GPARA: GRAPH-POSTERIOR-ALIGNED REFINEMENT AND ACTIVE ACQUISITION FOR GROUNDING DIFFUSION PRIORS

Anonymous authors
Paper under double-blind review

ABSTRACT

Active grounding of a frozen diffusion prior requires jointly determining where new measurements should be taken and how they should be used to refine the current reconstruction. Posterior-ensemble-based methods can estimate acquisition utility from generated samples, but require repeated ensemble generation as observations accumulate and capture posterior geometry only through empirical statistics. This paper proposes GPARA, which learns a context-dependent graph surrogate over diffusion prediction residuals, inducing an explicitly reusable posterior response operator that propagates measurement innovations to unobserved variables and evaluates candidate measurements through weighted posterior-risk reduction. Under the matched surrogate, we show that the same response operator also determines expected one-step acquisition benefit and yields an analytic ranking consistent with expected reconstruction improvement. A bounded learned residual calibrates the analytic utility to account for surrogate mismatch, while a small prior ensemble is generated once and reconditioned to update risk weights without repeated diffusion posterior sampling during acquisition. Experiments on two reconstruction tasks spanning wireless sensing and computer vision show consistent improvements in refinement and active acquisition over the evaluated baselines. Ablations further support the complementary roles of step-wise graph refinement, adaptive risk weighting, and analytically anchored calibration.

1 INTRODUCTION

Diffusion models provide expressive priors for recovering high-dimensional signals, images, and physical fields from heterogeneous information (Graikos et al., 2022; Song et al., 2021; Rombach et al., 2022; Fei et al., 2023). Given auxiliary context $\boldsymbol{c}$, such as image, feature representations, geometry layout, or physical metadata, the learned prior $p_\theta(\boldsymbol{x} \mid \boldsymbol{c})$ can generate a plausible target $\boldsymbol{x}$. This approach has been applied to imaging, remote sensing, field reconstruction, and geometric perception (Wu et al., 2024a; Khanna et al., 2024; Li et al., 2025; Ke et al., 2024). However, conditional plausibility does not ensure agreement with the real target realization. Information absent from the context can lead to instance-specific discrepancies in the predicted structure and values.

Direct target measurements provide complementary evidence about the current realization. Existing inverse solvers incorporate observations through spectral or null-space decomposition (Kawar et al., 2022; Song et al., 2023; Wang et al., 2023), likelihood guidance (Chung et al., 2023; Aali et al., 2025), variational or moment-based posterior approximations (Mardani et al., 2024; Boys et al., 2024; Rout et al., 2024), data-consistency (Zhu et al., 2023; Zirvi et al., 2025; Song et al., 2024), filtering (Dou & Song, 2024), or posterior sampling (Zhang et al., 2025; Wu et al., 2024b). These methods primarily address reconstruction under a prescribed observation operator.

Since measurement acquisition incurs sensing energy, bandwidth, or other resources, the measurements are often sparse. The problem is thus not only how to reconstruct from a prescribed observation set, but also which observations are the most useful to acquire. We refer to this setting as *active grounding*: actively acquiring a limited number of measurements to refine a frozen diffusion prior. This requires evaluating how each candidate measurement would affect the unresolved target, a quantity that conventional inverse solvers do not generally expose as an explicit, reusable response

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.

