0) = p(\boldsymbol{\theta}_0 \mid H_{k-1}, \mathbf{y}, \boldsymbol{\xi}).$$

By Bayes' rule,

$$q_0(\boldsymbol{\theta}_0) = \frac{p_0(\boldsymbol{\theta}_0)\, p(\mathbf{y} \mid \boldsymbol{\theta}_0, \boldsymbol{\xi})}{p(\mathbf{y} \mid \boldsymbol{\xi})}, \qquad p(\mathbf{y} \mid \boldsymbol{\xi}) = \int p(\mathbf{y} \mid \boldsymbol{\theta}_0, \boldsymbol{\xi})\, p_0(\boldsymbol{\theta}_0)\, \mathrm{d}\boldsymbol{\theta}_0. \tag{15}$$

Step 2 (Forward diffusion and time-t marginals). Let $K_{0,t}(\boldsymbol{\theta}_0, \cdot)$ be the Markov kernel of the forward diffusion (as in Proposition 1). Define

$$p_t(\boldsymbol{\theta}_t) := \int p_0(\boldsymbol{\theta}_0)\, K_{0,t}(\boldsymbol{\theta}_0, \boldsymbol{\theta}_t)\, \mathrm{d}\boldsymbol{\theta}_0, \qquad q_t(\boldsymbol{\theta}_t) := \int q_0(\boldsymbol{\theta}_0)\, K_{0,t}(\boldsymbol{\theta}_0, \boldsymbol{\theta}_t)\, \mathrm{d}\boldsymbol{\theta}_0. \tag{16}$$

Substituting (15) into q_t yields

$$q_t(\boldsymbol{\theta}_t) = \frac{1}{p(\mathbf{y} \mid \boldsymbol{\xi})} \int p_0(\boldsymbol{\theta}_0)\, p(\mathbf{y} \mid \boldsymbol{\theta}_0, \boldsymbol{\xi})\, K_{0,t}(\boldsymbol{\theta}_0, \boldsymbol{\theta}_t)\, \mathrm{d}\boldsymbol{\theta}_0. \tag{17}$$

Step 3 (Diffusion smoothing distribution). Define the forward-diffusion *smoothing* distribution under p_0:

$$p_0(\boldsymbol{\theta}_0 \mid \boldsymbol{\theta}_t) := p(\boldsymbol{\theta}_0 \mid \boldsymbol{\theta}_t, H_{k-1}) = \frac{p_0(\boldsymbol{\theta}_0)\, K_{0,t}(\boldsymbol{\theta}_0, \boldsymbol{\theta}_t)}{p_t(\boldsymbol{\theta}_t)}, \tag{18}$$

which is well-defined whenever $p_t(\boldsymbol{\theta}_t) > 0$.

Step 4 (Induced time-t likelihood). Define the *induced* time-t likelihood by averaging the clean likelihood over $p_0(\boldsymbol{\theta}_0 \mid \boldsymbol{\theta}_t)$:

$$p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}) := \int p(\mathbf{y} \mid \boldsymbol{\theta}_0, \boldsymbol{\xi})\, p_0(\boldsymbol{\theta}_0 \mid \boldsymbol{\theta}_t)\, \mathrm{d}\boldsymbol{\theta}_0. \tag{19}$$

Substituting (18) into (19) gives the equivalent ratio form

$$p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}) = \frac{1}{p_t(\boldsymbol{\theta}_t)} \int p_0(\boldsymbol{\theta}_0)\, p(\mathbf{y} \mid \boldsymbol{\theta}_0, \boldsymbol{\xi})\, K_{0,t}(\boldsymbol{\theta}_0, \boldsymbol{\theta}_t)\, \mathrm{d}\boldsymbol{\theta}_0. \tag{20}$$

Step 5 (Factorization of q_t). Comparing (17) and (20) yields

$$q_t(\boldsymbol{\theta}_t) = \frac{p_t(\boldsymbol{\theta}_t)\, p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})}{p(\mathbf{y} \mid \boldsymbol{\xi})} \;\propto\; p_t(\boldsymbol{\theta}_t)\, p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}), \tag{21}$$

where the proportionality constant $p(\mathbf{y} \mid \boldsymbol{\xi})^{-1}$ does not depend on $\boldsymbol{\theta}_t$.

Step 6 (Score decomposition at time t). Taking logarithms in (21) and differentiating with respect to (w.r.t.) $\boldsymbol{\theta}_t$ gives

$$\nabla_{\boldsymbol{\theta}_t} \log q_t(\boldsymbol{\theta}_t) = \nabla_{\boldsymbol{\theta}_t} \log p_t(\boldsymbol{\theta}_t) + \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}), \tag{22}$$

since $\nabla_{\boldsymbol{\theta}_t} \log p(\mathbf{y} \mid \boldsymbol{\xi}) = 0$. (Interchanging differentiation and integration is justified under the regularity assumptions in Proposition 1.)

Step 7 (MINDE-based estimator). From (22),

$$\nabla_{\boldsymbol{\theta}_t} \log q_t(\boldsymbol{\theta}_t) - \nabla_{\boldsymbol{\theta}_t} \log p_t(\boldsymbol{\theta}_t) = \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}).$$

Therefore the integrand in (5) becomes

$$\Big\| \nabla_{\boldsymbol{\theta}_t} \log q_t(\boldsymbol{\theta}_t) - \nabla_{\boldsymbol{\theta}_t} \log p_t(\boldsymbol{\theta}_t) \Big\|_2^2 = \Big\| \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}) \Big\|_2^2.$$

Applying Proposition 1 with these p_0, q_0 yields

$$D_{\mathrm{KL}}(q_0 \| p_0) = D_{\mathrm{KL}}(q_T \| p_T) + \frac{1}{2} \int_0^T a(t)^2\, \mathbb{E}_{\boldsymbol{\theta}_t \sim q_t} \Big[\| \nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}) \|_2^2 \Big]\, \mathrm{d}t. \tag{23}$$

Finally, $D_{\mathrm{KL}}(q_0 \| p_0) = D_{\mathrm{KL}}\big(p(\boldsymbol{\theta} \mid H_{k-1}, \mathbf{y}, \boldsymbol{\xi}) \,\|\, p(\boldsymbol{\theta} \mid H_{k-1})\big)$ by construction. For sufficiently large terminal noise/time, the forward diffusion washes out dependence on the initial law so that $p_T \approx q_T$ and $D_{\mathrm{KL}}(q_T \| p_T) \approx 0$, yielding (10).

