BISIP: BIAS-GUIDED STOCHASTIC SEARCH FOR SIMULATION-BASED INVERSE PROBLEMS

Anonymous authors
Paper under double-blind review

ABSTRACT

Complex simulations, common in science and engineering, lead to challenging inverse problems. Bayesian methods, in particular deep learning-based ones, are widely used but are limited by data scarcity, computational cost of simulations, and distribution shift during testing. To address these challenges, we develop a stochastic search-based method to iteratively refine the inverse estimate. We use forward-consistency to estimate model bias, which, together with the model variance, guides the active acquisition of training samples to improve inverse estimation during testing. We evaluate the proposed approach on inverse problems of increasing complexity, including linear inverse with heteroscedastic noise, inverse Laplace transform, and non-linear reconstruction of neutral hydrogen distribution in the early universe from sparse cosmological observations. Across the evaluated settings, our method consistently improves inverse estimation compared with the considered active-learning and sequential inference baselines. In the cosmological reconstruction task, BISIP achieves the lowest estimation error among the evaluated methods while requiring substantially fewer simulator evaluations.

1 INTRODUCTION

Computer simulation of physical systems is common in science and engineering. Given a cause u, these simulations generate an observation $y = \mathcal{F}(u)$. The inverse problem, inferring u given y, is fundamental to broad scientific domains, including digital image processing (Blackledge, 2005; Song et al., 2022; Moliner et al., 2023), astronomy (Craig & Brown, 1985), geophysics (Virieux & Operto, 2009), optical microscopy (Mahmood et al., 2008), and medical imaging (Lustig et al., 2008). These inverse problems are inherently challenging because the forward map $\mathcal{F}$ is often non-linear, high-dimensional, and computationally expensive. Furthermore, many inverse problems are ill-posed, requiring prior knowledge or regularisation to obtain stable solutions. We focus specifically on simulation-based inverse problems (Cranmer et al., 2020; Deistler et al., 2025), where $\mathcal{F}$ can be executed as a forward simulator, but the underlying likelihood $p(y \mid u)$ is implicit and intractable.

To tackle ill-posedness, classical approaches (Kaipio & Somersalo, 2004) incorporate handcrafted regularisers, such as smoothness priors, sparsity constraints, or physics-based differential constraints (Raissi et al., 2019). While these methods stabilise inversion, they yield over-smoothed point estimates with little insight into reconstruction confidence. Bayesian frameworks instead model the full posterior $p(u \mid y) \propto p(y \mid u)\, p(u)$ using methods, such as variational inference (Blei et al., 2017), MCMC (Brooks et al., 2011), and normalizing flows (Papamakarios et al., 2021). Some recent approaches focus on strengthening $p(u)$ by pre-training expressive generative priors, such as diffusion (Wu et al., 2024; Zhao et al., 2026) or flow (Erbach et al., 2025) models, on curated offline datasets. However, these offline approaches hinge on the critical assumption that the prior $p(u)$ is known in advance and remains static between training and deployment; in many real-world scientific applications, this assumption breaks down (Ovadia et al., 2019). Representative ground truth fields u are scarce, and models trained offline routinely encounter severe distribution shifts at deployment, i.e., $p_{\text{train}}(u) \neq p_{\text{test}}(u)$.

Another line of research, also known as sequential neural posterior estimation, uses active learning principles to iteratively acquire new data $(u, \mathcal{F}(u))$ at test time to refine the inverse model or the estimated posterior (Griesemer et al., 2024; Papamakarios & Murray, 2016). Reducing the simula-

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.

tion cost of the forward map $\mathcal{F}$ is a major challenge, for which these methods acquire u for training either by sampling from the estimated posterior $\hat{p}(u|y)$ or by using heuristic rules.

However, these approaches struggle to scale with dimensionality. Moreover, they do not leverage forward simulator explicitly. In this work, we frame the problem as a stochastic search problem aimed at minimising bias in posterior estimation. We propose estimating the bias via local linearisation of the forward map. We use bias as an indicator of the model's inconsistency and, together with the posterior, to obtain $(u, \mathcal{F}(u))$ samples for active learning. Our approach, BISIP (Bias-Guided Stochastic Search for Simulation-Based Inverse Problems), reduces computational cost and scales well while handling shifted priors for moderately ill-posed inverse problems.

We provide an extensive evaluation across inverse problems of increasing complexity, viz., linear inverse, inverse Laplace transform, and a high-dimensional field reconstruction problem, namely cosmological tomography. BISIP demonstrates consistent improvements over existing baselines, including state-of-the-art performance on the field reconstruction task.

2 PRELIMINARIES

2.0.1 PROBLEM FORMULATION

Let $u \in \mathcal{U} \subseteq \mathbb{R}^{d_u}$ denote the unknown causal parameters and $y \in \mathcal{Y} \subseteq \mathbb{R}^{d_y}$ denote the observed measurements. The forward model $\mathcal{F} : \mathcal{U} \to \mathcal{Y}$ maps parameters to observations according to

$$y = \mathcal{F}(u) + \boldsymbol{\varepsilon}, \qquad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{u})) \tag{1}$$

where $\boldsymbol{\Sigma}(u) \in \mathbb{R}^{d_y \times d_y}$ is a heteroscedastic noise covariance, allowing measurement uncertainty to depend on $\boldsymbol{u}$. Given observations $\boldsymbol{y}$, the inverse problem at hand seeks to recover the unknown $\boldsymbol{u}$ by estimating the posterior distribution

$$p(u \mid y) \propto p(y \mid u)\, p(u) \tag{2}$$

2.1 PROBABILISTIC INVERSE MODEL

To approximate this posterior at scale, we use the available training data $D = \{(u_i, y_i)\}_i$ to learn an amortised inverse network

$$q_\phi : \mathcal{Y} \to \mathcal{U} \times \mathbb{R}_+^{d_u},$$

that maps an observation directly to the parameters of a factored Gaussian approximation,

$$q_\phi(u|y) = \mathcal{N}\left(u; \mu_\phi(y), \mathrm{diag}(\sigma_\phi^2(y))\right) \tag{3}$$

The network is trained by minimising the β-NLL objective (Seitzer et al., 2022),

$$\mathcal{L}_\beta(\phi) = \mathbb{E}_{(u,y)\sim\mathcal{D}}\left[\frac{1}{2}\sum_{d=1}^{d_u} \mathrm{sg}\big(\sigma_{\phi,d}^2(y)\big)^\beta \left(\frac{(u_d - \mu_{\phi,d}(y))^2}{\sigma_{\phi,d}^2(y)} + \log \sigma_{\phi,d}^2(y)\right)\right], \tag{4}$$

where $\mathrm{sg}(\cdot)$ denotes the stop gradient operator. We tune $\beta \in [0.4, 0.8]$ per problem.

3 METHOD

Neural inverse model estimates $q_\phi(u|y)$ as an approximation to $p(u|y)$. However, it does not model inconsistency arising from the model's limitations. This uncertainty can be overcome by providing sufficient training data, and provided the model is expressive enough to learn from it. To acquire such informative samples, active learning strategies have been devised (Papamakarios et al., 2019; Griesemer et al., 2024). Given the costs of simulations and the high dimensionality of u and y, it is not practical to minimise the inconsistency for the entire $\mathcal{U}$. Nevertheless, if we could generate sufficient data in the neighbourhood of the true u^*, we could minimise the inconsistency locally.

Below, we describe how we estimate posterior bias, which we then use as an indicator of inconsistency for acquiring samples for active learning.

3.1 Estimating Posterior Bias

We now derive an estimate of posterior bias in the neighbourhood of the true u, denoted by u^*. Let $y^* = \mathcal{F}(u^*)$. For u in a neighbourhood of u^*, a first-order expansion of $\mathcal{F}$ about u gives

$$y^* = \mathcal{F}(u^*) = \mathcal{F}(u) + \nabla\mathcal{F}(u)(u^* - u) + \mathcal{O}(\|u^* - u\|^2) \tag{5}$$

$$\implies \|u - u^*\|^2 \approx (y - y^*)^\top \left(\nabla\mathcal{F}(u)^{-1}\right)^\top \nabla\mathcal{F}(u)^{-1}(y - y^*), \quad y = \mathcal{F}(u) \tag{6}$$

The expected estimation error can be decomposed as

$$\mathbb{E}_{u\sim q_\phi(u|y^*)}\|u - u^*\|^2 = \mathrm{Tr}(\mathrm{Cov}_{q_\phi}) + \mathbb{E}_{u\sim q_\phi(u|y^*)}\|\mathbb{E}_{q_\phi}[u] - u^*\|^2 \tag{7}$$

For a Gaussian posterior, $q_\phi(u \mid y^*)$ with mean μ_ϕ and diagonal covariance matrix $\sigma_\phi^2 = [\sigma_d^2]_{d=1}^{d_y}$, it simplifies to

$$\mathbb{E}_{u\sim q_\phi(u|y^*)}\|u - u^*\|^2 = \sum_d \sigma_d^2 + \mathbb{E}_{u\sim q_\phi(u|y^*)}\|\mu_\phi - u^*\|^2 \tag{8}$$

The term $S(\phi, y^*) := \mathbb{E}_{u\sim q_\phi(u|y^*)}\|\mu_\phi - u^*\|^2$ is the bias term that is an indicator of inconsistency. It can be estimated from above equations as

$$S(\phi, y^*) = \mathbb{E}_{u\sim q_\phi}\left[(y - y^*)^\top \left(\nabla\mathcal{F}(u)^{-1}\right)^\top \nabla\mathcal{F}(u)^{-1}(y - y^*)\right] - \sum_d \sigma_d^2 \tag{9}$$

Theorem 3.1. *$S(\phi, y^*) \to 0$ when the model is unbiased and the posterior mass is sufficiently concentrated, for a moderately ill-posed inverse problem.*

Proof. Substitute $(y - y^*) = \nabla\mathcal{F}(u)(u - u^*)$ to get $S(\phi, y^*) = \mathbb{E}_{u\sim q_\phi}[\|u - u^*\|^2] - \sum_d \sigma_d^2 + \mathcal{O}(\|u - u^*\|^3)$. When the model is unbiased, $\mathbb{E}_{u\sim q_\phi(u|y^*)} = u^*$. Higher-order terms tend to 0 when the posterior mass is sufficiently concentrated. □

If $\mathcal{F}$ is known and differentiable, one can directly obtain the Jacobian $J := \nabla\mathcal{F}$. Otherwise, as in many practical situations, we approximate $\mathcal{F}$ by a neural surrogate $\hat{\mathcal{F}}_\psi$, a lightweight, differentiable model, trained using the available data D by minimising $\mathbb{E}_{(u,y)\sim D}[\|y - \hat{\mathcal{F}}(u)\|^2]$. The Jacobian is then obtained via autograd. For stability and when $d_u \neq d_y$, $\nabla\mathcal{F}^{-1}$ can be replaced with Tikhonov-regularised inverse $(J^\top J + \lambda I)^{-1}J^\top$, with $\lambda \to 0^+$.

In high-dimensional field reconstruction problems, computing J is computationally expensive, but $\mathcal{F}$ is local in dimensions (space-time), resulting in J being a band matrix, which can be inverted efficiently. Further, if $d_u = d_y$, the computations may be simplified by considering only the diagonal terms. Assuming that $\mathcal{F}$ is linear in the neighbourhood where the posterior is concentrated, each component of the diagonal of J^{-1} can be approximated as $(u_d - \bar{u}_d)/(y_d - \bar{y}_d)$, where $\bar{u}_d = \mathbb{E}_u[u_d]$ and $\bar{y}_d = \mathbb{E}_u[y_d]$. Thus, we get

$$S(\phi, y^*) = \sum_d s_d;\, s_d = \mathbb{E}_{u\sim q_\phi(u|y^*)}\left[\frac{(u_d - \bar{u}_d)^2}{(y_d - \bar{y}_d)^2 + \delta}\,(y_d - y_d^*)^2\right] - \sigma_d^2 \tag{10}$$

For numerical stability, we add a small $\delta > 0$ to the denominator when the local output variation $(y_d - \bar{y}_d)^2$ is small. To estimate s_d, we use Monte Carlo samples. Finite samples may yield smaller estimates, making $s_d < 0$; to avoid this irregularity, we clip s_d to 0. Further, to minimise computational costs, we use the surrogate forward model to compute $y = \hat{\mathcal{F}}(u)$. The stability of $S(\phi, y^*)$ under this approximation is discussed in Appendix A.1. We provide empirical evidence for the relationship of bias and $S(\phi, y^*)$ in Section 5.2.