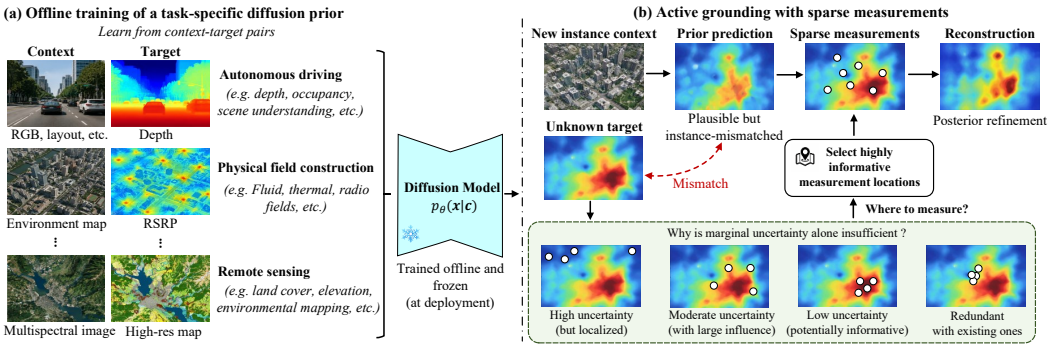


Figure 1: Active grounding of a frozen diffusion prior. A task-specific conditional prior is pretrained offline and frozen. At deployment, context alone may produce an instance-mismatched prediction, and sparse target measurements are thus acquired to ground the prior to the current realization.

for active acquisition. Fig. 1 illustrates this setting and highlights a key acquisition challenge, where a candidate's marginal uncertainty alone does not reflect its value for reconstructing the full target.

Active acquisition has been studied through Bayesian experimental design (Foster et al., 2019; 2021; Polyzos et al., 2024; Ivanova et al., 2021). Reconstruction-oriented active sensing uses uncertainty- or information-driven acquisition (Zhang et al., 2019; Morssy et al., 2024), and recent diffusion-based methods derive such criteria from predictive distributions through variance, diversity, or entropy (Barba et al., 2025; Gu et al., 2025; Nolan et al., 2025). Purely marginal criteria prioritize individually uncertain candidates but do not capture how cross-variable dependencies affect full-target reconstruction. More structured criteria account for posterior dependence (Iollo et al., 2025; Krause et al., 2008). Notably, AdaSense (Elata et al., 2024) estimates posterior covariance from empirical samples for linear minimum mean-squared error (LMMSE)-based criterion, coupling covariance estimation and acquisition cost to repeated ensemble regeneration as observations accumulate.

We address these limitations with GPARA, a graph-posterior-aligned refinement and acquisition framework. Building on graph precision models (Shuman et al., 2013; Rue & Held, 2005; Pinti et al., 2025), GPARA learns a graph surrogate over prediction residuals, rather than directly smoothing target values. Unlike AdaSense that infers posterior covariance from empirical ensembles, GPARA explicitly models residual dependence through a learned sparse precision, yielding a shared structured posterior response that both propagates local measurement innovations and evaluates prospective measurements, with risk weights specifying where uncertainty reduction is currently valuable. The analytic utility is retained as an anchor and augmented by a bounded residual calibration.

Our contributions are threefold: (i) we propose a structured graph residual refinement model that explicitly propagates local measurement innovations across related variables while keeping the diffusion backbone frozen; (ii) we develop a response-aligned acquisition that couples the same graph-induced posterior response with adaptive weighted Bayes risk, and show that its analytic utility is consistent with expected one-step reconstruction improvement; and (iii) We combine the analytic anchor with bounded residual calibration and matrix-free posterior computation using cached prior samples, stochastic probing, and sparse linear solves, avoiding dense covariance construction and repeated posterior-ensemble generation at every acquisition round.

2 Background and Problem Formulation

2.1 Conditional Diffusion Prior

We instantiate the formulation with a denoising diffusion probabilistic model (DDPM) (Ho et al., 2020; Nichol & Dhariwal, 2021), while the framework is not restricted to this particular formulation.

Let $\boldsymbol{x}_0 \in \mathbb{R}^M$ denote the clean target and $\boldsymbol{c}$ the available context. A DDPM defines a T-step forward diffusion process that progressively corrupts $\boldsymbol{x}_0$ toward standard Gaussian noise:

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = \mathcal{N}\left(\boldsymbol{x}_t; \sqrt{\bar{\alpha}_t}\boldsymbol{x}_0, (1-\bar{\alpha}_t)\boldsymbol{I}\right), \tag{1}$$

where $\{\beta_t\}_{t=1}^T$ is the noise schedule, $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and $\boldsymbol{I}$ is the identity matrix. A conditional denoising network $\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t, \boldsymbol{c})$ is trained to predict the noise injected into $\boldsymbol{x}_t$.