Step 8 (DPS plug-in for computability). The induced score $\nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi})$ is generally intractable for simulator likelihoods. DPS replaces it by evaluating the clean likelihood at the denoised prediction $\hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t)$ and backpropagating to $\boldsymbol{\theta}_t$:

$$\nabla_{\boldsymbol{\theta}_t} \log p_t(\mathbf{y} \mid \boldsymbol{\theta}_t, \boldsymbol{\xi}) \;\approx\; \nabla_{\boldsymbol{\theta}_t} \log p\Big(\mathbf{y} \mid \hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t), \boldsymbol{\xi}\Big) \;=:\; \boldsymbol{g}_{\mathrm{DPS}}(\boldsymbol{\theta}_t, t; \mathbf{y}, \boldsymbol{\xi}),$$

which leads to the computable steering-cost surrogate used in our estimator by DPS plug-in guidance. Averaging the resulting DPS steering cost over $\mathbf{y} \sim p(\cdot \mid H_{k-1}, \boldsymbol{\xi})$ yields a computable surrogate for $\mathrm{EIG}_k(\boldsymbol{\xi})$. □

C KL DIVERGENCE COMPUTATION

Fix $T < \infty$ and consider the forward Itô diffusion on $\mathbb{R}^N$

$$dX_t = f_t X_t\, dt + g_t\, dW_t,$$

with continuous $f_t \leq 0$ and $g_t > 0$. For any initial law μ on $\mathbb{R}^N$, let P^μ be the path measure of the forward SDE with $X_0 \sim \mu$, and define the time-t marginal

$$\nu_t^\mu := (X_t)_\# P^\mu, \qquad \nu_0^\mu = \mu.$$

Assume ν_t^μ admits a density $\bar{\nu}_t^\mu$ and define the (Stein) score

$$s_t^\mu(x) := \nabla_x \log \bar{\nu}_t^\mu(x).$$

Given two initial laws p_0 and q_0, set

$$p_t := \nu_t^{p_0}, \qquad q_t := \nu_t^{q_0}.$$

Assume the scores exist and are square-integrable along the forward process. Then

$$D_{\mathrm{KL}}(q_0 \| p_0) = \mathbb{E}_{P^{q_0}} \left[\int_0^T \frac{g_t^2}{2} \left\| s_t^{q_0}(X_t) - s_t^{p_0}(X_t) \right\|_2^2 dt \right] + D_{\mathrm{KL}}(q_T \| p_T).$$

D Derivation of Gaussian-to-Gaussian Diffusion

Consider a diffusion process where the data prior $p(x_0)$ is a Gaussian distribution:

$$x_0 \sim \mathcal{N}(\mu_{data}, \Sigma_{data}) \tag{24}$$

The forward diffusion process at any timestep t adds Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$ according to the schedule $\bar{\alpha}_t$:

$$x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon \tag{25}$$

Since x_t is a linear combination of two independent Gaussians (x_0 and ϵ), the marginal distribution $p_t(x_t)$ is also Gaussian:

$$p_t(x_t) = \mathcal{N}(x_t; \mu_t, \Sigma_t) \tag{26}$$

We can derive the exact mean μ_t and covariance Σ_t analytically:

$$\mu_t = \mathbb{E}[x_t] = \sqrt{\bar{\alpha}_t} \mathbb{E}[x_0] = \sqrt{\bar{\alpha}_t} \mu_{data} \tag{27}$$

$$\Sigma_t = \mathrm{Cov}[x_t] = \bar{\alpha}_t \Sigma_{data} + (1 - \bar{\alpha}_t) I \tag{28}$$

D.1 The Analytic Score Function

The score function is defined as the gradient of the log-probability density with respect to the input x_t. For a multivariate Gaussian $p_t(x_t)$, the log-density is:

$$\log p_t(x_t) = -\frac{1}{2}(x_t - \mu_t)^\top \Sigma_t^{-1} (x_t - \mu_t) + C \tag{29}$$

where C is a constant independent of x_t. Taking the gradient ∇_{x_t}:

$$\nabla_{x_t} \log p_t(x_t) = -\Sigma_t^{-1}(x_t - \mu_t) \tag{30}$$

Substituting μ_t and Σ_t back gives the explicit score for the Gaussian prior:

$$\nabla_{x_t} \log p_t(x_t) = -(\bar{\alpha}_t \Sigma_{data} + (1 - \bar{\alpha}_t) I)^{-1} (x_t - \sqrt{\bar{\alpha}_t} \mu_{data}) \tag{31}$$

D.2 FunDPS Architecture and details

We train an unconditional diffusion model with the ddpmpp-uno backbone, which combines a DDPM++-style U-Net with a U-shaped neural-operator design. The network is configured with base channel width $\mathcal{C}_0^{\mathrm{FunDPS}} = 64$ and per-resolution channel multipliers $\boldsymbol{\mathcal{C}}^{\mathrm{FunDPS}} = \{1, 2, 4, 4\}$. Self-attention is applied at resolution $\mathcal{R}_{\mathrm{attn}}^{\mathrm{FunDPS}} = \{8\}$, and each resolution stage contains $\mathcal{N}_{\mathrm{blk}}^{\mathrm{FunDPS}} = 4$ blocks. For the neural-operator component, we set the feature multiplier $\mathcal{M}^{\mathrm{FunDPS}} = 0.5$ and operator rank $\mathcal{K}^{\mathrm{FunDPS}} = 0.1$. Function-space (GRF) noise is modeled using an RBF kernel with

correlation scale $\mathcal{L}^{\mathrm{FunDPS}} = 0.05$. We use EDM preconditioning, which specifies the noise parameterization and denoising loss across noise levels. The model is trained without class conditioning.

Training uses 100,000 Darcy triplets at resolution 64×64, constructed from 1,000 permeability fields with 100 randomly sampled forcing designs per field. The batch size is 64. We use an Adam optimizer with learning rate $\eta^{\mathrm{FunDPS}} = 10^{-4}$ and a learning-rate ramp-up over the first $\mathcal{T}_{\mathrm{warm}}^{\mathrm{FunDPS}} = 5$ million images. We maintain an exponential moving average of model parameters with half-life $\mathcal{H}^{\mathrm{FunDPS}} = 0.5$ million images. Dropout is set to $\mathcal{P}^{\mathrm{FunDPS}} = 0.13$.