3.2 Sequential Data Acquisition and Learning

We treat the inverse problem as a stochastic search problem (Gutmann & Corander, 2016), with the goal of finding u that minimises the reconstruction error $\|\mathcal{F}(u) - y^*\|$. The estimated bias guides this search by iteratively refining the estimated posterior $q_\phi(u|y)$. Now we describe how we use the

estimated bias $S(\phi, y^*)$ to first acquire y^* with a large bias and then sample u for forward simulation for further learning.

We operate in an active learning setting to minimise labelling costs. Given an initial labelled set $D_0 = \{(u_i, y_i)\}_{i=1}^{N_0}$ and an unlabelled test pool $\mathcal{Y}^* = \{y_j^*\}_{j=1}^{M}$, with $M \gg N_0$, each acquisition round scores every pool candidate under an acquisition function $\alpha(y^*; \phi) := S(\phi, y^*)$ and selects the top-k samples,

$$\mathcal{B}_t = \arg \max_{\mathcal{S} \subseteq \mathcal{Y}^*, |\mathcal{S}|=k} \sum_{y^* \in \mathcal{S}} \alpha(y^*; \phi). \tag{11}$$

Next, for the selected observations $y_j^* \in \mathcal{B}_t$, we need to acquire the candidate u. We sample u from a Gaussian with mean $\mu_\phi(y^*)$ and covariance matrix $\mathbb{E}_{u \sim q_\phi}[(u - u^*)(u - u^*)^\top]$. Using the diagonal Gaussian approximation along with equation 8 and equation 10, we can write the diagonal terms of this covariance matrix as $s_d + \sigma_d^2$. The bias term s_d widens the search region, encouraging exploration. Thus, $u_j \sim \mathcal{N}(\mu_\phi(y_j^*), \mathrm{diag}(s_d + \sigma_d^2))$.

These selected candidates u_j are then labelled via $\mathcal{F}$ and incorporated into the training set, $D_{t+1} = D_t \cup \{(u_j, \mathcal{F}(u_j)\}_{j \in \mathcal{S}_t}$. The process repeats for T acquisition rounds, with a total labelling budget of Tk, where k is the number of samples acquired in each round.

The updated training set D_{t+1} is used to update both q_ϕ and $\hat{\mathcal{F}}_\psi$. While the bias-based acquisition strategy helps the Bayesian stochastic search, we also make use of the differentiable forward surrogate $\hat{\mathcal{F}}_\psi$ to guide the search by updating q_ϕ to minimise the reconstruction error. In other words, ϕ is updated to minimise

$$\mathcal{L}_{\mathrm{f}}(\phi; \psi) = \mathbb{E}_{y \sim \mathcal{B}_t} \| \, y - \mathbb{E}_{u \sim q_\phi(\cdot|y)} \hat{\mathcal{F}}_\psi(u) \|^2 \tag{12}$$

while ψ is kept frozen. Overall, ϕ is updated to minimise $\mathcal{L}_\beta(\phi) + \gamma \mathcal{L}_{\mathrm{f}}(\phi; \psi)$, with $\gamma > 0$ being a hyperparameter.

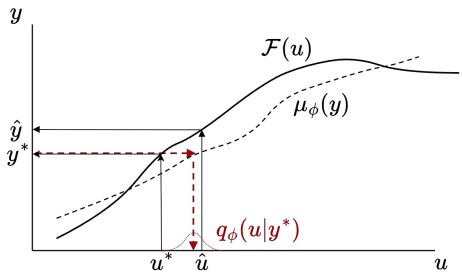


Figure 1: **Stochastic search via active acquisition:** True u^* (unknown) maps to y^* (known). The inverse model gives $q_\phi(u|y^*)$, which, along with estimated bias $S(\phi, y^*)$, generates a sample $\hat{u}$. If $S(\phi, y^*)$ is large, $\hat{u}$ may be selected to acquire a training pair $(\hat{u}, \hat{y} = \mathcal{F}(\hat{u}))$.

The algorithm for BISIP is summarised in Appendix (algorithm 1). Fig. 1 shows a schematic for a one-dimensional case with a continuous and differentiable $\mathcal{F}$. When the likelihood $q_\phi(u^*|y^*)$ is low, the samples $\hat{u} \sim q_\phi(u|y^*)$ indicate a large model bias $S(\phi, y^*)$. This, in turn, leads to sample acquisition from a wider neighbourhood, encouraging exploration. These samples help q_ϕ learn the inverse map better in a wider neighbourhood, iteratively converging to the true posterior.

Factors such as prior shift and data scarcity could lead the posterior to be confidently wrong. Such low-variance, high-bias posteriors would fail to acquire good samples, but $S(\phi, y^*)$ would be high in such cases and encourage acquisition of better samples.

4 Related Work

Inverse Problems Inverse problems can be treated as regression or classification problems, providing a point estimate of u, given y. This approach is agnostic to the forward map. Such approaches are common for inverse problems such as image reconstruction (Chung et al., 2023). Nevertheless, physical constraints could be used to regularise model training (Raissi et al., 2019). Bayesian approach could also be used to decouple prior and likelihood, which are modelled using a generative model, such as a diffusion model, conditioned on y (Chung et al., 2023; Erbach et al., 2025). It is possible to work with weak priors also (Jia et al., 2026) when the inverse problem is moderately ill-conditioned.

Inverse problems, where $p(y|u)$ is explicitly known, are studied under the name likelihood-based inference, while those, where $y = \mathcal{F}(u)$ is explicitly known, are categorised under simulation-based

inference. As an example of the former, Zhao et al. (2026) use a Bayesian approach with known $p(y|u)$ and a diffusion model to learn the prior for a cosmology inverse problem.

Simulation-Based Inference Simulation-based inference trains density estimators to approximate intractable posterior or likelihood distributions using simulated input-output pairs (Cranmer et al., 2020; Deistler et al., 2025; Lueckmann et al., 2021). Non-parametric Gaussian process (Villani et al., 2024) provides a rigorous framework for such problems, but fails to scale to high dimensions. Neural posterior estimation (Greenberg et al., 2019) learns the distribution using neural networks. While these methods enable fast amortised inference, static density estimators systematically fail to produce accurate reconstructions under test-time distribution shifts (Hermans et al., 2022; Huang et al., 2023). Sequential neural posterior estimation (SNPE) methods attempt to improve sample efficiency by targeting simulator queries across multiple rounds (Papamakarios et al., 2019), yet they still depend on a fixed prior and demand heavy simulator budgets at deployment. Active SNPE (Griesemer et al. (2024)) acquires new training samples $(u, \mathcal{F}(u))$ by sampling u that minimises epistemic uncertainty estimated by Monte Carlo dropout-based sampling of the posterior (Gal & Ghahramani, 2016; Kandasamy et al., 2015). In contrast, BISIP explicitly uses the reconstruction error and the forward simulator to estimate the bias in posterior estimation. We hypothesise that this bias is a better indicator of inconsistency than a perturbation-based uncertainty estimate. Moreover, (Griesemer et al., 2024) estimate the posterior using a flow model, which does not scale for high-dimensional problems such as field reconstruction.

Active Acquisition Strategies Active learning acquisition strategies span uncertainty-based, diversity-based, and gradient-based paradigms. *Posterior variance* (Kendall & Gal, 2017) scores candidates by predictive variance, while *Core-Set* (Sener & Savarese, 2018) pursues coverage directly in observation space, scoring each candidate $\boldsymbol{y} \in \mathcal{Y}$ by its distance to the nearest labelled point and building the batch via farthest-first traversal. *BADGE* (Ash et al., 2020) instead computes a gradient embedding per unlabelled sample and selects a diverse batch via k-means++ initialization in this space. Gaussian Processes offer a well-calibrated, non-parametric alternative, but independent per-dimension GPs (via GPyTorch (Gardner et al., 2018)) scale poorly with output dimensionality. BISIP resembles Bayesian optimisation (Gutmann & Corander, 2016), but targets improved posterior estimation rather than optimizing $\mathcal{F}$'s output.

Cosmological Field Reconstruction Cosmological field reconstruction infers latent physical fields from sparse observations. During reionization, radiation from the first galaxies transformed neutral intergalactic hydrogen into an ionised state, producing a spatially heterogeneous reionization field. Existing field-level approaches have mainly targeted the 21-cm signal from neutral hydrogen, using SegU-Net for neutral–ionised segmentation (Bianco et al., 2021), U-Nets for recovering signal lost to foreground filtering (Kennedy et al., 2024), and SERENEt for reconstructing the spatial 21-cm field (Bianco et al., 2025). Since direct 21-cm tomography remains observationally difficult, Lyα emitters (LAEs) offer a complementary observable, as neutral-hydrogen attenuation of Lyα emission makes their abundance and spatial distribution sensitive to local reionization structure (Santos, 2004; McQuinn et al., 2007; Mesinger & Furlanetto, 2008; Jensen et al., 2013). Building on this, Maitra et al. (2026a) used a 3D U-Net to reconstruct the neutral-hydrogen field from sparse galaxy observations. More broadly, ML-based inversion has addressed high-redshift matter-density reconstruction with 3D VAE–U-Nets (Maitra et al., 2026b), weak-lensing dark-matter mapping with U-Nets (Jeffrey et al., 2020), CMB-lensing reconstruction with ResUNets (Caldeira et al., 2019), density/velocity reconstruction from galaxy distributions with autoencoders (Ganeshaiah Veena et al., 2023), primordial-density reconstruction with CNNs (Shallue & Eisenstein, 2023), cluster-velocity inference with CNNs (Tanimura et al., 2022), and weak-lensing mass-map reconstruction with diffusion models (Boruah et al., 2025).

5 Experiments and Results

We evaluate our approach (BISIP) across simulation-based inverse problems of increasing complexity, detailed in the following subsections. In all the problems, to test the robustness under test-time prior shift, our training prior $p_{\text{train}}(u)$ differs explicitly from that during evaluation $p_{\text{test}}(u)$. Also, we use the standard evaluation metrics, namely, estimation error $E_{\text{est}} = \mathbb{E}_{(u,y)}[u - \mathbb{E}_{\hat{u}}[q_\phi(\hat{u}|y)]]$, negative log-likelihood $\text{NLL} = \mathbb{E}_{(u,y)}[-\log q_\phi(u|y)]$, and reconstruction error, $E_{\text{rec}} = \mathbb{E}_{(u,y)}[\|y - \mathcal{F}(\mathbb{E}_{\hat{u}}[q_\phi(\hat{u}|y)])\|^2]$.

Baselines To rigorously evaluate BISIP, we compare against a diverse set of baseline strategies spanning active learning, sequential simulation-based inference (SBI), and non-active inverse problem solvers. For active learning baselines, we compare against random acquisition (Rand) as a model-independent reference, MC-dropout disagreement (MCD) (Gal & Ghahramani, 2016), posterior variance uncertainty (PVU) (Kendall & Gal, 2017), BADGE (Ash et al., 2020), Coreset selection (Sener & Savarese, 2018) and GP (Gardner et al., 2018) In addition to conventional active learning methods, we compare against sequential SBI baselines that adapt the simulation distribution over successive rounds. Specifically, we utilise standard SNPE-C (Greenberg et al., 2019) and an adaptive sequential SBI baseline ASNPE (Griesemer et al., 2024). For non-active methods specifically tailored to inverse problems, we evaluate Physics-Informed Neural Networks (PINNs) (Raissi et al., 2019) and FLAIR (Erbach et al., 2025), which optimise solutions directly or leverage physics-constrained architectures without iterative active sample acquisition. All active learning and sequential SBI methods operate under a shared simulation budget across rounds, whereas non-active baselines are evaluated using samples acquired via random acquisition at the final acquisition budget, with all comparisons averaged consistently over multiple random seeds. All neural-based methods use identical network architectures. More details are in Appendix A.2, and implementation of PINN in A.3.1.We exclude FLAIR Erbach et al. (2025) from the low-dimensional benchmarks because these experiments are designed to isolate the effect of the inference and acquisition strategies under controlled forward models. FLAIR additionally introduces a high-capacity flow-based generative prior, which constitutes an additional inductive bias whose contribution would be difficult to disentangle from that of the inference procedure in these controlled settings. We therefore evaluate FLAIR in the cosmological reconstruction experiment, where a high-dimensional structured prior is more directly relevant. Implementation details for FLAIR (Erbach et al., 2025) on field reconstruction are provided in Section A.3.3.