At inference time, the reverse process starts from Gaussian noise $\boldsymbol{x}_T \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$. At each reverse step t, the conditional denoiser can first estimate the clean target as

$$\widehat{\boldsymbol{x}}_{0|t} = \frac{\boldsymbol{x}_t - \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_\theta(\boldsymbol{x}_t, t, \boldsymbol{c})}{\sqrt{\bar{\alpha}_t}}. \tag{2}$$

The standard reverse transition then uses $\widehat{\boldsymbol{x}}_{0|t}$ and $\boldsymbol{x}_t$ to predict $\boldsymbol{x}_{t-1}$:

$$\boldsymbol{x}_{t-1} = a_t \widehat{\boldsymbol{x}}_{0|t} + b_t \boldsymbol{x}_t + \sqrt{\tilde{\beta}_t}\boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}), \tag{3}$$

where a_t, b_t, and $\tilde{\beta}_t$ are determined by the predefined diffusion schedule. Once trained, the diffusion model is frozen and provides the conditional generative prior $p_\theta(\boldsymbol{x}_0 \mid \boldsymbol{c})$ throughout deployment.

2.2 Problem Formulation

Although the context $\boldsymbol{c}$ provides rich information cues about $\boldsymbol{x}_0$, environment dynamics, sensing imperfections, transient states, or fine-scale structures may induce a non-negligible mismatch between the context-conditioned prior and the actual target. We retain the pretrained diffusion model $p_\theta(\boldsymbol{x}_0 \mid \boldsymbol{c})$ as an offline prior and use direct measurements to ground it to the current target.

Since it is costly and often infeasible to obtain dense target measurements, only a limited budget of $\Gamma \ll M$ measurements can be acquired. Let $\mathcal{S}$ be the selected observation set with $|\mathcal{S}| \leq \Gamma$, and let $\boldsymbol{H}_\mathcal{S} \in \mathbb{R}^{|\mathcal{S}| \times M}$ be the corresponding observation matrix. The measurement model is given by

$$\boldsymbol{y}_\mathcal{S} = \boldsymbol{H}_\mathcal{S}\boldsymbol{x}_0 + \boldsymbol{\eta}_\mathcal{S}, \quad \boldsymbol{\eta}_\mathcal{S} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_\mathcal{S}), \tag{4}$$

where $\boldsymbol{R}_\mathcal{S}$ is the noise covariance. We consider equal-cost point observations in this work.

We jointly seek highly informative acquisition and reconstruction under the sensing budget Γ. Let $\mathcal{S}_\Gamma$ be the observation set induced by an acquisition policy $\pi_{\Phi,\psi}$, which may condition subsequent selections on the information already obtained. The population-level active grounding objective is

$$(\Phi^*, \psi^*) \in \arg\min_{\Phi,\psi} \mathbb{E}\left[\mathcal{L}\left(\boldsymbol{x}_0, \widehat{\boldsymbol{x}}_\Phi\left(\boldsymbol{c}, \boldsymbol{y}_{\mathcal{S}_\Gamma}, \mathcal{S}_\Gamma\right)\right)\right], \tag{5}$$

where Φ and ψ parameterize refinement and acquisition, respectively, $\widehat{\boldsymbol{x}}_\Phi$ is the target reconstruction, and $\mathcal{L}(\cdot,\cdot)$ measures reconstruction error.

The resulting problem thus couples *posterior refinement*, which determines how newly acquired measurements modify the structured belief induced by the diffusion prior, with *active acquisition*, which determines where measurements should be collected to maximize reconstruction benefit.

3 Method

Fig. 2 provides an overview of GPARA, which proceeds in two stages. During active acquisition, a cached prior ensemble specifies the current spatial risk, and the graph posterior provides a shared response for conditioning acquired measurements and evaluating prospective ones through weighted posterior-risk reduction, with a bounded residual scorer compensating for surrogate mismatch. Newly obtained measurements are thus incorporated through graph conditioning without regenerating diffusion samples. After the sensing budget is reached, final refinement applies the resulting graph-posterior conditioner to each prediction throughout the reverse diffusion process.