For parameter estimation, we perform posterior sampling with DPS. Unless stated otherwise, we use a fixed random seed $\mathcal{Z}^{\mathrm{DPS}} = 1$ and run $\mathcal{N}_{\mathrm{step}}^{\mathrm{DPS}} = 500$ sampling iterations with an EDM noise schedule $(\underline{\sigma}^{\mathrm{DPS}}, \overline{\sigma}^{\mathrm{DPS}}, \varrho^{\mathrm{DPS}}) = (0.002, 80, 7)$. We initialize latents with RBF noise using correlation scale $\mathcal{L}^{\mathrm{DPS}} = 0.05$. The RBF Gaussian field is used to initialize the diffusion latent. It is distinct from the measurement noise in (47). For observation conditioning, we use the squared mismatch between the PDE forward-model prediction and the observed solution values at the acquired pixels, evaluated in the normalized space. The corresponding likelihood guidance is defined in Equation (47).

E Proofs for Error Analysis

Throughout this section, define the finite-terminal-time MINDE residual by

$$\mathcal{R}_T(\boldsymbol{\xi}) := \mathbb{E}_{\mathbf{y}}\left[D_{\mathrm{KL}}(q_T \| p_T)\right].$$

Proposition 4 (Vanishing MINDE terminal residual)**.** *Under the forward-diffusion conditions ensuring vanishing terminal KL divergence in Franzese et al. (2024), and assuming* $\mathrm{EIG}_k(\boldsymbol{\xi}) < \infty$*,*

$$\mathcal{R}_T(\boldsymbol{\xi}) \longrightarrow 0 \qquad \textit{as } T \to \infty.$$

Proof. For each hypothetical observation $\mathbf{y}$, the terminal KL divergence between the posterior marginal q_T and the prior marginal p_T satisfies

$$D_{\mathrm{KL}}(q_T \| p_T) \longrightarrow 0 \qquad \text{as } T \to \infty,$$

under the stated diffusion conditions (Franzese et al., 2024). Since $0 \le D_{\mathrm{KL}}(q_T \| p_T) \le D_{\mathrm{KL}}(q_0 \| p_0)$ and $\mathbb{E}_{\mathbf{y}}[D_{\mathrm{KL}}(q_0 \| p_0)] = \mathrm{EIG}_k(\boldsymbol{\xi}) < \infty$, dominated convergence allows taking the expectation over $\mathbf{y}$, proving the result. □

Define the population DPS steering cost by

$$\widehat{\mathcal{J}}_T(\boldsymbol{\xi}) := \frac{1}{2}\mathbb{E}_{\mathbf{y}}\left[\int_0^T a(t)^2\, \mathbb{E}_{\widehat{q}_t}\left[\|\boldsymbol{g}_{\mathrm{DPS}}\|_2^2\right] \mathrm{d}t\right].$$

Here, $\widehat{q}_t$ is obtained by applying the forward-noising kernel to the clean DPS posterior-sample distribution $\widehat{q}_0$.

Proposition 5 (Exact MINDE–DPS decomposition)**.** *Under the regularity assumptions of Proposition 1,*

$$\mathrm{EIG}_k(\boldsymbol{\xi}) - \widehat{\mathcal{J}}_T(\boldsymbol{\xi}) = \mathcal{R}_T(\boldsymbol{\xi}) + \mathcal{B}_{\mathrm{int}}(\boldsymbol{\xi}) + \mathcal{B}_{\mathrm{path}}(\boldsymbol{\xi}), \tag{32}$$

where

$$\mathcal{B}_{\mathrm{int}}(\boldsymbol{\xi}) := \frac{1}{2}\mathbb{E}_{\mathbf{y}}\left[\int_0^T a(t)^2\, \mathbb{E}_{q_t}\left[\|\boldsymbol{g}_t^\star\|_2^2 - \|\boldsymbol{g}_{\mathrm{DPS}}\|_2^2\right] \mathrm{d}t\right],$$

and

$$\mathcal{B}_{\mathrm{path}}(\boldsymbol{\xi}) := \frac{1}{2}\mathbb{E}_{\mathbf{y}}\left[\int_0^T a(t)^2\, \left(\mathbb{E}_{q_t} - \mathbb{E}_{\widehat{q}_t}\right)\left[\|\boldsymbol{g}_{\mathrm{DPS}}\|_2^2\right] \mathrm{d}t\right].$$

Proof. The MINDE identity gives

$$\mathrm{EIG}_k(\boldsymbol{\xi}) = \mathcal{R}_T(\boldsymbol{\xi}) + \mathcal{J}_T(\boldsymbol{\xi}),$$

where $\mathcal{J}_T$ is the steering-cost functional evaluated using the exact marginals q_t and guidance $\boldsymbol{g}_t^\star$. Adding and subtracting $\mathbb{E}_{q_t}[\|\boldsymbol{g}_{\mathrm{DPS}}\|_2^2]$ inside the steering-cost integral yields (32). □

For the error-rate proof, fix $T < \infty$ and $\boldsymbol{\xi}$, and assume $\mathrm{EIG}_k(\boldsymbol{\xi}) < \infty$. Let $A_T := \int_0^T a(t)^2 \, \mathrm{d}t < \infty$. Assume that Girsanov's theorem applies to the exact and approximate reverse samplers initialized from q_T and p_T, respectively, with common diffusion coefficient $a(t)$ and drift difference $a(t)^2(\boldsymbol{g}_t^\star - \boldsymbol{g}_{\mathrm{DPS}})$. Assume that both guidance fields are bounded in norm by a deterministic constant $G < \infty$, uniformly in latent state and time and for predictive-almost every $\mathbf{y}$. The constants in the bounds may depend on T and G, but not on M.

Proof of Theorem 3. By Proposition 5, it remains to bound $\mathcal{B}_{\mathrm{int}}(\boldsymbol{\xi})$ and $\mathcal{B}_{\mathrm{path}}(\boldsymbol{\xi})$.