5.1 Linear Inverse Problem

First, we consider a toy linear inverse problem, $y = Au + \boldsymbol{\epsilon}$, where $u \in \mathbb{R}^d$ is the latent signal, $A \in \mathbb{R}^{m\times d}$ the forward operator, and $\boldsymbol{\epsilon}$ additive zero-mean Gaussian noise. The forward operator A is generated via Singular Value Decomposition, $A = U\Sigma V^\top$, with singular values controlling the condition number $\kappa(A)$, which corresponds to problem ill-posedness. Three experiments are performed with $\kappa(A) = 10, 10^2, 10^3$. Sweeping $\kappa(A)$ systematically controls problem ill-posedness, as $\kappa(A)$ grows, the inverse becomes increasingly sensitive to perturbations in the observations.

The inverse network predicts posterior mean and variance and is trained to minimise $\mathcal{L} = \mathcal{L}_\beta(u, \hat{\mu}_\phi(y), \log\hat{\sigma}^2_\phi(y)) + \lambda\mathcal{L}_f(y, \mathcal{F}(u))$ using definitions in 4 and 12 for all active baseline methods. $\lambda = 0.08$ is used consistently throughout.

Results From Table 1, across all condition numbers, BISIP outperforms the baselines, achieving small estimation and reconstruction errors, especially for large κ. Both SNPE-C and ASNPE exhibit degraded performance on the considered linear inverse problem. Although both methods access $\mathcal{F}$, their inference prior, $p_{\text{train}}(x)$, is estimated from a finite initial training set and deviates from the true data-generating distribution. This prior mismatch is especially consequential in ill-conditioned settings, where weak observational constraints make the posterior heavily reliant on the assumed prior, leaving it inaccurate in poorly constrained directions. For ASNPE, this error compounds through the sequential acquisition process because the acquisition criterion depends on the variance of the current posterior density, early inaccuracies misguide subsequent simulation selections and hinder refinement. Further details of the implementation and data generation is in Appendix A.2.1. PINN shows a competitive performance but provides only a point estimate, rather than a posterior, and hence, NLL cannot be computed. Implementation for PINN is summarised in Appenndix A.3.1.

5.2 Inverse Laplace Transform

Laplace transform (Kaipio & Somersalo, 2004) is $(Lu)(s) = \int_0^\infty e^{-st}u(t)\,dt$, and the goal of inverse problem is to recover the unknown function $u(t)$ given its transform. This is a classic example of severe ill-posedness due to the smoothing nature of the Laplace kernel. The integral is approximated via Gauss–Legendre quadrature on $t \in [0, 5]$ with 40 nodes t_k and weights w_k, evaluated at 40 observation points s_j placed on a logarithmic grid, $s_j = 10^{-1+(j-1)/20}$. This yields the discretised forward model

$$\boldsymbol{y} = A\boldsymbol{u} + \epsilon, \qquad A_{jk} = w_k\, e^{-s_j t_k}, \tag{13}$$

Table 1: Comparison of acquisition methods across increasing ill-posedness (κ) of the linear inverse problem. Values are mean $\pm$ std over 10 seeds. Errors E_{est} and E_{rec} are scaled by $\times 10^3$.

Method	$\kappa = 10$			$\kappa = 100$			$\kappa = 1000$		
	NLL ↓	E_{est} ↓	E_{rec} ↓	NLL ↓	E_{est} ↓	E_{rec} ↓	NLL ↓	E_{est} ↓	E_{rec} ↓
BISIP	$\mathbf{-0.734 \pm 0.016}$	$\mathbf{32 \pm 1}$	28 ± 0	$\mathbf{-0.947 \pm 0.027}$	$\mathbf{20 \pm 1}$	15 ± 0	$\mathbf{-0.870 \pm 0.025}$	$\mathbf{26 \pm 1}$	$\mathbf{9 \pm 0}$
PVU	-0.435 ± 0.060	41 ± 4	30 ± 1	-0.313 ± 0.021	42 ± 2	16 ± 0	0.768 ± 0.178	31 ± 1	10 ± 0
Rand	-0.458 ± 0.028	33 ± 2	31 ± 0	-0.730 ± 0.032	43 ± 2	16 ± 0	-0.221 ± 0.089	33 ± 2	10 ± 0
MCD	-0.550 ± 0.007	43 ± 3	32 ± 1	-0.858 ± 0.008	38 ± 1	16 ± 0	-0.737 ± 0.017	33 ± 1	10 ± 0
GP	-0.382 ± 0.015	214 ± 3	$\mathbf{18 \pm 1}$	0.103 ± 0.001	404 ± 4	23 ± 1	0.174 ± 0.013	543 ± 15	17 ± 0
BADGE	-0.482 ± 0.013	42 ± 3	32 ± 1	-0.735 ± 0.032	54 ± 2	16 ± 0	-0.406 ± 0.043	42 ± 2	11 ± 0
Coreset	-0.625 ± 0.037	40 ± 3	31 ± 0	-0.528 ± 0.052	55 ± 2	16 ± 0	-0.094 ± 0.103	39 ± 4	11 ± 0
ASNPE	34.546 ± 8.573	320 ± 80	85 ± 16	49.968 ± 6.66	346 ± 60	40.357 ± 5.761	40 ± 5	360 ± 69	24 ± 4
SNPE-C	78.625 ± 4.038	554 ± 36	134 ± 8	80.051 ± 3.942	595 ± 31	65 ± 2	81.580 ± 4.353	614 ± 37	40 ± 3
PINN	–	38 ± 5	30 ± 0	–	23 ± 1	$\mathbf{14 \pm 0}$	–	35 ± 3	10 ± 0

with $A \in \mathbb{R}^{40\times40}$, $\boldsymbol{u} = [u(t_1), \ldots, u(t_{40})]^\top$, and $\boldsymbol{y} = [(Lu)(s_1), \ldots, (Lu)(s_{40})]^\top$. ϵ is isotropic zero-mean Gaussian noise. The operator is severely ill-conditioned ($\kappa(A) \approx 5.6 \times 10^{16}$), serving as an aggressive stress test for acquisition in the extreme ill-posed regime. The signal $u(t)$ is drawn from the triangular pulse family with its parameters sampled from different train and test distributions.

Results As shown in Table 2 and Figure 2, we observe a consistent scale difference between the E_{est} ($\sim 10^{-2}$) and E_{rec} ($\sim 10^{-4}$). This discrepancy arises from the ill-posed nature of the linear operator A: because A attenuates high-frequency and null-space components, errors in u space are heavily smoothed out in the y-space. While PINN minimises E_{rec} slightly better than BISIP, E_{est} captures unobservable directions that cannot be uniquely resolved from limited data alone, over which BISIP provides marginal gains. Full parameter sampling distributions, network architecture, and hyperparameter details are provided in Appendix A.2.2.

Table 2: Results on Inverse Laplace Transform Problem

Method	E_{est} ($\times 10^{-3}$)	NLL	E_{rec} ($\times 10^{-3}$)
BISIP	$\mathbf{9.1 \pm 0.4}$	$\mathbf{-2.6 \pm 0.05}$	2.5 ± 0.7
Rand	10.9 ± 0.8	-2.42 ± 0.08	4.3 ± 2.0
PVU	10.9 ± 1.0	-2.41 ± 0.08	4.7 ± 1.9
MCD	11.2 ± 0.7	-2.24 ± 0.04	4.3 ± 1.3
Coreset	11.4 ± 0.2	-2.42 ± 0.09	3.9 ± 0.7
BADGE	11.6 ± 0.9	-2.39 ± 0.06	3.4 ± 1.3
GP	12.2 ± 0.2	-2.0 ± 0.02	4.2 ± 0.4
PINN	9.8 ± 0.5	–	$\mathbf{1.1 \pm 0.02}$
SNPE-C	25.2 ± 1.5	73.26 ± 2.24	16.2 ± 6.1
ASNPE	12.8 ± 1.0	30.13 ± 9.66	5.3 ± 1.0

Empirical Validation of $S(\phi, y)$ To assess whether the $S(\phi, y)$ captures information about reconstruction bias, we compare it with the ground-truth latent error $B(y) = \|\mu_\phi(y) - u^*\|_2^2$ on this benchmark, where u^* is available. $S(\phi, y)$ achieves a Pearson correlation of 0.646 with B. More importantly for acquisition, it recovers 40.0%, 35.0%, and 41.25% of the highest-bias samples among the top $5\%, 10\%, 20\%$ candidates, respectively. These results provide empirical evidence that $S(\phi, y)$ supports for the claims made in section 3.1.

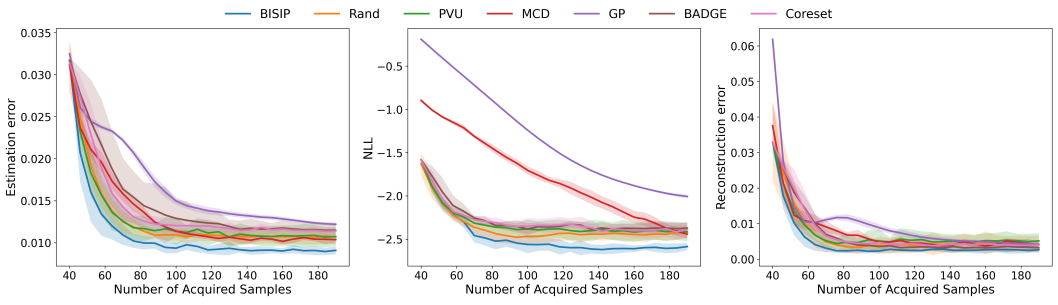


Figure 2: Performance of different acquisition strategies on the Inverse Laplace Transform.

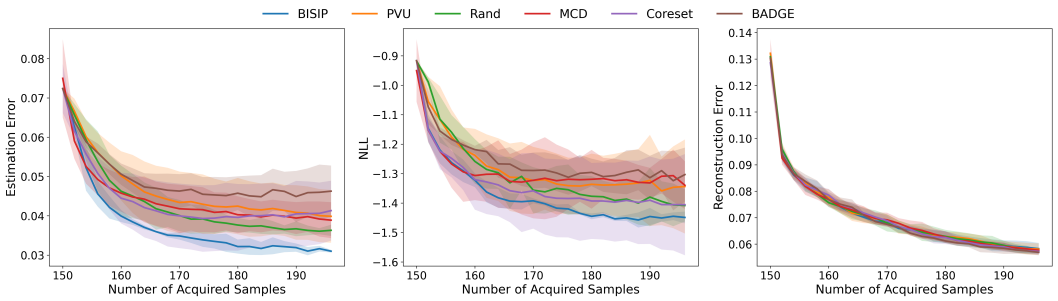


Figure 3: Performance of different acquisition strategies on the Reionization Field Reconstruction.

5.3 Cosmological Reionization Field Reconstruction

As a real-world, nonlinear instance of the general problem, we now consider an inverse problem from cosmology. One of the major goals in cosmology is to infer the patchy distribution of neutral hydrogen in the early Universe, which is otherwise difficult to observe directly. A promising approach is to use the spatial distribution of high-redshift galaxies, particularly objects known as Lyα emitters (LAEs), to tomographically reconstruct this hidden neutral hydrogen field. Here, we consider this reconstruction problem as a representative astrophysical inverse problem, using sparse galaxy observations to infer the underlying neutral hydrogen distribution.

During a key phase transition in the early Universe, known as the epoch of reionization, the Universe evolved from being filled largely with neutral hydrogen to becoming ionised by radiation from the first galaxies, producing a complex spatial pattern of neutral and ionised regions. Because high-redshift galaxies trace both the underlying matter distribution and the progress of reionization, their observed locations can be used to infer this otherwise difficult-to-observe three-dimensional neutral hydrogen field. We therefore cast this as an inverse problem in which the neutral hydrogen field is reconstructed from sparse galaxy tracers.

We consider reconstructing the three-dimensional neutral hydrogen field $u = x_{\mathrm{HI}}(r)$ during the epoch of reionization from these sparse, indirect observational tracers under a total labeling budget of $B = Tk$, with $T = 24$ acquisition rounds of $k = 2$ samples each, starting from an initial labelled set of $N_0 = 150$ examples. At each acquisition round, BISIP samples candidate x_{HI} fields from the bias-inflated posterior and queries the forward simulator on these newly generated samples to obtain additional training pairs. We report results averaged over 6 independent random seeds. Unlike the synthetic experiments above, $\mathcal{F}$ here is a full reionization simulation, expensive and non-differentiable, making this the setting $\mathcal{F}_\psi$ is designed for.