3.1 Graph-Structured Bayesian Residual Refinement

For the clean-state estimate $\widehat{\boldsymbol{x}}_{0|t}$ at reverse step t, we write the remaining discrepancy as

$$\boldsymbol{x}_0 = \widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{\delta}, \tag{6}$$

where $\boldsymbol{\delta} \in \mathbb{R}^M$ is the residual correction relative to the prediction. GPARA constructs a shared residual-correction geometry from a fixed graph-conditioning descriptor $\boldsymbol{c}_g$, including the context $\boldsymbol{c}$ and the known prior $p_\theta(\boldsymbol{x}_0 \mid \boldsymbol{c})$. A graph network G_ϕ parameterizes an undirected graph

$$\mathcal{G} = \big(\mathcal{V}, \mathcal{E}, \boldsymbol{C}_g\big), \quad \boldsymbol{C}_g = G_\phi(\boldsymbol{c}_g), \tag{7}$$

where node $i \in \mathcal{V}$ represents a target variable, $\mathcal{E}$ captures edges, and $\boldsymbol{C}_g = [c_{ij}] \in \mathbb{R}^{M\times M}$ is the symmetric and nonnegative adjacency matrix. A larger c_{ij} encourages measurement corrections to propagate coherently between coupled variables. Let $\boldsymbol{D}_g = \mathrm{diag}(\boldsymbol{C}_g\mathbf{1})$ and $\boldsymbol{L}_g = \boldsymbol{D}_g - \boldsymbol{C}_g$ be the degree matrix and Laplacian, respectively. We define a graph-Gaussian residual surrogate as

$$q_\Phi(\boldsymbol{\delta}|\boldsymbol{c}_g) = \mathcal{N}(\mathbf{0}, \boldsymbol{Q}_g^{-1}), \quad \boldsymbol{Q}_g = \tau\boldsymbol{I} + \lambda\boldsymbol{L}_g, \tag{8}$$

where $\tau > 0$ controls the residual precision and $\lambda \geq 0$ controls the graph coupling. The graph precision is shared across reverse steps as a reusable correction operator that is learned jointly across reverse timesteps, rather than parameterized separately for each noise level.

Proposition 1. *(Graph-induced residual posterior). Given the graph residual surrogate (8) and the observation model (4) with $\boldsymbol{R}_\mathcal{S} \succ 0$, the conditional residual posterior is*

$$q_\Phi\left(\boldsymbol{\delta} \mid \widehat{\boldsymbol{x}}_{0|t}, \boldsymbol{y}_\mathcal{S}, \boldsymbol{c}_g\right) = \mathcal{N}\left(\boldsymbol{m}_{\mathcal{S},t}, \boldsymbol{\Sigma}_\mathcal{S}\right), \tag{9}$$

where

$$\boldsymbol{m}_{\mathcal{S},t} = \boldsymbol{A}_\mathcal{S}^{-1}\boldsymbol{H}_\mathcal{S}^\mathsf{T}\boldsymbol{R}_\mathcal{S}^{-1}\left(\boldsymbol{y}_\mathcal{S} - \boldsymbol{H}_\mathcal{S}\widehat{\boldsymbol{x}}_{0|t}\right), \quad \boldsymbol{\Sigma}_\mathcal{S} = \boldsymbol{A}_\mathcal{S}^{-1}, \tag{10}$$

and

$$\boldsymbol{A}_\mathcal{S} = \boldsymbol{Q}_g + \boldsymbol{H}_\mathcal{S}^\mathsf{T}\boldsymbol{R}_\mathcal{S}^{-1}\boldsymbol{H}_\mathcal{S}. \tag{11}$$

Here, $\mathbf{m}_{\mathcal{S},t} \in \mathbb{R}^M$, $\boldsymbol{\Sigma}_\mathcal{S} \in \mathbb{R}^{M\times M}$ are the residual posterior mean and covariance conditioned on $\boldsymbol{y}_\mathcal{S}$. Moreover, $\boldsymbol{A}_\mathcal{S} \succ 0$, and the posterior is uniquely defined (see Appendix B.1 for the proof).