First,

$$\|\boldsymbol{g}_t^\star\|_2^2 - \|\boldsymbol{g}_{\mathrm{DPS}}\|_2^2 = (\|\boldsymbol{g}_t^\star\|_2 - \|\boldsymbol{g}_{\mathrm{DPS}}\|_2)\,(\|\boldsymbol{g}_t^\star\|_2 + \|\boldsymbol{g}_{\mathrm{DPS}}\|_2)\,.$$

The reverse triangle inequality and Cauchy–Schwarz imply

$$|\mathcal{B}_{\mathrm{int}}(\boldsymbol{\xi})| \le G\sqrt{A_T}\,\varepsilon_{\mathrm{DPS}}(\boldsymbol{\xi}, T).$$

Next, applying Girsanov's theorem to the reverse path measures, followed by data processing at the clean endpoint and under the common forward-noising kernel, yields

$$D_{\mathrm{KL}}(q_t\|\widehat{q}_t) \le D_{\mathrm{KL}}(q_0\|\widehat{q}_0) \le D_{\mathrm{KL}}(q_T\|p_T) + \frac{1}{2}\int_0^T a(s)^2\,\mathbb{E}_{q_s}\left[\|\boldsymbol{g}_s^\star - \boldsymbol{g}_{\mathrm{DPS}}\|_2^2\right] \mathrm{d}s.$$

Using $\mathrm{TV}(\mu,\nu) = \sup_A |\mu(A) - \nu(A)|$, the bound $0 \le \|\boldsymbol{g}_{\mathrm{DPS}}\|_2^2 \le G^2$, Pinsker's inequality, and Jensen's inequality give

$$|\mathcal{B}_{\mathrm{path}}(\boldsymbol{\xi})| \le \frac{G^2 A_T}{2\sqrt{2}}\sqrt{\mathcal{R}_T(\boldsymbol{\xi}) + \frac{1}{2}\varepsilon_{\mathrm{DPS}}(\boldsymbol{\xi}, T)^2}.$$

By Proposition 4, $\mathcal{R}_T(\boldsymbol{\xi}) \to 0$ as $T \to \infty$. We retain the terminal residual explicitly for finite T.

Combining the bounds above with (32) gives

$$\left|\mathrm{EIG}_k(\boldsymbol{\xi}) - \widehat{\mathcal{J}}_T(\boldsymbol{\xi})\right| = O\Big(\varepsilon_{\mathrm{DPS}}(\boldsymbol{\xi}, T) + \sqrt{\mathcal{R}_T(\boldsymbol{\xi})} + \mathcal{R}_T(\boldsymbol{\xi})\Big)\,. \tag{33}$$

Finally, for fixed T and a fixed trained model, assume that the estimator averages M independent and identically distributed contributions with mean $\widehat{\mathcal{J}}_T(\boldsymbol{\xi})$ and finite variance. Its variance is therefore of order M^{-1}, and Chebyshev's inequality implies

$$\widehat{\mathrm{EIG}}_{k,M}(\boldsymbol{\xi}) - \widehat{\mathcal{J}}_T(\boldsymbol{\xi}) = O_p(M^{-1/2}).$$

Combining this Monte Carlo error with (33) proves (13). □

F Experimental Details

F.1 Source Localization via Sequential Sensor Placement

We follow the four-dimensional source-localization benchmark of Ivanova et al. (2021). The two source locations are independently distributed as $\boldsymbol{\theta}_j \sim \mathcal{N}(0, \mathbf{I}_2)$, $j = 1, 2$, so that $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^\top, \boldsymbol{\theta}_2^\top)^\top \in \mathbb{R}^4$. Each run comprises ten sequential measurements. For a sensor location $\boldsymbol{\xi} \in \mathbb{R}^2$,

$$\mu(\boldsymbol{\theta}, \boldsymbol{\xi}) = b + \sum_{j=1}^{2} \frac{\alpha_j}{m + \|\boldsymbol{\theta}_j - \boldsymbol{\xi}\|_2^2}, \qquad \log y \mid \boldsymbol{\theta}, \boldsymbol{\xi} \sim \mathcal{N}\big(\log \mu(\boldsymbol{\theta}, \boldsymbol{\xi}), \sigma^2\big)\,, \tag{34}$$

where $\alpha_1 = \alpha_2 = 1$, $b = 0.1$, $m = 10^{-4}$, and $\sigma = 0.5$. For MINDE-DPS, we use a pool of 200 sensor locations sampled from $\mathcal{N}(0, \mathbf{I}_2)$.

Prior model. We use an analytic flow for the Gaussian source-location prior. The prior model is evaluated in closed form and requires no neural-network training.

Design scoring. Candidate scoring uses $M = 8$ predictive observations and $S = 300$ posterior particles per predictive observation. We use $N_{\text{diff}} = 100$ reverse steps and 500 particles for full-history posterior updates. The steering cost is accumulated over the reverse steps.

F.2 Constant Elasticity of Substitution

For the second example, we follow Huang et al. (2025) in testing performance on the CES problem. This economics problem considers the utility difference of two baskets of goods described by parameters ρ and α. The utility of a basket is defined as

$$U(z) = \left(\sum_{i=1}^{3} z_i^{\rho} \alpha_i \right)^{\frac{1}{\rho}}. \tag{35}$$

The utility difference η depends on a third parameter u as

$$\eta \sim \mathcal{N}(u \cdot (U(z) - U(z')), u^2 \cdot \tau^2 \cdot (1 + \|z - z'\|)^2). \tag{36}$$

The observation is $y = \text{clip}(\text{sigmoid}(\eta), \epsilon, 1 - \epsilon)$, where $\epsilon = 2^{-22}$. The priors are $\rho \sim \text{Uniform}(0, 1)$, $\boldsymbol{\alpha} \sim \text{Dirichlet}(1, 1, 1)$, and $\log u \sim \mathcal{N}(1, 3^2)$, where $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3)$. Each design comprises two three-good baskets, $\boldsymbol{\xi} = (z, z') \in [0, 100]^6$. We use ten sequential acquisitions and set $\tau = 0.005$.

Prior model. We train a flow model on 4,000 CES parameter samples for 19,000 epochs and keep the trained prior fixed throughout sequential design.

Design scoring. Candidate scoring uses $M = 8$ predictive observations and $S = 256$ posterior particles per predictive observation. We use $N_{\text{diff}} = 100$ reverse steps and 512 particles for full-history posterior updates. The steering cost is accumulated over the reverse steps.

F.3 Darcy Flow

The unknown parameter is a spatially varying permeability field $\chi(x)$. Observations are noisy samples of the steady-state solution u under the design-dependent forcing or observation operator:

$$\mathbf{y} = \mathcal{F}(\boldsymbol{\theta}, \boldsymbol{\xi}) + \mathbf{n}, \qquad \mathbf{n} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_y). \tag{37}$$

Observation noise is independent across acquired pixels, with standard deviation 0.001.

Diffusion prior. Following FunDPS (Yao et al., 2025), we train a three-channel diffusion model jointly on (f, χ, u) using EDM preconditioning. Our BOED procedure uses only the permeability channel χ for posterior inference and design scoring. We keep the pretrained model fixed throughout sequential design and use $N_{\text{diff}} = 500$ reverse steps for posterior sampling. During training and inference, fields are preprocessed by applying $\log(\max(\cdot, 1 \times 10^{-10}))$ pointwise and then normalized using fixed dataset statistics. For inverse parameter field estimation problems, we follow the Diffusion-PDE style posterior sampling procedure (Huang et al., 2024): starting from an RBF-initialized latent, we iteratively denoise with DPS guidance, using the squared mismatch between the PDE forward-model prediction and the observed solution values. More implementation details can be found in Appendix D.2.