Simulator and data The unknown $u \in [0,1]^{64^3}$ reionization field and $y \in (\mathbb{R}^+)^{64^3}$ is a sparse tracer map derived from it. The observed tracer $\boldsymbol{y}$ is a sparse, binary-like proxy for the underlying continuous ionization field $\boldsymbol{u}$. The posterior estimation model and forward emulator $\hat{\mathcal{F}}_\psi$ follow the architectures described later, with $\hat{\mathcal{F}}_\psi$ retrained each round as the labelled set grows and the true simulator SiMPLE-GEN queried only for candidates selected via the S-score. Full details of the simulation pipeline, dataset sizes and training hyperparameters are provided in Appendix A.2.3.

Reconstruction Network The inverse network q_ϕ is a 3D U-Net-based conditional volumetric encoder-decoder featuring a stochastic latent bottleneck $\boldsymbol{z} \sim \mathcal{N}(\mu_z(y), \sigma_z^2(y))$. The decoder branches into mean and log-variance heads with skip connections, predicting a voxel-wise factored Gaussian posterior $q_\phi(u \mid y) = \mathcal{N}(\mu_\phi(y), \mathrm{diag}(\sigma_\phi^2(y)))$ from which samples $u^{(k)}$ are drawn directly. Architectural specifications and network parameters are detailed in Appendix A.2.3.

Forward Surrogate Network To avoid querying the computationally expensive simulator during acquisition, $\hat{\mathcal{F}}_\psi$ is instantiated as a SegResNet-based 3D encoder-decoder (Myronenko, 2018) implemented via MONAI. Adapted for volume-to-volume regression, it maps reconstructed ionization fields to predicted tracer observations.

Results Figure 3 shows results on the test set as a function of the number of labelled samples acquired, and Table 3 reports the corresponding metrics after final acquisition for direct comparison across all methods. Among the methods evaluated in this high-dimensional setting, BISIP achieves

Table 3: Evaluation metrics for Ionization field Reconstruction across acquisition strategies and baseline methods at the full labelling budget. Values are mean ± std over 6 seeds.

Method	E_{est} ($\times 10^{-2}$) ↓	NLL ↓	E_{rec} ($\times 10^{-2}$) ↓
BISIP	**3.57 ± 0.08**	**-1.3133 ± 0.1693**	**5.60 ± 0.16**
Rand	3.94 ± 0.11	-1.2489 ± 0.1589	6.15 ± 0.63
MCD	4.08 ± 0.08	-1.2518 ± 0.1246	5.77 ± 0.05
PVU	4.12 ± 0.11	-1.2194 ± 0.1584	5.80 ± 0.25
Coreset	4.46 ± 0.16	-1.1283 ± 0.1561	5.88 ± 0.15
BADGE	5.02 ± 0.12	-1.1612 ± 0.0736	5.67 ± 0.15
TORRCH	10.25 ± 0.39	-0.6415 ± 0.0168	13.17 ± 0.57
PINN	5.35 ± 0.80	–	7.29 ± 0.73
FLAIR	10.98 ± 11.23	–	13.49 ± 10.13

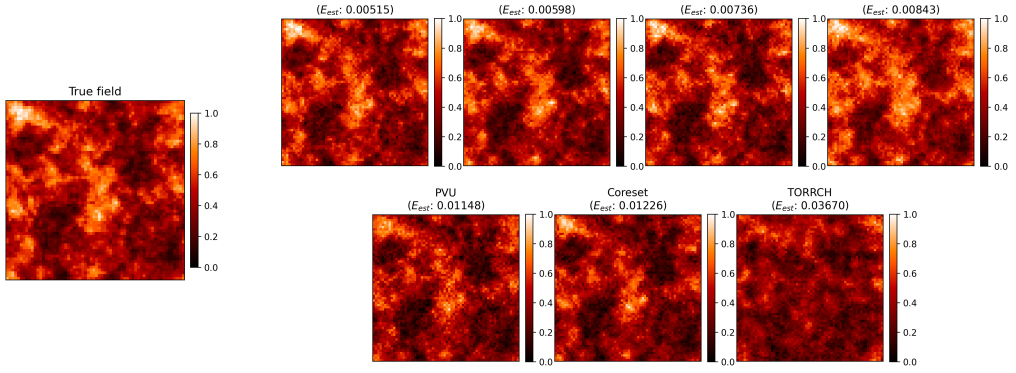


Figure 4: Ionization Field Reconstruction ranked by Fidelity to True Field

the lowest estimation error at the final acquisition budget. SNPE-C and ASNPE are not included in this experiment because their architectures and computational requirements do not readily scale to the $\mathbb{R}^{64^3}$ field reconstruction setting under the available simulation budget also considering prior misspecification. BISIP achieves 9.39% and 12.5% lower estimation error at the final budget than Rand and MCD, respectively. For a representative test sample, Figure 4 shows the reconstructed field $\hat{u}$ under each method alongside the corresponding estimation error E_{est}. BISIP and MCD yield sharper structural definition and capture the high-intensity hot spots reasonably well, whereas reconstructions from TORRCH (Maitra et al., 2026a) and Rand are comparatively smoother. BISIP achieves superior performance overall, better preserving the fine-scale structure of the true field relative to competing acquisition strategies. FLAIR, as a non-active baseline, underperforms in our setting given the pretrained generative prior is subject to substantial distributional shift from the target cosmological ionization fields, and this mismatch likely limits the effectiveness of flow-based regularization, particularly given the sparse, highly non-linear nature of our observations.

6 Conclusion

We presented BISIP as a stochastic search method to iteratively improve the posterior estimate in a simulation-based inverse problem. At each iteration, new samples are selected that minimise the model bias, estimated by local linearisation of the forward map. Our empirical findings demonstrate that BISIP outperforms popular single-point or fixed-posterior benchmarks as well as state-of-the-art active learning-based benchmarks for standard inverse problems. It also sets a new benchmark on inferring neutral hydrogen distribution in the early universe, a high-dimensional field reconstruction problem in cosmology. Future work includes extending the current framework to incorporate known observational uncertainties and to problems where the dimensionalities of cause and observation are large and unequal. Moreover, the current method estimates inconsistency for sample acquisition; its theoretical connections to uncertainty could inform further development.

AI USE STATEMENT

In this work, we used generative AI tools to search relevant literature during the initial review phase and to refine portions of the manuscript for clarity and tone. We did not use generative AI tools for research idea generation, algorithm design, experimental execution, or synthetic data generation. We take full responsibility for the final content of this work, including all text, mathematical formulations, research claims, and code.

REPRODUCIBILITY STATEMENT

While source code is not provided with this initial submission, comprehensive efforts have been made to ensure full reproducibility through the text. A complete description of our network architectures, layer dimensions, and initialization schemes is detailed in Appendix provided. All optimisation settings, including learning rate schedules, batch sizes, and hardware configurations, are explicitly listed. Pseudocode mapping out the core training loop is provided in Algorithm 1.

REFERENCES

Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *International Conference on Learning Representations (ICLR)*, 2020.

Michele Bianco, Sambit K. Giri, Ilian T. Iliev, and Garrelt Mellema. Deep learning approach for identification of H II regions during reionization in 21-cm observations. *Monthly Notices of the Royal Astronomical Society*, 505:3982–3997, 2021. doi: 10.1093/mnras/stab1518.

Michele Bianco, Sambit K. Giri, Rohit Sharma, Tianyue Chen, Shreyam Parth Krishna, Chris Finlay, Viraj Nistane, Philipp Denzel, Massimo De Santis, and Hatem Ghorbel. Deep learning approach for identification of H II regions during reionization in 21-cm observations - III. Image recovery. *Monthly Notices of the Royal Astronomical Society*, 541:234–250, 2025. doi: 10.1093/mnras/staf973.

Jonathan M. Blackledge. Digital image processing: Mathematical and computational methods. 2005.

David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

Supranta S. Boruah, Michael Jacob, and Bhuvnesh Jain. Diffusion-based mass map reconstruction from weak lensing data. *Physical Review D*, 111:083542, 2025. doi: 10.1103/PhysRevD.111.083542.

R. J. Bouwens, G. D. Illingworth, P. A. Oesch, M. Trenti, I. Labbé, L. Bradley, M. Carollo, P. G. van Dokkum, V. Gonzalez, B. Holwerda, M. Franx, L. Spitler, R. Smit, and D. Magee. UV luminosity functions at redshifts $z \sim 4$ to $z \sim 10$: 10,000 galaxies from HST legacy fields. *The Astrophysical Journal*, 803(1):34, 2015. doi: 10.1088/0004-637X/803/1/34.

Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, 2011.

João Caldeira, W. L. K. Wu, Brian Nord, Camille Avestruz, Shubhendu Trivedi, and K. T. Story. DeepCMB: Lensing reconstruction of the cosmic microwave background with deep neural networks. *Astronomy and Computing*, 28:100307, 2019. doi: 10.1016/j.ascom.2019.100307.

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 (ICLR)*, 2023.

Youngseog Chung, Ian Char, Han Guo, Jeff Schneider, and Willie Neiswanger. Uncertainty toolbox: an open-source library for assessing, visualizing, and improving uncertainty quantification. *arXiv preprint arXiv:2109.10254*, 2021.

I. J. D. Craig and J. C. Brown. *Inverse Problems in Astronomy: A Guide to Inversion Strategies for Remotely Sensed Data*. Adam Hilger Ltd., Accord, MA, 12 1985.

Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

Michael Deistler, Janne Boelts, Peter Steinbach, Richard Moss, Thomas Moreau, Marcel Gloeckler, Pedro Luiz Coelho Rodrigues, Alexander Linhart, et al. Simulation-based inference: A practical guide. *arXiv preprint arXiv:2508.12939*, 2025.

Mark Dijkstra and J. Stuart B. Wyithe. An empirical study of the relationship between lyα and UV-selected galaxies: do theorists and observers "select" the same objects? *Monthly Notices of the Royal Astronomical Society*, 419(4):3181–3193, 2012. doi: 10.1111/j.1365-2966.2011.19958.x.

Julius Erbach, Dominik Narnhofer, Andreas Dombos, Bernt Schiele, Jan Eric Lenssen, and Konrad Schindler. Solving inverse problems with flair. 2025. URL https://arxiv.org/abs/2506.02680.

Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1050–1059. PMLR, 2016.

Punyakoti Ganeshaiah Veena, Robert Lilow, and Adi Nusser. Large-scale density and velocity field reconstructions with neural networks. *Monthly Notices of the Royal Astronomical Society*, 522: 5291–5307, 2023. doi: 10.1093/mnras/stad1222.

Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

David Greenberg, Marcel Nonnenmacher, and Jakob Macke. Automatic posterior transformation for likelihood-free inference. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2404–2414. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/greenberg19a.html.

Sam Griesemer, Defu Cao, Zijun Cui, Carolina Osorio, and Yan Liu. Active sequential posterior estimation for sample-efficient simulation-based inference. *Advances in Neural Information Processing Systems*, 37:127907–127936, 2024.

Michael U. Gutmann and Jukka Corander. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *JOURNAL OF MACHINE LEARNING RESEARCH*, 17, 2016. ISSN 1532-4435.

Per Christian Hansen. Discrete inverse problems: Insight and algorithms. 2010. URL https://api.semanticscholar.org/CorpusID:120990370.

Joeri Hermans, Arnaud Delaunoy, François Rozet, Antoine Wehenkel, Volodimir Begy, and Gilles Louppe. A crisis in simulation-based inference? beware, your posterior approximations can be unfaithful. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL https://openreview.net/forum?id=LHAbHkt6Aq.

Daolang Huang, Ayush Bharti, Amauri H Souza, Luigi Acerbi, and Samuel Kaski. Learning robust statistics for simulation-based inference under model misspecification. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=STrXsSIEiq.

Niall Jeffrey, François Lanusse, Ofer Lahav, and Jean-Luc Starck. Deep learning dark matter map reconstructions from DES SV weak lensing data. *Monthly Notices of the Royal Astronomical Society*, 492:5023–5029, 2020. doi: 10.1093/mnras/staa127.

Hannes Jensen, Peter Laursen, Garrelt Mellema, Ilian T. Iliev, Jesper Sommer-Larsen, and Paul R. Shapiro. On the use of Ly-alpha emitters as probes of reionization. *Monthly Notices of the Royal Astronomical Society*, 428:1366–1381, 2013. doi: 10.1093/mnras/sts116.

Jing Jia, Wei Yuan, Sifan Liu, Liyue Shen, and Guanyang Wang. Weak diffusion priors can still achieve strong inverse-problem performance. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=fdkSA4F0lN.