Proposition 1 establishes an explicit refinement mechanism that jointly performs observation consistency and graph propagation within a single update. While measurement innovations enter at observed locations, $\boldsymbol{A}_\mathcal{S}^{-1}$ propagates their effects to graph-coupled unobserved variables. Although $\boldsymbol{Q}_g$ is shared across reverse steps, $\boldsymbol{m}_{\mathcal{S},t}$ remains timestep dependent through the innovation.

The posterior mean also admits a variational interpretation

$$\boldsymbol{m}_{\mathcal{S},t} = \arg\min_{\boldsymbol{\delta}} \ \frac{\tau}{2}\|\boldsymbol{\delta}\|_2^2 + \frac{\lambda}{2}\sum_{(i,j)\in\mathcal{E}} c_{ij}(\delta_i - \delta_j)^2 + \frac{1}{2}\left\|\boldsymbol{H}_\mathcal{S}\left(\widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{\delta}\right) - \boldsymbol{y}_\mathcal{S}\right\|_{\mathbf{R}_\mathcal{S}^{-1}}^2, \tag{12}$$

where $\|\mathbf{v}\|_\mathbf{B}^2 = \mathbf{v}^\mathrm{T}\mathbf{B}\mathbf{v}$. It balances correction magnitude, graph-structured residual consistency, and measurement fidelity. The graph-posterior correction is then injected as

$$\widetilde{\boldsymbol{x}}_{0|t} = \Pi_{[x_{\min}, x_{\max}]}\left(\widehat{\boldsymbol{x}}_{0|t} + \rho_g\boldsymbol{m}_{\mathcal{S},t}\right), \tag{13}$$

where $\rho_g > 0$ controls correction strength, and $\Pi_{[\cdot,\cdot]}$ is element-wise clipping to the normalized range. The refined estimate $\widetilde{\boldsymbol{x}}_{0|t}$ then replaces the original $\widehat{\boldsymbol{x}}_{0|t}$ in the subsequent reverse transition.

3.2 Posterior-Aligned Active Acquisition

Let $\mathcal{S}_k$ be the observation set of acquisition round k. GPARA evaluates prospective measurements through the same graph-posterior covariance. Let $\boldsymbol{W}_k = \mathrm{diag}(w_{k,1}, ..., w_{k,M}) \succeq 0$ denote a diagonal risk-weighting matrix to specify the risk emphasis of target variables, as detailed in Appendix B.6. We define the weighted posterior risk as

$$\mathcal{R}(\mathcal{S}_k) = \mathrm{tr}\left(\boldsymbol{W}_k\boldsymbol{\Sigma}_{\mathcal{S}_k}\right). \tag{14}$$

While $\boldsymbol{\Sigma}_{\mathcal{S}_k}$ characterizes how a prospective measurement propagates across the target through cross-variable posterior dependence, $\boldsymbol{W}_k$ specifies where such reduction is currently most valuable.

Consider an unobserved candidate $i \notin \mathcal{S}_k$ with measurement noise variance σ_y^2, and let $\boldsymbol{e}_i \in \mathbb{R}^M$ be the ith canonical basis vector. Holding the current graph and risk weights, its one-step utility is

$$U_i(\mathcal{S}_k) = \mathcal{R}(\mathcal{S}_k) - \mathcal{R}(\mathcal{S}_k\cup\{i\}) = \frac{\boldsymbol{e}_i^\mathsf{T}\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{W}_k\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i}{\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}} = \underbrace{\frac{\Sigma_{\mathcal{S}_k,ii}^2}{\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}}}_{\text{local reducibility } \ell_i} \times \underbrace{\frac{\sum_{j=1}^M w_{k,j}\Sigma_{\mathcal{S}_k,ji}^2}{\Sigma_{\mathcal{S}_k,ii}^2}}_{\text{weight response reach } g_i}, \tag{15}$$

where $\Sigma_{\mathcal{S},ji}$ denotes the posterior covariance between node i and j. The local factor ℓ_i measures uncertainty removable at the candidate itself; g_i captures how broadly that information propagates through their weight correlations. A moderately uncertain location can be more valuable than a highly uncertain one if it influences a broader unresolved region. When $\boldsymbol{W}_k = \boldsymbol{I}$ and $\sigma_y^2 = 0$, it reduces to the LMMSE criterion used in AdaSense (Elata et al., 2024). We next characterize how the shared graph-posterior response relates posterior refinement to prospective measurement value.