Forward problem. We consider steady Darcy flow on the unit square $\Omega = (0, 1)^2$:

$$-\nabla \cdot \big(\chi(x) \nabla u(x)\big) = f(x) \quad \text{in } \Omega, \tag{38}$$

with homogeneous Dirichlet boundary conditions

$$u(x) = 0 \quad \text{on } \partial\Omega. \tag{39}$$

Here, $\chi(x) > 0$ is the permeability field, $f(x)$ is the forcing, and $u(x)$ is the steady-state solution.

Random coefficient model. The spatially varying coefficient $\chi(x)$ is generated from a Gaussian random field $\gamma(x)$ with Matérn-type covariance defined via the inverse of a squared elliptic operator:

$$\gamma \sim \mathcal{N}\big(0,\, (-\Delta + \tau^2 I)^{-2}\big). \tag{40}$$

In practice, we sample a discrete white noise ξ and solve the SPDE

$$(-\Delta + \tau^2 I)\,\gamma = \xi, \tag{41}$$

using standard finite elements. The positive parameter field is then defined through a lognormal map,

$$\chi(x) = \exp(\beta\,\gamma(x)), \tag{42}$$

and is projected to the finite element space used in assembly. In our implementation, the stored coefficient array corresponds to the evaluated $\chi(x)$ used in the PDE solve.

Source term with random location. The forcing term is a localized Gaussian centered at a random location (x_c, y_c):

$$f(x) = A \exp\left(-\frac{(x - x_c)^2 + (y - y_c)^2}{\varepsilon^2}\right), \tag{43}$$

where the source center is sampled uniformly in an interior window,

$$(x_c, y_c) \sim \mathrm{Unif}\big([m, 1-m]^2\big), \tag{44}$$

with fixed margin $m > 0$ to reduce boundary-induced variability under Dirichlet conditions. We use $\tau = 3$ and $\beta = 10$ for the permeability prior, and set $A = 1$, $\varepsilon = 0.5$, and $m = 0.2$ for the forcing. Thus, $\chi(x) = \exp(10\gamma(x))$ and source locations are sampled uniformly from $[0.2, 0.8]^2$.

Discretization. We use continuous piecewise-linear Lagrange finite elements (CG1) on a uniform triangular mesh generated by `UnitSquareMesh` with resolution $N \times N$. After imposing the Dirichlet boundary conditions, the discrete system is

$$K(\chi)\,\mathbf{u} = \mathbf{F}, \tag{45}$$

where $K(\chi)$ is the stiffness matrix and $\mathbf{F}$ is the load vector associated with f.

Dataset generation and storage. We independently sample 1,000 permeability fields. For each field, we sample 100 forcing designs by drawing random source locations and solve the steady Darcy equation. This produces 100,000 triplets (f, χ, u), all used for training. Evaluation uses a separate test set of 2,000 examples that are not used for training. Each stored example contains the permeability field $\chi(x)$, the forcing field $f(x)$, the steady-state solution $u(x)$, and the source location (x_c, y_c). To ensure a consistent ordering for machine learning, all scalar fields are exported on an $N \times N$ cell-centered grid:

$$x_i = \frac{i + 1/2}{N}, \qquad y_j = \frac{j + 1/2}{N}, \qquad \texttt{array}[j, i] \leftrightarrow (x_i, y_j), \tag{46}$$

which avoids ambiguity due to finite-element degree-of-freedom ordering.

Coordinate convention. The saved arrays follow the "physics" convention in which the row index j increases with the physical y coordinate. When visualizing with image routines, we use `origin="lower"` (or equivalently flip the vertical axis) to align array indices with physical coordinates. If an image-style convention is desired (row 0 at the top), one may flip arrays along the first axis and transform the stored source coordinate via $y_c \leftarrow 1 - y_c$.

Observation guidance and design scoring. For Darcy flow, $\mathcal{F}(\boldsymbol{\theta}, \boldsymbol{\xi})$ denotes the PDE forward prediction restricted to the observed solution values specified by the design. Under the Gaussian observation model, the data misfit is

$$\mathcal{L}_{\mathrm{obs}}(\boldsymbol{\theta}; \mathbf{y}, \boldsymbol{\xi}) = \frac{1}{2} \left\| \boldsymbol{\Sigma}_y^{-1/2} \left(\mathcal{F}(\boldsymbol{\theta}, \boldsymbol{\xi}) - \mathbf{y} \right) \right\|_2^2. \tag{47}$$

The DPS plug-in guidance is $-\nabla_{\boldsymbol{\theta}_t} \mathcal{L}_{\mathrm{obs}}(\hat{\boldsymbol{\theta}}_0(\boldsymbol{\theta}_t, t); \mathbf{y}, \boldsymbol{\xi})$. For isotropic observation noise, this is proportional to the gradient of the squared residual between the PDE solution and the observations. All acquired observations are used for posterior conditioning, while the EIG score uses the guidance contribution of the newly proposed observation.

Source-position comparison and evaluation. For Darcy source-position design, ALINE and MINDE-DPS use the same dataset of 1,000 permeability fields and the same pool of 100 candidate source locations. We compare MINDE-DPS with the 8×8 and 64×64 ALINE configurations. The ALINE (8×8) configuration uses an 8×8 grid subsampled from the original 64×64 parameter field.

For test case r, let $\chi^{(r)}$ be the true permeability field and $\widehat{\chi}^{(r,s)}$ its s-th reconstruction, both expressed on the original physical scale. We compute the relative L_2 error of each reconstruction and average these errors within each case:

$$e_{r,s} = \frac{\|\widehat{\chi}^{(r,s)} - \chi^{(r)}\|_2}{\|\chi^{(r)}\|_2}, \qquad e_r = \frac{1}{S_{\mathrm{rec}}} \sum_{s=1}^{S_{\mathrm{rec}}} e_{r,s}.$$

We use $S_{\mathrm{rec}} = 5$ reconstruction samples per test case and acquisition stage. The reported mean and standard deviation are computed across the 20 case-level values e_r. Reconstruction samples within a case are not treated as independent test cases. Stage 1 uses the first acquired measurement, while Stage 2 uses both acquired measurements.