Jari Kaipio and Erkki Somersalo. *Statistical and Computational Inverse Problems*, volume 160. Springer Science & Business Media, 2004.

Kirthevasan Kandasamy, Jeff Schneider, and Barnabás Póczos. Bayesian active learning for posterior estimation. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pp. 3605–3611. AAAI Press, 2015. ISBN 9781577357384.

Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Jacob Kennedy, Jonathan Colaço Carr, Samuel Gagnon-Hartman, Adrian Liu, Jordan Mirocha, and Yue Cui. Machine-learning recovery of foreground wedge-removed 21-cm light cones for high-z galaxy mapping. *Monthly Notices of the Royal Astronomical Society*, 529:3684–3698, 2024. doi: 10.1093/mnras/stae760.

Jan-Matthis Lueckmann, Jan Boelts, David Greenberg, Pedro Goncalves, and Jakob Macke. Benchmarking simulation-based inference. In Arindam Banerjee and Kenji Fukumizu (eds.), *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 343–351. PMLR, 13–15 Apr 2021. URL https://proceedings.mlr.press/v130/lueckmann21a.html.

Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing mri. *IEEE Signal Processing Mgazine*, 25(2):72–82, 2008.

M. T. Mahmood, T. S. Choi, and S. B. Kim. Pca-based method for 3d shape recovery of microscopic objects from image focus using discrete cosine transform. *Journal of Microscopy*, 232(3):514–525, 2008. doi: 10.1002/jemt.20635.

Soumak Maitra, Girish Kulkarni, Shikhar Asthana, James S. Bolton, Martin G. Haehnelt, and Laura Keating. The lyman-α emitter bispectrum as a probe of reionization morphology. *Monthly Notices of the Royal Astronomical Society*, 542(1):486–507, 2025. doi: 10.1093/mnras/staf1262.

Soumak Maitra, Girish Kulkarni, Vipul Arora, Matteo Viel, Shikhar Asthana, James S. Bolton, Martin G. Haehnelt, and Laura Keating. Torrch: Tomographic reconstruction of the reionization of cosmic hydrogen with lyα emitters and non-lyα-selected galaxies, 2026a. URL https://arxiv.org/abs/2601.16263.

Soumak Maitra, Matteo Viel, and Girish Kulkarni. Deepchart: Mapping the 3d dark matter density field from lyα forest surveys using deep learning. *Monthly Notices of the Royal Astronomical Society*, 549(2):stag851, 06 2026b. ISSN 0035-8711. doi: 10.1093/mnras/stag851. URL https://doi.org/10.1093/mnras/stag851.

Matthew McQuinn, Lars Hernquist, Matias Zaldarriaga, and Suvendra Dutta. Studying reionization with Ly-alpha emitters. *Monthly Notices of the Royal Astronomical Society*, 381:75–96, 2007. doi: 10.1111/j.1365-2966.2007.12085.x.

Andrei Mesinger and Steven R. Furlanetto. Ly-alpha emitters during the early stages of reionization. *Monthly Notices of the Royal Astronomical Society*, 386:1990–2002, 2008. doi: 10.1111/j.1365-2966.2008.13039.x.

Andrei Mesinger, Steven Furlanetto, and Renyue Cen. 21cmFAST: A fast, semi-numerical simulation of the high-redshift 21-cm signal. *Monthly Notices of the Royal Astronomical Society*, 411 (2):955–972, 2011. doi: 10.1111/j.1365-2966.2010.17731.x.

Eloi Moliner, Jaakko Lehtinen, and Vesa Välimäki. Solving audio inverse problems with a diffusion model. In *ICASSP 2023 – IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 1–5. IEEE, 2023.

Steven G. Murray, Bradley Greig, Andrei Mesinger, Julian B. Muñoz, Yuxiang Qin, Jaehong Park, and Catherine A. Watkinson. 21cmFAST v3: A python-integrated c code for generating 3d realizations of the cosmic 21cm signal. *Journal of Open Source Software*, 5(54):2582, 2020. doi: 10.21105/joss.02582.

Andriy Myronenko. 3d mri brain tumor segmentation using autoencoder regularization. In *International MICCAI Brainlesion Workshop*, pp. 311–320. Springer, 2018.

Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

George Papamakarios and Iain Murray. Fast ε-free inference of simulation models with bayesian conditional density estimation. *Advances in neural information processing systems*, 29, 2016.

George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 837–848. PMLR, 16–18 Apr 2019. URL https://proceedings.mlr.press/v89/papamakarios19a.html.

George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.

Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

Michael R. Santos. Probing reionization with Lyman-alpha emission lines. *Monthly Notices of the Royal Astronomical Society*, 349:1137–1152, 2004. doi: 10.1111/j.1365-2966.2004.07594.x.

Maximilian Seitzer, Arash Tavakoli, Dimitrije Antic, and Georg Martius. On the pitfalls of heteroscedastic uncertainty estimation with probabilistic neural networks. In *International Conference on Learning Representations*, April 2022. URL https://openreview.net/forum?id=aPOpXlnV1T.

Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations (ICLR)*, 2018.

Christopher J. Shallue and Daniel J. Eisenstein. Reconstructing cosmological initial conditions from late-time structure with convolutional neural networks. *Monthly Notices of the Royal Astronomical Society*, 520:6256–6267, 2023. doi: 10.1093/mnras/stad528.

Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=vaRCHVj0uGI.

Hideki Tanimura, Nabila Aghanim, Victor Bonjean, and Saleem Zaroubi. Convolutional neural network-reconstructed velocity for kinetic SZ detection. *Astronomy & Astrophysics*, 662:A48, 2022. doi: 10.1051/0004-6361/202243046.

M. Trenti, M. Stiavelli, R. J. Bouwens, P. Oesch, J. M. Shull, G. D. Illingworth, L. D. Bradley, and C. M. Carollo. The galaxy luminosity function during the reionization epoch. *The Astrophysical Journal Letters*, 714(2):L202–L207, 2010. doi: 10.1088/2041-8205/714/2/L202.

Paolo Villani, Jörg Unger, and Martin Weiser. Adaptive gaussian process regression for bayesian inverse problems. *arXiv preprint arXiv:2404.19459*, 2024.

Jean Virieux and Stéphane Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6):WCC1–WCC26, 2009.

Lewis H. Weinberger, Martin G. Haehnelt, and Girish Kulkarni. Modelling the observed luminosity function and clustering evolution of lyα emitters: growing evidence for late reionization. *Monthly Notices of the Royal Astronomical Society*, 485(1):1350–1366, 2019. doi: 10.1093/mnras/stz481.

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

Brandon Zhao, Diana Scognamiglio, Olivier Doré, and Katherine L. Bouman. Generative diffusion priors for 3d mapping of the dark universe. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 23581–23590, June 2026.

A Appendix

A.1 Stability of $S(\phi, y^*)$ under surrogate approximation

Let $S(\phi, y^*)$ (equation 9) be computed using true $\mathcal{F}$ and $\hat{S}(\phi, y^*)$ be computed using the surrogate $\hat{\mathcal{F}}$. For continuously differentiable $\mathcal{F}$, we can state

Theorem A.1. *Assume $\|\mathcal{F} - \hat{\mathcal{F}}\|^2 \leq \epsilon$ and $\|\nabla\mathcal{F} - \nabla\hat{\mathcal{F}}\| \leq \epsilon_J$ over the regions where posterior mass is concentrated, and that $\nabla\mathcal{F}$ and $\nabla\mathcal{F}^{-1}$ are bounded. Then, there exist constants $C_1, C_2 > 0$ such that $|S(\phi, y^*) - \hat{S}(\phi, y^*)| \leq C_1\epsilon + C_2\epsilon_J$*

Proof. $|S(\phi, y^*) - \hat{S}(\phi, y^*)| = |r^\top A r + \hat{r}^\top \hat{A}\hat{r}|$, where $r = y - y^*, \hat{r} = \hat{\mathcal{F}} - y^*, A = (\nabla\mathcal{F}^{-1})^\top \nabla\mathcal{F}^{-1}$ and $\hat{A} = (\nabla\hat{\mathcal{F}}^{-1})^\top \nabla\hat{\mathcal{F}}^{-1}$.

By triangle inequality, $|S(\phi, y^*) - \hat{S}(\phi, y^*)| \leq \mathbb{E}_u[|r^\top A r - \hat{r}^\top A\hat{r}|] + \mathbb{E}_u[|\hat{r}^\top (A - \hat{A})\hat{r}|]$.

The first term on the right side of the equality is $\leq C_1\epsilon$, and the second term is $\leq C_2\epsilon_J$. Here, C_1, C_2 are constants that depend on the local behaviour of $\mathcal{F}$. □

Hence, small approximation errors due to using the surrogate $\hat{\mathcal{F}}$ induce small errors in the bias estimate $S(\phi, y^*)$.

A.2 Experiment Details

A.2.1 Linear Inverse Problem

Forward Operator Construction The discrete forward operator $A \in \mathbb{R}^{64\times 64}$ models an ill-posed linear mapping acting on a discretized continuous latent field $u \in \mathbb{R}^{64}$. To systematically control the degree of ill-posedness, A is constructed via Singular Value Decomposition (Hansen, 2010):

$$A = U\Sigma V^\top, \tag{14}$$

where $U, V \in \mathbb{R}^{64\times 64}$ are orthogonal matrices generated via QR factorization of independent standard Gaussian matrices $G_U, G_V \sim \mathcal{N}(0,1)^{64\times 64}$. The diagonal matrix $\Sigma = \mathrm{diag}(s_1, s_2, \ldots, s_{64})$ contains singular values that decay geometrically according to a targeted conditioning

$$\kappa(A) = s_1/s_{64}, \quad s_k = s_1 \cdot \left(\frac{1}{\kappa(A)}\right)^{\frac{k-1}{63}} \quad k \in \{1, \ldots, 64\}, \tag{15}$$

with $s_1 = 1.0$. We generate distinct forward operators across three condition number regimes $\kappa(A) \in \{10, 10^2, 10^3\}$.

Latent Field Generation Latent field signals $u \in \mathbb{R}^{64}$ are generated on a uniform spatial grid $x \in [0, 1]$ using a finite sum of continuous sinusoidal basis functions with randomized phase and modal amplitudes:

$$u(x) = \sum_{k\in\mathcal{F}} \alpha_k \sin\left(2\pi f_k x + \phi_k\right), \quad x_j = \frac{j-1}{63},\ j \in \{1, \ldots, 64\}, \tag{16}$$

where random phases are drawn as $\phi_k \sim \mathcal{U}(0, 2\pi)$ and modal coefficients as $\alpha_k \sim \mathcal{U}(0.5, 1.5)$.

To evaluate active learning acquisition performance, the spatial frequency sets f used to generate latent fields differ between dataset splits. Signals drawn for training comprises $N_{\text{train}} = 50$ initial training examples and $N_{\text{pool}} = 2000$ candidate pool points constructed using low spatial frequencies $f_{\text{train}} = \{1, 2, 3, 4, 5\}$. In contrast, evaluation signals ($N_{\text{test}} = 2000$) are generated using higher spatial frequencies $f_{\text{test}} = \{6, 7, 8, 9, 10, 11, 12\}$. Prior to applying the forward operator, all generated signals across splits are centered to zero mean and normalized to unit variance.

Noise Pipeline Observation vectors $y_i \in \mathbb{R}^{64}$ are generated as $y_i = Au_i + \boldsymbol{\epsilon}_i$, with $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \Sigma_{\boldsymbol{\epsilon},i})$ and diagonal covariance $\Sigma_{\boldsymbol{\epsilon},i} = \mathrm{diag}(\sigma_{i,1}^2, \ldots, \sigma_{i,64}^2)$ combining a fixed channel profile with a per-sample noise level:

$$\sigma_{i,j}^2 = \sigma_i^2 \cdot c_j, \qquad c_j = \frac{\exp\left(\gamma \cdot \frac{j-1}{63}\right)}{\frac{1}{64}\sum_{k=1}^{64} \exp\left(\gamma \cdot \frac{k-1}{63}\right)}, \qquad \gamma = 4.0, \tag{17}$$

where c_j is fixed and shared across all samples and splits. The per-sample level σ_i^2 is set so that the realized SNR depends deterministically on the mean modal amplitude $\bar{\alpha}_i = \frac{1}{|\mathcal{F}|}\sum_{k\in\mathcal{F}} \alpha_k$ of u_i itself, rather than on an independently drawn quantity:

$$\mathrm{SNR}_i^{\mathrm{dB}} = \mathrm{SNR}_{\max}^{\mathrm{dB}} - \rho_i \left(\mathrm{SNR}_{\max}^{\mathrm{dB}} - \mathrm{SNR}_{\min}^{\mathrm{dB}}\right), \quad \rho_i = \frac{\bar{\alpha}_i - 0.5}{1.5 - 0.5}, \quad \sigma_i^2 = \frac{P_i^{\mathrm{sig}}}{10^{\mathrm{SNR}_i^{\mathrm{dB}}/10}}, \tag{18}$$

with $P_i^{\mathrm{sig}} = \frac{1}{64}\|Au_i\|^2$, $\mathrm{SNR}_{\min}^{\mathrm{dB}} = 0$, $\mathrm{SNR}_{\max}^{\mathrm{dB}} = 30$. Since $\bar{\alpha}_i$ depends only on u_i's own amplitude draw.