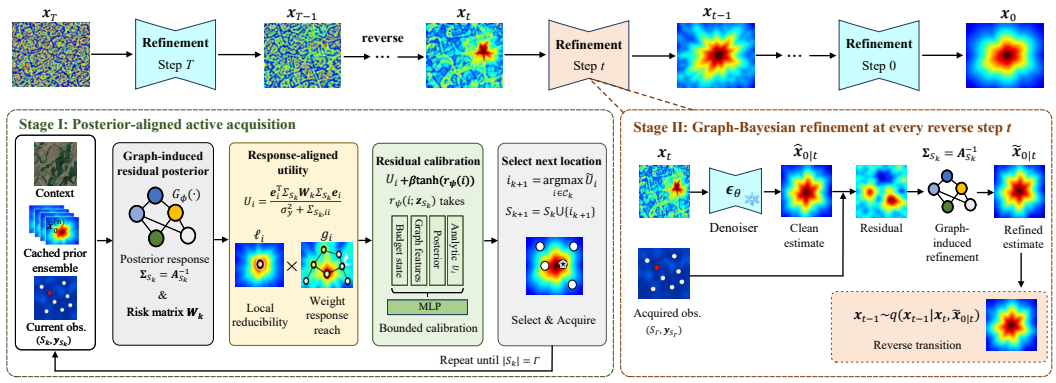


Figure 2: Overview of GPARA. Stage I performs response-aligned active acquisition: the graph posterior specifies how information propagates, while a state-adaptive risk matrix specifies where uncertainty reduction is valuable. A bounded residual scorer calibrates the analytic utility. After reaching the sensing budget, Stage II fixes the acquired observations and reuses the same graph-posterior conditioner to refine prediction throughout the reverse diffusion process.

Proposition 2. *(Posterior Propagation-Value Identity). For an arbitrary reverse step t, define the posterior response vector $\boldsymbol{u}_i = \boldsymbol{\Sigma}_{\mathcal{S}_k} \boldsymbol{e}_i / (\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii})$ and the prospective innovation $\nu_i = \boldsymbol{y}_i - \boldsymbol{e}_i^{\mathsf{T}} \left(\widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{m}_{\mathcal{S}_k,t} \right)$. The posterior-mean change induced by measuring i is thus $\Delta \widehat{\boldsymbol{x}}_i = \boldsymbol{u}_i \nu_i$, and*

$$U_i(\mathcal{S}_k) = (\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}) \, \|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2 = \mathbb{E}_{q_\Phi} \left[\|\Delta \widehat{\boldsymbol{x}}_i\|_{\boldsymbol{W}_k}^2 \mid \mathcal{F}_k \right],$$

where $\mathcal{F}_k$ denotes the acquisition state, including all known quantities such as $\mathbf{c}_g$ (see Appendix B.2).

Proposition 2 shows acquisition-refinement alignment: the same posterior response that propagates a newly acquired innovation also determines its expected weighted acquisition value.

Corollary 1. *(Matched-Model Ranking Consistency). For the unclipped graph-refined estimate $\widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k}$, define the one-step reconstruction improvement induced by acquiring candidate i as*

$$\Delta_i(\rho_g) = \left\| \boldsymbol{x}_0 - \widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k} \right\|_{\boldsymbol{W}_k}^2 - \left\| \boldsymbol{x}_0 - \widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k \cup \{i\}} \right\|_{\boldsymbol{W}_k}^2 .$$

Under the matched graph-Gaussian surrogate,

$$\mathbb{E}_{q_\Phi} \left[\Delta_i(\rho_g) \mid \mathcal{F}_k \right] = (2\rho_g - \rho_g^2) U_i(\mathcal{S}_k).$$

For $0 < \rho_g < 2$, maximizing $U_i(\mathcal{S}_k)$ thus maximizes its expected improvement (see Appendix B.3).