Monte Carlo settings for source-position design. For each candidate source location, the outer predictive expectation is approximated using $M = 100$ prospective observations. For each prospective observation, we use $S = 3$ DPS trajectories with $N_{\mathrm{diff}} = 500$ reverse steps. We subsample 100 states from each saved trajectory, defining 99 adjacent intervals. After excluding the final 10 intervals in the saved trajectory order, 89 intervals are retained for steering-cost estimation. At each retained interval, we generate $L_{\mathrm{dir}} = 20$ direct-noising states and $L_{\mathrm{traj}} = 20$ perturbations of the saved DPS anchor. Both components are evaluated using the same pretrained denoiser.

Pixel-index observation patches. The observation patches are specified by pixel-index ranges on the 64×64 grid. Each selected patch adds its observed pixels to the accumulated observation operator. We retain the same pixel partition throughout a sequential run.

F.4 Benchmark Evaluation

Reported benchmark results. The EIG entries in Table 1 are evaluated using sPCE after ten acquisitions. The source-localization results for Random, SG-BOED, and DAD are taken from Ivanova et al. (2021), with $L_{\mathrm{PCE}} = 5 \times 10^5$ contrastive samples and mean $\pm$ one standard error over 4096 histories. The CES baseline results are taken from Huang et al. (2025), with $L_{\mathrm{PCE}} = 10^7$ and 95% confidence intervals over 2000 histories, except VPCE, which uses 200 histories. We reevaluate ALINE under our source-localization setting. This ALINE result and all MINDE-DPS results are from our experiments.

Evaluation units. Our source-localization and CES results are evaluated over 2000 complete BOED histories per method, using the same contrastive-sample budgets as the corresponding benchmarks. We report one standard error for source localization and 95% confidence intervals for CES. For CES, the interval is $\bar{I} \pm 1.96\, s_I/\sqrt{2000}$, where $\bar{I}$ and s_I are the mean and sample standard deviation of the final cumulative sPCE estimates across the 2,000 independent complete histories.

G Additional Results and Ablations

G.1 Source-Localization Posterior and EIG Validation

We visualize posterior samples conditioned on the accumulated measurements in Figure 2. For the EIG maps, we use an initial pool of 200 candidate sensor locations sampled from a standard Gaussian distribution and exclude previously selected locations at later steps. The selected sensor is marked with a star in Figure 4. True source locations are shown for reference. We compare candidate rankings within each step rather than interpreting the absolute map scales across steps. Table 3 reports agreement with the sPCE reference estimates.

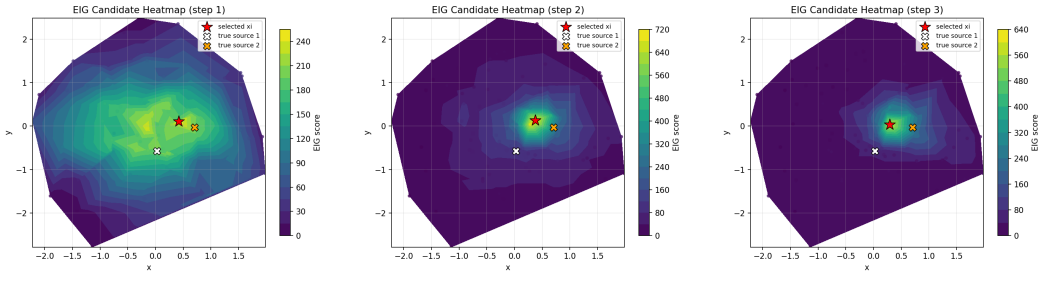

(a) Step 1

(b) Step 2

(c) Step 3

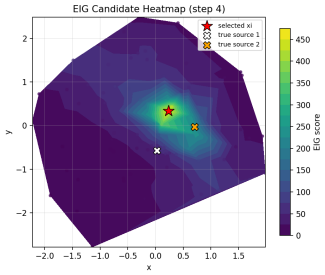

(d) Step 4

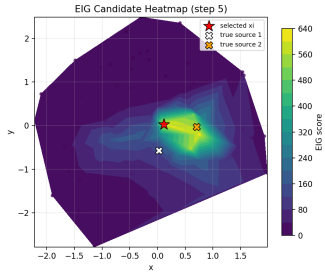

(e) Step 5

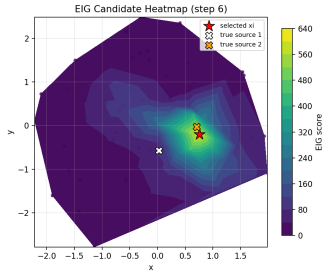

(f) Step 6

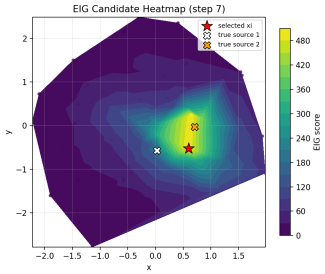

(g) Step 7

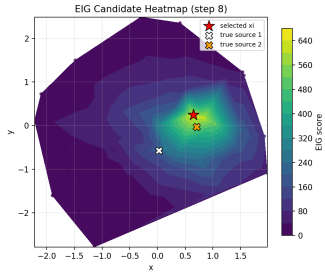

(h) Step 8

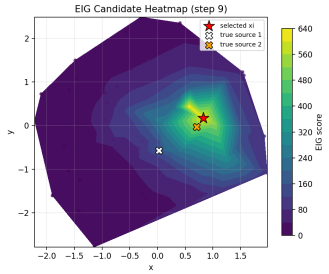

(i) Step 9

Figure 4: EIG surrogate heatmaps for steps $k = 1, \ldots, 9$ in the source-location problem. The initial candidate pool contains 200 locations sampled from a standard Gaussian distribution; previously selected locations are excluded at later steps. The selected design ξ_k is marked with a star; true source locations are shown for reference. Scores are shown up to a design-independent scaling (constants that do not depend on ξ are omitted), so selection is based on within-step ranking.

Table 3: Per-step calibration metrics, sequential BOED (K=10). The pool decreases from 200→191 as the picked sensor is removed. Spearman ρ is the global rank correlation between MINDE-DPS and the sPCE reference estimates. "MINDE top-1 %ile" is the rank, in the sPCE reference distribution, of the design that MINDE-DPS would have selected. Simple regret is $(\mathrm{sPCE}_{\max} - \mathrm{sPCE}(\xi^{\star}_{\mathrm{minde}}))/(\mathrm{sPCE}_{\max} - \mathrm{sPCE}_{\min})$.