Training The posterior estimation neural network $q_\phi(y) = \left(\hat{\mu}_\phi(y), \log \hat{\sigma}_\phi^2(y)\right) \in \mathbb{R}^{64} \times \mathbb{R}^{64}$ is implemented as a 4-layer fully connected MLP with 128 hidden units per layer and GELU activations. Models are trained using the Adam optimizer with a learning rate $\eta = 4e-5$, minimizing the joint supervision loss $\mathcal{L}_\beta(\phi) + \gamma\mathcal{L}_{\mathrm{f}}(\phi;\psi)$, with $\gamma > 0$ with $\gamma = 0.08$ over 200 epochs per active learning step. All experiments are executed on a single NVIDIA GeForce RTX 3080 GPU. The acquisition curves are presented in Figure 5. Note that because the exact Gaussian Process baseline exhibited significantly larger relative errors compared to the neural inverse models, its performance trajectories are displayed on separate plots (Figure 6) to preserve visual clarity across the primary model comparisons.

A.2.2 Inverse Laplace Transform Implementation Details

We consider the severely ill-conditioned Laplace inversion problem discretized with $N_s = N_t = 40$ log-spaced measurement points, yielding a forward operator condition number of $\kappa(A) \approx 5.6\times10^{16}$. The dataset split consists of an initial labelled training set $N_{\mathrm{train}} = 40$, a candidate active pool $N_{\mathrm{pool}} = 2000$, and an evaluation test set $N_{\mathrm{test}} = 2000$. To evaluate robustness against distribution shift, the training/pool sets and the evaluation test set are drawn from non-overlapping parameter ranges within the same continuous piecewise-triangular pulse family.

Each spatial latent field realization $u_i \in \mathbb{R}^{40}$ is generated as

$$u(t) = \begin{cases} h\left(\frac{t}{t_p}\right), & 0 \le t < t_p, \\ h\left(\frac{t_e - t}{t_e - t_p}\right), & t_p \le t < t_e, \\ 0, & t \ge t_e, \end{cases} \qquad 0 \le u(t) \le 1, \tag{19}$$

where the peak location t_p, peak height h, and terminal boundary t_e are sampled independently for each example i. For the training set N_{train} the parameters are drawn from

$$t_{p,i} \sim \mathcal{U}(0.5, 1.8), \tag{20}$$
$$h_i \sim \mathcal{U}(0.2, 0.7), \tag{21}$$
$$t_{e,i} = t_{p,i} + \Delta t_i, \quad \text{where } \Delta t_i \sim \mathcal{U}(0.5, 1.5), \tag{22}$$

whereas the evaluation test set N_{test} features a shifted distribution with parameters sampled from non-overlapping ranges

$$t_{p,i} \sim \mathcal{U}(1.8, 3.1), \tag{23}$$
$$h_i \sim \mathcal{U}(0.6, 1.0), \tag{24}$$
$$t_{e,i} = t_{p,i} + \Delta t_i, \quad \text{where } \Delta t_i \sim \mathcal{U}(1.5, 2.5). \tag{25}$$

Measurements are corrupted by additive isotropic Gaussian noise with fixed standard deviation $\sigma_n = 0.01$.

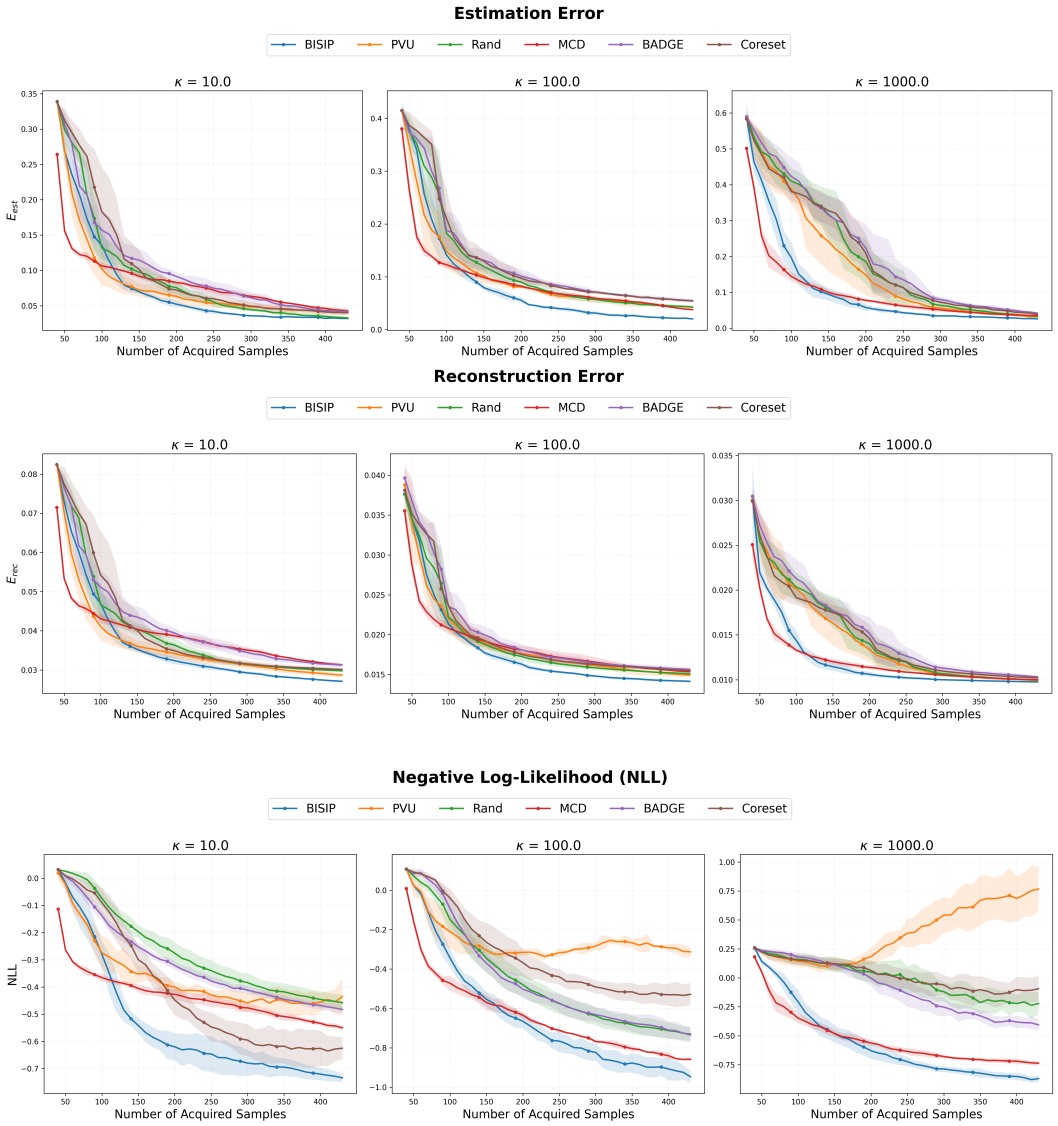


Figure 5: Active learning performance trajectories for neural inverse models across acquired samples for condition numbers $\kappa \in \{10, 10^2, 10^3\}$ on Linear Inverse Problem benchmark. Panels show (a) Estimation Error (E_{est}), (b) Forward Reconstruction Error (E_{rec}), and (c) Negative Log-Likelihood (NLL).

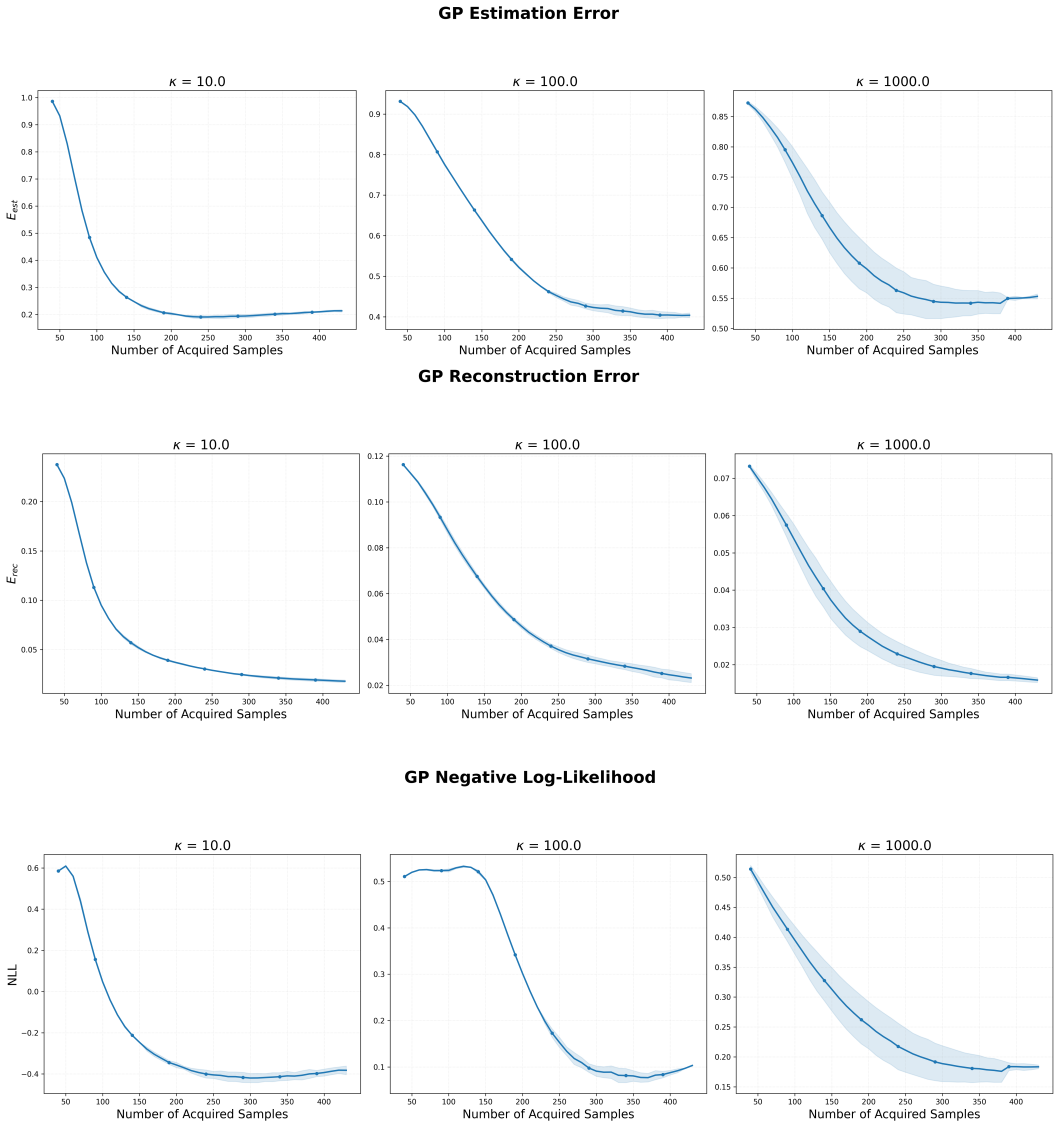
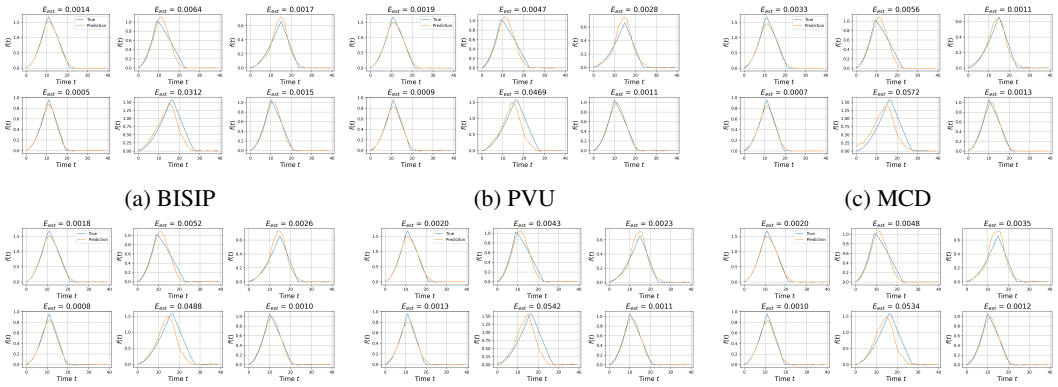


Figure 6: GP active selection trajectories across acquired samples for condition numbers $\kappa \in \{10, 10^2, 10^3\}$. Panels show (a) Estimation Error (E_{est}), (b) Forward Reconstruction Error (E_{rec}), and (c) Negative Log-Likelihood (NLL) respectively.