Corollary 1 further shows that, under the matched surrogate, the candidate with the largest utility $U_i(\mathcal{S}_k)$ is expected to maximize the one-step improvement in the same loss. This result motivates using such utility as the acquisition anchor. In the diffusion process, where clipping, finite-step denoising, approximate linear solves, and surrogate mismatch may break the matched-model assumptions, we use U_i as an explicit acquisition anchor and permit residual calibration from it.

Accordingly, GPARA retains $U_i(\mathcal{S}_k)$ as the analytic anchor and defines the final acquisition score as

$$\widetilde{U}_i(\mathcal{S}_k) = \bar{U}_i(\mathcal{S}_k) + \beta_{\text{acq}} \tanh \left(r_\psi(i; \mathbf{z}_{\mathcal{S}_k}) \right),$$

where $r_\psi(\cdot)$ is the scorer network, $\bar{U}_i$ is the normalized analytic utility, $\mathbf{z}_{\mathcal{S}_k}$ denotes the acquisition state, including context, graph-posterior statistics, observation information, and the budget, and $\beta_{\text{acq}} > 0$ bounds its magnitude. When $r_\psi(i; \mathbf{z}_{\mathcal{S}_k}) = 0$, it reduces exactly to the analytic anchor.

Starting from $\mathcal{S}_0 = \varnothing$, GPARA progressively acquires measurements according to the current calibrated utility. The underlying single-point selection rule is

$$i_{k+1} = \arg\max_{i \in \mathcal{C}_k} \widetilde{U}_i(\mathcal{S}_k), \quad \mathcal{S}_{k+1} = \mathcal{S}_k \cup \{i_{k+1}\},$$

where $\mathcal{C}_k \subseteq \mathcal{V} \setminus \mathcal{S}_k$ denotes the admissible unobserved candidates at step k. For computational efficiency, we additionally provide a batched acquisition strategy that selects a small group of candidates at each round k and jointly assimilates them before refreshing the posterior and acquisition scores. The single-point utility remains the primary acquisition criterion; batching is introduced only as a computational tradeoff to reduce update frequency. Further details are provided in Appendix B.7.

Algorithm 1 GPARA: Graph-Posterior-Aligned Active Acquisition and Reconstruction

Require: Context $\boldsymbol{c}$; frozen networks $\boldsymbol{\epsilon}_\theta$, G_ϕ, and r_ψ; sensing budget Γ; ensemble size L; number of diffusion samples N; measurement interface and noise model
Return: Reconstruction of final target $\hat{\boldsymbol{x}}_0$ and the acquired measurement set $\mathcal{S}_\Gamma$