Step	Pool	Spearman ρ	Kendall τ	MINDE top-1 %ile	Regret (%)
1	200	0.855	0.674	80.0	5.3
2	199	0.891	0.743	76.4	17.7
3	198	0.888	0.724	74.2	33.7
4	197	0.886	0.707	85.8	25.5
5	196	0.889	0.707	96.9	5.8
6	195	0.893	0.712	90.3	23.7
7	194	0.911	0.739	91.8	20.7
8	193	0.889	0.705	93.8	11.2
9	192	0.924	0.757	96.9	12.0
10	191	0.912	0.739	97.9	17.4
static (t=0, 200 cand.)		0.853	0.677	71.5	7.8

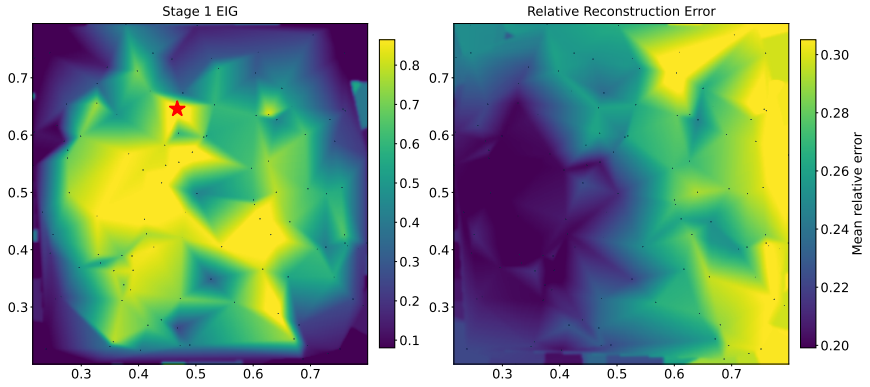

Figure 5: EIG map and reconstruction-error map on the Darcy-flow benchmark. Higher-EIG regions generally coincide with lower DPS reconstruction error. Red stars mark the coordinates with the largest EIG.

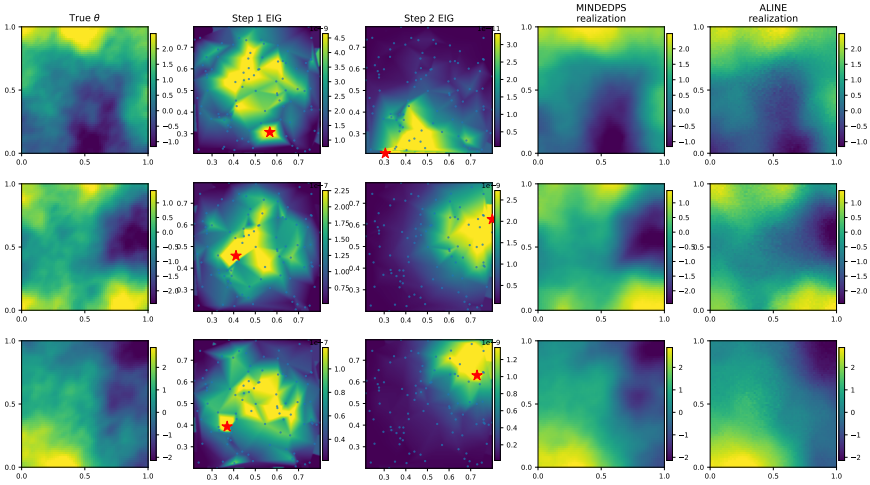

Figure 6: Two-stage sequential EIG for Darcy source-position design. Stage 1 scores candidate forcing locations. Stage 2 recomputes EIG after the first measurement. Red stars mark the location selected at each stage.

G.2 Additional Darcy Results

EIG map and reconstruction error. We keep the underlying parameter field fixed across test cases, and generate different observation sets by changing the source locations at random. For each test instance, we create 10 Gaussian-noise versions of the observation, and for each one, we draw 5 posterior samples using DPS-guided sampling. We then average the resulting estimates to build an EIG map, and compare it with a metric map computed from the average DPS reconstruction error. In Figure 5, high-EIG regions generally correspond to lower reconstruction error.

Sequential Source-position Design. Figure 6 provides additional examples of two-stage source-position design.

Fixed-budget comparison. In the main text, we evaluate sequential patch acquisition on a test set where the parameter field is fixed across instances. Here we consider a more challenging setting in which the full Darcy triplet (f, χ, u) varies across test instances, while the goal remains to reconstruct the parameter field from partial and noisy observations. In this section, we compare an EIG-guided design against a random baseline under an identical observation budget. For each test instance, we partition the 64×64 field into a 5×5 grid (25 non-overlapping patches) and select