Training The inverse network q_ϕ maps observation vectors $y \in \mathbb{R}^{40}$ to predictive parameter targets $(\boldsymbol{\mu}_\phi, \log \boldsymbol{\sigma}^2_\phi) \in \mathbb{R}^{40} \times \mathbb{R}^{40}$. The architecture consists of a linear input projection, residual blocks with a hidden width of 256, and separate linear output heads for posterior mean and variance predictions.

Models are trained using the β-NLL objective 4, regularised with a total-variation (TV) penalty applied directly to the predicted variance field. Starting from the first acquisition round, a forward-consistency penalty is introduced into the objective function with a geometrically decaying weight profile $w_{\text{fc}} = 0.05 \times 0.9^{\text{round}}$. Optimisation is performed using AdamW (weight_decay $= 10^{-4}$) with an initial learning rate of 10^{-4} dynamically annealed to approximately 3×10^{-5} via `ReduceLROnPlateau`. Active learning acquisition follows the method in Algorithm 1.

To provide further evaluation, the predicted latent reconstructions $\hat{\boldsymbol{u}}(t)$ across active acquisition strategies against the true latent pulse realizations $u(t)$ is shown in Fig 7.



(a) BISIP (b) PVU (c) MCD

(d) Rand (e) BADGE (f) Coreset

Figure 7: Qualitative reconstruction comparison on the Inverse Laplace Transform benchmark with the same set of test samples across different acquisition strategies. BISIP achieves tighter variance estimates and preserves peak features and edge boundaries under severe operator ill-conditioning ($\kappa(A) \approx 5.6 \times 10^{16}$)

A.2.3 Reionization Field Reconstruction

The posterior estimation network q_ϕ is a conditional volumetric encoder-decoder built on a 3D U-Net backbone processing $\mathbb{R}^{64^3}$ spatial inputs y. The encoder consists of four resolution levels with feature channels doubling at each downsampling step, utilizing $3 \times 3 \times 3$ convolutional blocks with Instance Normalization and LeakyReLU activations ($\alpha = 0.2$). At the bottleneck, feature maps are projected via the reparameterization trick to parameterize the mean $\mu_z(y)$ and log-variance $\log \sigma^2_z(y)$ of a spatially-resolved latent distribution $\boldsymbol{z} \sim \mathcal{N}(\mu_z(\boldsymbol{y}), \sigma^2_z(\boldsymbol{y})) \in \mathbb{R}^{18\times8\times8\times8}$. The decoding pathway splits into two parallel branches with skip connections to the encoder: a mean branch $\mu_\phi(\boldsymbol{y})$ predicting expected reionization values, and a variance branch predicting the diagonal log-variance $\log \sigma^2_\phi(\boldsymbol{y})$. For the MCD baseline, a location-specific dropout schedule is implemented in place of a single global rate. Monte Carlo posterior sampling utilizes $M = 15$ draws during training and same draws during acquisition.

$\hat{\mathcal{F}}_\psi$ leverages a modified SegResNet architecture (Myronenko, 2018) implemented via the MONAI framework. The architecture is built with repeating residual blocks using 3^3 convolutions and ReLU non-linearities across four encoding levels, with spatial reduction performed via strided convolutions and upsampling via trilinear interpolation. The standard medical segmentation head is replaced with a linear 1^3 regression layer mapping internal representations directly to the predicted tracer volume under the surrogate loss in Eq. (12). Rather than being retrained from scratch each round, $\hat{\mathcal{F}}_\psi$ is warm-started from a pretrained checkpoint and fine-tuned at a learning rate of $5e-5$ using the AdamW optimizer after every acquisition step.

Algorithm 1 BISIP: Bias-guided Stochastic Search for Inverse Problems

Input: Initial labelled set $D_0 = \{(u_i, y_i)\}$, unlabelled pool $Y^* = \{y_j^*\}$, number of rounds T, batch size k, hyperparameters β, γ
Output: q_ϕ, D_T
Initialize q_ϕ and $\hat{\mathcal{F}}_\psi$ using D_0
Train q_ϕ by minimizing $\mathcal{L}_\beta(\phi)$ (Eq. 4)
Train $\hat{\mathcal{F}}_\psi$ on D_0 to approximate the forward map by minimizing

$$\mathbb{E}_{(u,y)\sim D_0}\left[\left\|y - \hat{\mathcal{F}}_\psi(u)\right\|^2\right]$$

for $t = 1$ to T **do**
  for each $y_j^* \in Y^*$ **do**
    Draw M Monte Carlo samples $u^{(1)}, \ldots, u^{(M)} \sim q_\phi(u \mid y_j^*)$
    Compute predicted observations $y^{(m)} = \hat{\mathcal{F}}_\psi(u^{(m)})$ for $m = 1, \ldots, M$
    Estimate bias score $S(\phi, y_j^*)$ from $\{u^{(k)}, y^{(k)}\}_{k=1}^K$ using Eq. 9 or Eq. 10
  end for
  Select top-k pool candidates $\mathcal{B}_t$ according to the acquisition score $\alpha(y;\phi) := S(\phi, y)$ (Eq. 11)
  for each $y_j^* \in \mathcal{B}_t$ **do**
    Sample

$$u_j \sim \mathcal{N}\left(\mu_\phi(y_j^*), \mathrm{diag}(s_d + \sigma_d^2)\right)$$

    Query the simulator to obtain $y_j = \mathcal{F}(u_j)$
  end for
  $D_{t+1} \leftarrow D_t \cup \{(u_j, y_j)\}_{j\in\mathcal{B}_t}$
  Update $\hat{\mathcal{F}}_\psi$ on D_{t+1}
  Update ϕ by minimizing $\mathcal{L}_\beta(\phi) + \gamma\mathcal{L}_\mathrm{f}(\phi;\psi)$ (Eq. 12)
end for
return q_ϕ

Training proceeds in stages rather than through a single joint objective. Firstly, only the mean branch is supervised, with gradients to the variance head masked followed by joint optimisation of both.

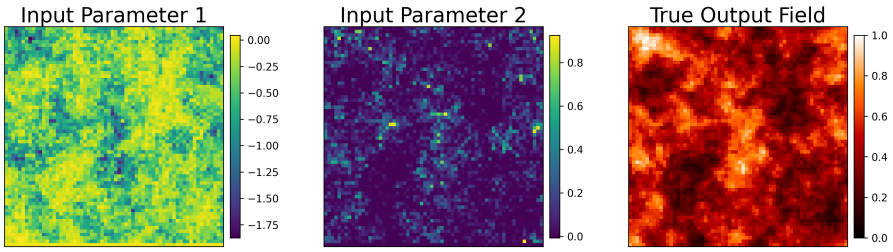
$$\mathcal{L} = \mathcal{L}_\beta(u, \mu_\phi, \log\sigma_\phi^2; \beta{=}0.6) + kl_\beta \cdot \mathrm{KLD}(z) + \mathcal{L}_\mathrm{f} \tag{26}$$

where KLD is kl-divergence term in latent space used as a regulariser. $\mathcal{L}_\mathrm{f}$ dropped entirely in initial training, since no acquired samples yet exist to make it reliable. All stages are optimized with AdamW each with its own optimizer instance and learning rate; 3e-5 for on a cold-started model or 1e-5 on a warm-started model, each following a cosine annealing schedule.

Implemenation of baselines is discussed section A.3. For Coreset and BADGE, we did use approximations considering the high dimensional nature of the problem. Coreset performs greedy k-center selection, but in a pooled 96-dimensional feature-embedding space rather than observation space. Finally, BADGE computes gradient embeddings of a self-labelled NLL loss with respect to the final mean-decoder convolutional layer, giving a 432-dimensional embedding per candidate, and selects a batch via k-means++ sampling seeding. All experiments use a single NVIDIA L4 GPU.

A.2.4 Dataset Generation for Reionization Benchmark

We generate synthetic training data using a three-stage pipeline consisting of 21cmFAST (Mesinger et al., 2011; Murray et al., 2020), SIMPLE-GEN (Maitra et al., 2025), and a final preprocessing stage. Each example pairs a sparse three-dimensional galaxy distribution with the corresponding hydrogen reionization field. The latter describes the spatial state of intergalactic hydrogen during the Epoch of Reionization (EoR), when radiation from the first galaxies progressively ionized the surrounding gas. Our procedure follows the forward-modelling strategy of TORRCH (Maitra et al., 2026a), but uses the computationally efficient semi-numerical 21cmFAST framework to generate a



(a) Input parameters and output field.

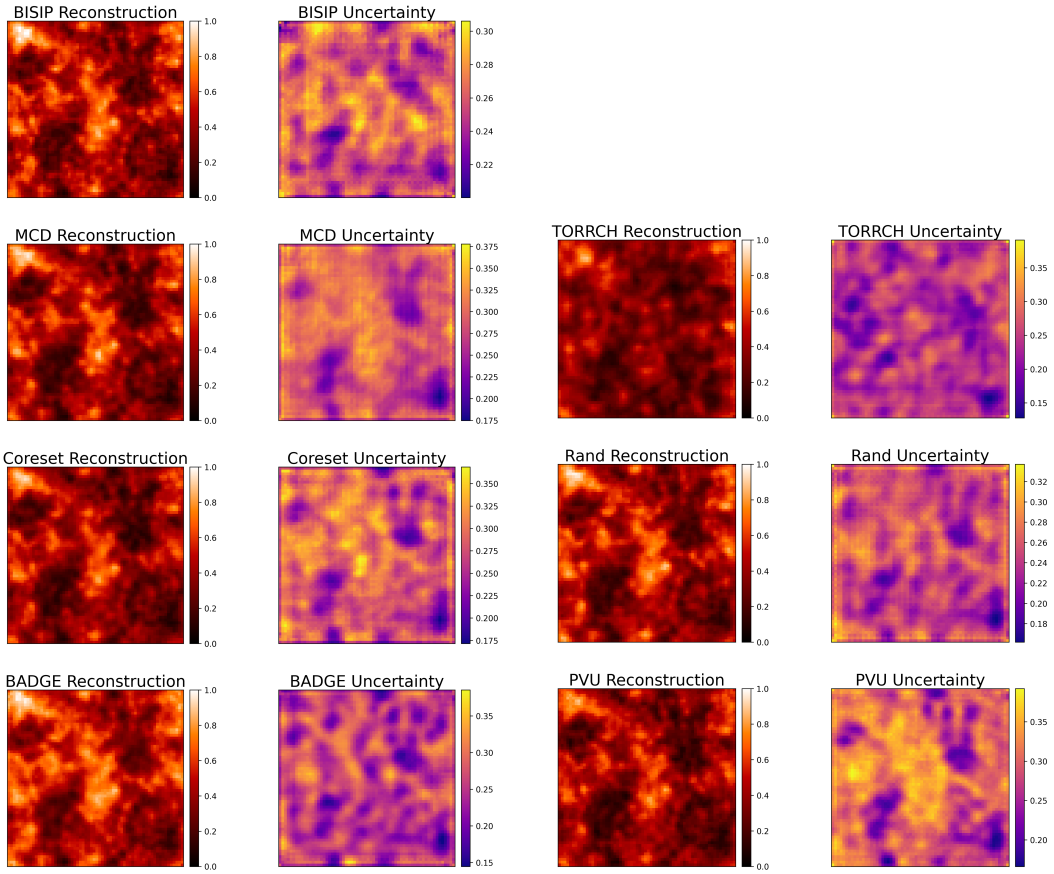


Figure 8: Comprehensive evaluation for test sample, showing (a) the conditioning input fields and ground truth, alongside method wise mean predictions and uncertainty estimates

larger ensemble of independent realizations. The forward model is parameterized by astrophysical inputs rather than arbitrary perturbations of the continuous ionization field. For each sampled parameter set, 21CMFAST jointly generates a self-consistent ionization field and halo catalogue, which are then passed to SIMPLE-GEN to construct the corresponding galaxy observables. Thus, the dataset contains only physically consistent field–catalogue pairs produced by the forward pipeline.