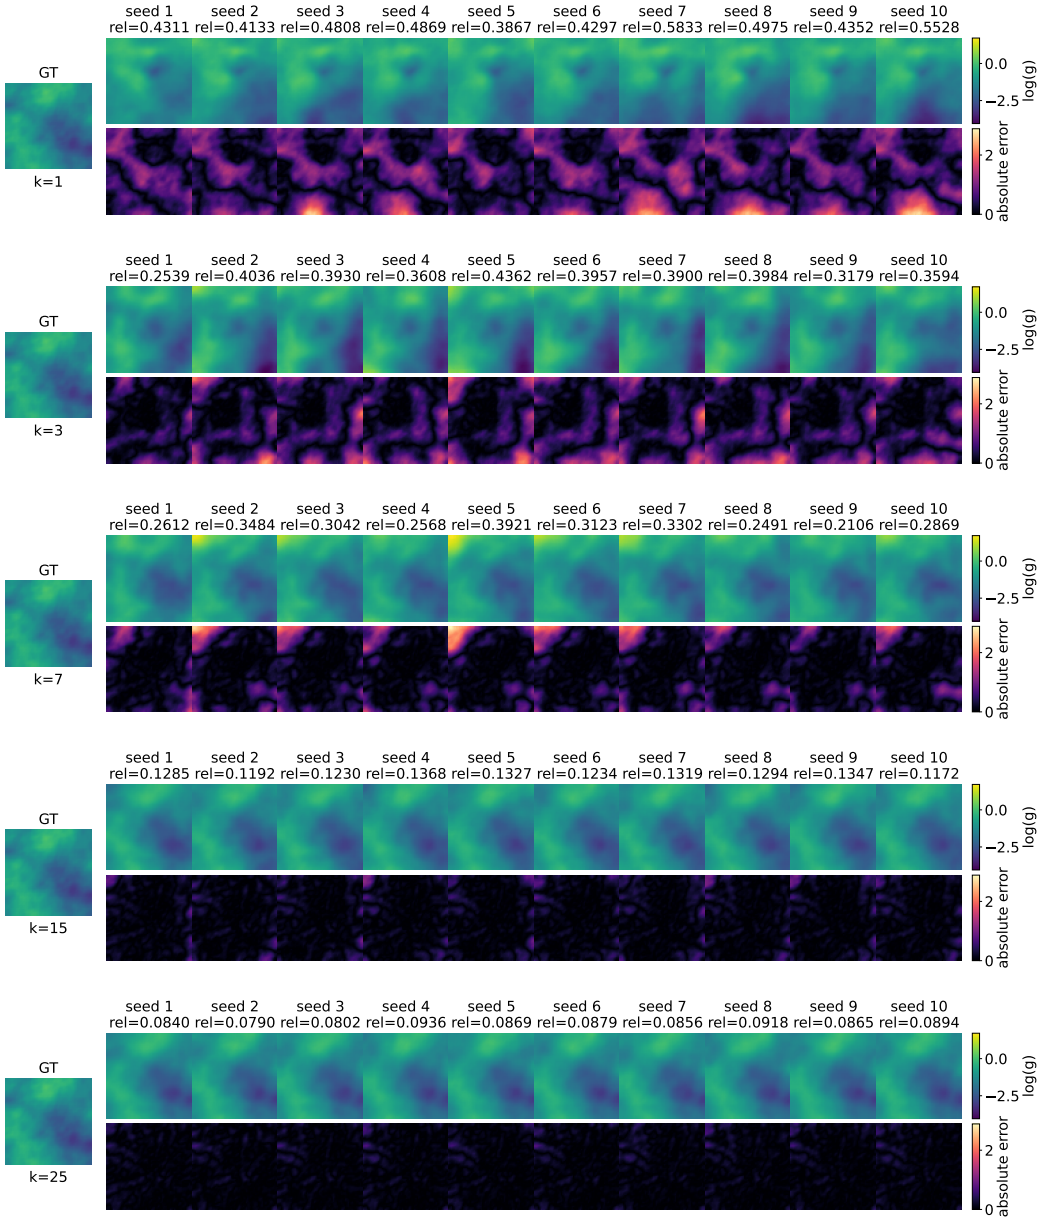

Figure 7: K-step DPS with adaptive patch observations on the Darcy-flow benchmark. The 64×64 field is divided into a 5×5 grid of pixel-index patches. After each acquisition, EIG is recomputed under the updated posterior. For each observation budget k, the top row shows reconstructions and the bottom row shows absolute-error maps.

exactly $k = 3$ observed patches. **EIG-guided:** We acquire three patches sequentially. After each acquisition, we update the posterior using the accumulated observations and recompute the EIG of the remaining patches before selecting the next one. **Random:** We acquire three distinct patches uniformly at random, using the same observation and reconstruction budgets. For each test example and selection strategy, we fix the selected patches and the acquired observations, and perform five independent DPS reconstruction runs with different sampler seeds. The acquisition sequence is not rerun for these repetitions.

Figure 8 shows representative results. In both examples, EIG-guided patch selection achieves a lower mean relative reconstruction error across the five sampler seeds than random selection under the same observation budget.

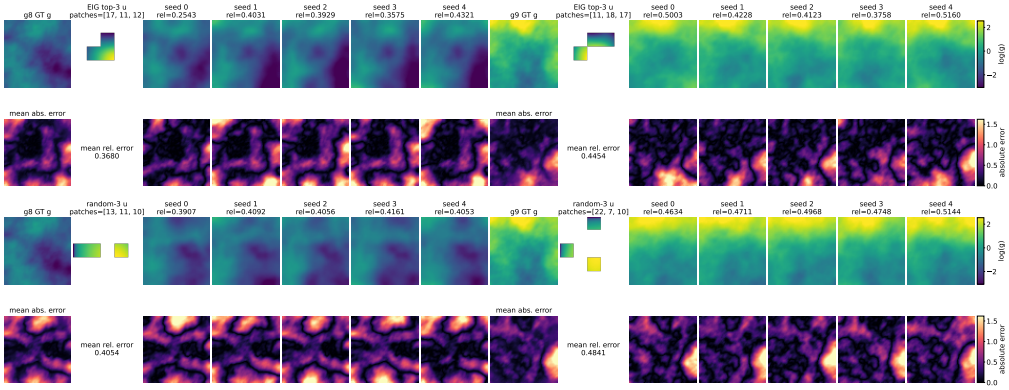

Figure 8: Qualitative comparison with a fixed budget of $k = 3$ patches. For each of two Darcy test examples, we show five independent DPS reconstructions under fixed acquired observations, using different sampler seeds. Adaptive EIG-guided selection and random selection are compared through the reconstructed permeability fields and absolute-error maps. The repetitions reflect reconstruction variability conditional on each selected observation set.

G.3 Guidance-Strength and Diffusion-Step Sensitivity

These ablations are conducted on the source-localization benchmark. The reported θ MSE measures the mean squared error between inferred and true source coordinates.

Table 4: Per-location sensitivity of MINDE-DPS to DPS guidance strength ζ and the number of reverse diffusion steps N_{diff}. We evaluate two fixed locations, using 500 test samples per location and 5 posterior samples per datum. Lower θ MSE indicates better inverse performance.

ζ	N_{diff}	Location	EIG mean	EIG median	θ MSE mean
0.1	50	$(0.304, 0.419)$	0.136	0.084	2.548
		$(0.469, 0.411)$	0.179	0.111	2.738
0.5	50	$(0.304, 0.419)$	0.116	0.063	0.453
		$(0.469, 0.411)$	0.155	0.086	0.555
1.0	50	$(0.304, 0.419)$	0.107	0.052	0.167
		$(0.469, 0.411)$	0.144	0.070	0.255
0.1	100	$(0.304, 0.419)$	0.123	0.075	1.519
		$(0.469, 0.411)$	0.163	0.099	1.690
0.5	100	$(0.304, 0.419)$	0.103	0.049	0.158
		$(0.469, 0.411)$	0.139	0.066	0.244
1.0	100	$(0.304, 0.419)$	0.094	0.039	0.095
		$(0.469, 0.411)$	0.127	0.050	0.149
0.1	200	$(0.304, 0.419)$	0.113	0.063	0.649
		$(0.469, 0.411)$	0.151	0.086	0.761
0.5	200	$(0.304, 0.419)$	0.092	0.038	0.088
		$(0.469, 0.411)$	0.125	0.049	0.140
1.0	200	$(0.304, 0.419)$	0.063	0.018	0.074
		$(0.469, 0.411)$	0.085	0.023	0.107