Simulation of ionization fields and dark-matter haloes. We use 21CMFAST to generate three-dimensional realizations of the matter and ionization structure at $z \simeq 7.14$. Each simulation covers a periodic volume of side length $80\,h^{-1}\mathrm{cMpc}$ and produces the hydrogen ionization field together with a catalogue of dark-matter haloes, the gravitationally collapsed structures in which galaxies form. We vary the astrophysical parameters M_{turn} and $f_{\mathrm{esc},10}$, which control the abundance and ionizing efficiency of galaxy-hosting haloes. Other cosmological and simulation parameters are kept

fixed. For each realization, we retain the three-dimensional ionization field and the positions and masses of haloes above $10^{9.5}\,M_\odot/h$. The ionization field serves as the target of the reconstruction problem.

Generation of mock galaxy observations. The halo catalogues are converted into synthetic galaxy catalogues using SIMPLE-GEN, following the forward-modelling framework of Weinberger et al. (2019). This step connects the simulated halo population to quantities accessible to galaxy surveys. Each halo is first assigned a rest-frame ultraviolet magnitude, $M_{\rm UV}$, through abundance matching to the observed high-redshift UV luminosity function (Bouwens et al., 2015), with a duty-cycle correction accounting for the fraction of haloes hosting observable galaxies at a given time (Trenti et al., 2010). A subset of these galaxies is then identified as Lyman-α emitters (LAEs), i.e. galaxies detected through their Lyα emission. Intrinsic Lyα luminosities are assigned by sampling a UV-luminosity-dependent distribution of Lyα equivalent widths (Dijkstra & Wyithe, 2012). Because Lyα photons interact strongly with neutral hydrogen, their observed luminosity depends on the surrounding ionization state. SIMPLE-GEN therefore computes the transmission of Lyα emission through the simulated gas distribution, defining

$$L_{\rm Ly\alpha}^{\rm obs} = T_{\rm IGM} L_{\rm Ly\alpha}^{\rm int}, \tag{27}$$

where $T_{\rm IGM}$ is the effective transmission factor. This forward model directly couples the observable LAE population to the underlying ionization morphology.

Catalogue preprocessing. We retain two complementary tracer populations. LAEs are selected using

$$L_{\rm Ly\alpha}^{\rm obs} > 10^{42}\,{\rm erg\,s^{-1}}, \tag{28}$$

while non-Lyα-selected galaxies (NLSGs), identified through their ultraviolet continuum rather than Lyα emission, are selected using

$$M_{\rm UV} < -17. \tag{29}$$

Objects outside the reconstruction volume are discarded. The ionization fields are represented on a common 64^3 Cartesian grid, while the corresponding LAE and NLSG catalogues retain their three-dimensional coordinates and source properties. Each realization therefore defines a paired example

$$\mathcal{D}_i = (\mathcal{G}_i, \mathbf{X}_i)\,, \tag{30}$$

where $\mathcal{G}_i$ denotes the sparse, variable-length galaxy catalogue and $\mathbf{X}_i \in \mathbb{R}^{64\times64\times64}$ denotes the corresponding ionization field.

A.2.5 Evaluation Under Distribution Shift and Additional Results

To assess generalization beyond the training distribution, we consider two independent evaluation settings that has distinct forms of distribution shift. The training and validation data are generated using a fixed simulation seed of 1441, with $(M_{\rm turn}, f_{\rm esc,10})$ randomly sampled over their respective parameter ranges. The training set comprises 150 realizations with $M_{\rm turn} \in [8.1, 9.2]$ and $f_{\rm esc,10} \in [-2.15, -1.5]$, while the validation set contains 200 realizations spanning $M_{\rm turn} \in [8.0, 9.25]$ and $f_{\rm esc,10} \in [-2.20, -1.55]$.

First, we evaluate generalization to an unseen region of the parameter space. We construct an evaluation set of 400 samples using parameter values that are non-overlapping with those used to construct $p_{\rm train}$, with $M_{\rm turn} \in [9.3, 10.2]$ and $f_{\rm esc,10} \in [-1.6, -0.4]$. This setting evaluates whether the learned model can generalize to previously unseen combinations of the underlying astrophysical parameters. The corresponding results are reported in the main-text Table 3.

Second, we evaluate generalization across an independent realization of the underlying density field while keeping the parameter distribution within the same broad range as the training distribution. Specifically, we generate a separate set of 400 samples using the training-set $(M_{\rm turn}, f_{\rm esc,10})$ parameter ranges, but with a different simulation seed, 1445. The simulation seed is a hyperparameter of the cosmological data-generation procedure that determines the realization of the underlying density field. Thus, this setting isolates generalization across different realizations of the simulated universe rather than extrapolation to unseen parameter values. Results for this independent-realization setting are reported in Table 4.

Table 4: Standard Evaluation Metrics along with Calibration Metrics across different acquisition strategies on Ionization Field Reconstruction Benchmark

Method	E_{est}	NLL	MACE	RMSCE	Miscalibration
BISIP	0.0518 ± 0.0061	-0.9679 ± 1.0063	0.0479 ± 0.0223	0.0583 ± 0.0259	0.0451 ± 0.0214
MCD	0.0667 ± 0.0126	-0.4619 ± 0.7169	0.0968 ± 0.0442	0.1181 ± 0.0585	0.0928 ± 0.0432
Rand	0.0943 ± 0.0225	-0.0119 ± 0.9613	0.1357 ± 0.0583	0.1476 ± 0.0613	0.1297 ± 0.0562
Coreset	0.0640 ± 0.0075	1.4363 ± 1.2527	0.0540 ± 0.0172	0.0628 ± 0.0231	0.0508 ± 0.0163
BADGE	0.1369 ± 0.0368	-0.3018 ± 0.1259	0.0837 ± 0.0357	0.0966 ± 0.0484	0.0801 ± 0.0344
PVU	0.1085 ± 0.0788	0.3432 ± 0.6731	0.0912 ± 0.0528	0.1022 ± 0.0589	0.0865 ± 0.0503
TORRCH	0.0790 ± 0.0098	0.7115 ± 0.5823	0.0513 ± 0.0203	0.0577 ± 0.0204	0.0482 ± 0.0193

We additionally observe that the proposed approach remains well calibrated across the other active-acquisition methods considered in our experiments shown in Table 4. For calibration assessment, we report the Mean Absolute Calibration Error (MACE) and Root Mean Squared Calibration Error (RMSCE), standard metrics as defined in Chung et al. (2021). However, we restrict this observation to the empirical settings evaluated here and do not make broader claims regarding calibration across acquisition strategies. To further examine the spatial quality of the predicted uncertainty estimates, we consider samples from the non-overlapping test set generated with the seed 1441. While the corresponding reconstructions are shown in Figure 4, here we additionally visualize the uncertainty maps predicted by each model alongside its reconstruction in Fig. 8. The predicted uncertainty bounds from PVU and BISIP exhibit a sensible spatial correspondence between the magnitude of σ and regions of increased structural complexity, with uncertainty changing more sharply near relevant spatial boundaries. In comparison, the uncertainty estimates from TORRCH and Rand exhibit more spatially diffuse distributions, with less correspondence to the boundaries and fine-scale structures of the reconstructed field.

A.3 Discussion on Baselines Implementation

A.3.1 PINN

Our benchmarks lacks a closed-form differential equation, hence the PINN baseline is formulated as a physics-regularised regression model. The training objective combines a standard supervised data-mismatch term in (u-space) and a physics inconsistency term in (y-space) target reference field $y_{\text{true}} : \mathcal{L}_\beta(\phi) + \gamma \mathcal{L}_{\text{f}}(\phi; \psi)$ where we set $\lambda = 0.8$ in all benchmark tasks and L_β and L_f are defined in Equations 4 and 12 respectively.

A.3.2 Active Acquisition Strategies

All baselines operate on the same heteroscedastic output head $(\mu_\phi(y), \sigma_\phi(y))$ used by our method, differing in (i) how the top-k pool candidates are scored and (ii) how the acquired candidate u is sampled once selected.

PVU. Using heteroscedastic output head of q_ϕ, the predictive covariance for candidate y^* is $\text{Cov}_{\text{pred}}(y^*) = \text{diag}\big(\hat{\sigma}^2(y^*)\big)$, where $\hat{\sigma}^2(y^*)$ is the predicted per-voxel aleatoric variance from a single forward pass. Kendall & Gal (2017). Candidates are scored by the mean predictive variance over the field and Eq 11 is defined as,

$$\alpha_{\text{PVU}}(y^*) = \frac{1}{D} \sum_{d=1}^{D} \hat{\sigma}_d^2(y^*), \tag{31}$$

where $D = |V|$ is the number of voxels. Unlike our $S(\phi, y^*)$, this criterion is computed entirely from the model's own predictive head and does not query the forward model.

Rand Candidates are assigned i.i.d. uniform random scores, and this is uniform sampling without replacement from the pool, independent of q_ϕ.

MCD. With dropout kept active at inference, we draw $T = 20$ stochastic forward passes to obtain the predictive mean $\bar{\mu}_\phi(y^*)$, the mean aleatoric log-variance $\overline{\log \sigma_\phi^2}(y^*)$, and the epistemic variance

across passes, $\sigma^2_{\text{epi}}(y^*) = \widehat{\text{Var}}_{t=1}^{T}\left[\mu_\phi^{(t)}(y^*)\right]$. Candidates are scored by the mean epistemic variance over the volume,

$$\alpha_{\text{MCD}}(y^*;\phi) = \frac{1}{|V|}\sum_v \sigma^2_{\text{epi}}(y^*)_v, \tag{32}$$

following Gal & Ghahramani (2016). For the top-k selected candidates, the acquired sample combines both uncertainty sources additively with independent noise draws.

Coreset. We extract a pooled feature embedding $z(y)$ for each candidate from an intermediate layer of q_ϕ, spatially pooled via adaptive average pooling and ℓ_2-normalised. Given the embeddings of the currently labelled set, $Z_t = \{z(y') : y' \in D_t\}$, each pool candidate y^* is scored by

$$\alpha_{\text{Coreset}}(y^*;\phi) = \min_{z' \in Z_t} \|z(y^*) - z'\|_2, \tag{33}$$

that is, candidates that are least well-represented by the current labelled set receive the highest score following the greedy k-center formulation of Sener & Savarese (2018).

BADGE. BADGE (Ash et al., 2020) selects a diverse, high-uncertainty batch via gradient embeddings and k-means++ seeding, originally formulated for classification. We adapt it to our regression setting by sampling a pseudo-target $\tilde{u} \sim \mathcal{N}(\mu_\phi(y), \hat{\sigma}^2_\phi(y^*))$ from predictive distribution OF q_ϕ and computing the gradient of the NLL loss with respect to the last-layer weights as the embedding $g(y^*)$. As in Ash et al. (2020), we then apply k-means++ seeding over $\{g(y^*)\}_{y^* \in \mathcal{Y}^*}$ to select a batch of size k that is jointly high-magnitude and diverse and deterministically initialise with the candidate of maximum embedding norm.

GP As an additional non-neural baseline for Linear Inverse and Laplace Transform benchmark, we fit independent Gaussian Processes over each output dimension using GPyTorch Gardner et al. (2018), and score pool candidates by the summed posterior predictive variance,

$$\alpha_{\text{GP}}(y^*) = \sum_d \sigma^2_{\text{GP},d}(y^*). \tag{34}$$

A.3.3 FLAIR

To evaluate FLAIR Erbach et al. (2025) for our field reconstruction benchmark, we independently implemented the method following the formulation described in the original work, as the authors' implementation was not available to us. FLAIR uses a pretrained flow-based generative model as a well-defined prior and performs training-free variational inference by alternating flow-matching-based prior updates with hard data-consistency steps. We adapted these components to our 3D cosmological reconstruction setting, using our pretrained generative model to represent the prior and our cosmological forward model to impose consistency with the observed measurements.To the best of our ability, we have attempted to faithfully reproduce FLAIR for field reconstruction based on the details reported in the original work; this constitutes an independent implementation rather than an evaluation using the authors' original codebase, and minor discrepancies from the original results may arise from implementation-level choices not fully specified in the source material.