Initialization
Set $k \leftarrow 0$, and $\mathcal{S}_0 \leftarrow \varnothing$. Generate $\mathbf{C}_g$ from G_ϕ, compute $\mathbf{L}_g$ and $\mathbf{Q}_g \leftarrow \tau\mathbf{I} + \lambda\mathbf{L}_g$
Generate and cache L DDPM prior samples $\{\widehat{\boldsymbol{x}}_0^{(l)}\}_{l=1}^{L}$ without measurements
Stage I: Posterior-aligned active acquisition
while $|\mathcal{S}_k| < \Gamma$ **do**
  Construct $\boldsymbol{A}_{\mathcal{S}_k} \leftarrow \mathbf{Q}_g + \mathbf{H}_{\mathcal{S}_k}^{\mathsf{T}}\mathbf{R}_{\mathcal{S}_k}^{-1}\mathbf{H}_{\mathcal{S}_k}$, and obtain $\{\widetilde{\boldsymbol{x}}_{0|\mathcal{S}_k}^{(l)}\}_{l=1}^{L}$
  Compute the ensemble mean and empirical variance, and construct $\boldsymbol{W}_k$ (see Appendix B.6)
  Estimate $U_i(\mathcal{S}_k)$ and covariance statistics using stochastic probing and PCG (see Appendix B.5)
  Compute, for each admissible candidate, $\widetilde{U}_i(\mathcal{S}_k) \leftarrow \bar{U}_i + \beta_{\mathrm{acq}} \tanh\big(r_\psi(i; \mathbf{z}_{\mathcal{S}_k})\big)$
  Select $i_{k+1} \leftarrow \arg\max_{i \in \mathcal{C}_k} \widetilde{U}_i(\mathcal{S}_k)$, acquire $y_{i_{k+1}}$ and $\mathcal{S}_{k+1} \leftarrow \mathcal{S}_k \cup \{i_{k+1}\}$
  $k \leftarrow k + 1$
end while
Stage II: Final graph-refined diffusion reconstruction
Reconstruct $\boldsymbol{A}_{\mathcal{S}_\Gamma}$ using the final observation set and keep it fixed throughout reconstruction
Initialize N reverse diffusion trajectories from Gaussian noise
for each reverse transition t **do**
  for $n = 1, \ldots, N$ **do**
    Obtain the clean-state prediction $\widehat{\boldsymbol{x}}_{0|t}^{(n)}$ from $\boldsymbol{\epsilon}_\theta(\hat{\boldsymbol{x}}_t^{(n)}, t, \mathbf{c})$
    Solve the graph-conditioning system for $\widehat{\mathbf{m}}_{\mathcal{S}_\Gamma,t}^{(n)}$ with $\mathbf{y}_{\mathcal{S}_\Gamma}$ by Proposition 1
    Update $\widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_\Gamma}^{(n)} \leftarrow \Pi_{[x_{\min}, x_{\max}]}\Big(\widehat{\boldsymbol{x}}_{0|t}^{(n)} + \rho_g \widehat{\mathbf{m}}_{\mathcal{S}_\Gamma,t}^{(n)}\Big)$
    Apply the DDPM transition with $\widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_\Gamma}^{(n)}$ to obtain $\hat{\boldsymbol{x}}_{t-1}^{(n)}$ by Eq. (3)
  end for
end for
Collect the terminal samples $\{\hat{\boldsymbol{x}}_0^{(n)}\}_{n=1}^{N}$ and compute $\hat{\boldsymbol{x}}_0 \leftarrow N^{-1}\sum_{n=1}^{N}\hat{\boldsymbol{x}}_0^{(n)}$
return $\hat{\boldsymbol{x}}_0, \mathcal{S}_\Gamma$

Implementation and training. GPARA consists of two networks trained offline. The graph predictor G_ϕ learns edge conductances through differentiable unrolled graph conditioning with reconstruction and graph-regularization objectives. The residual scorer r_ψ is trained to calibrate the analytic utility using a training-only one-step improvement teacher and distillation, ranking, and calibration losses. No parameters are updated at deployment. A small prior ensemble is generated once and cached for adaptive risk weighting. Posterior corrections use fixed-depth Jacobi iterations, and covariance-dependent utilities are evaluated matrix-free using stochastic probing and PCG without forming a dense posterior covariance matrix. Algorithm 1 summarizes the full procedure, with numerical settings and network training deferred to Appendix B and C.

4 Experiments

We evaluate GPARA on two representative tasks: radio field reconstruction (RFR) and RGB-guided depth completion (DC). Each task uses a separately pretrained DDPM as the frozen backbone, with a five-step reverse process for generation. All experiments are conducted on an NVIDIA RTX 4060 Ti GPU (16 GB) with an Intel Core i7-10700K CPU and 32 GB of RAM.

4.1 Experimental Setup

Datasets and metrics. RFR uses RadioMapSeer IRT4 (Yapar et al., 2024), with building geometry and transmitter configuration as context. Each target is a 128×128 radio map, with $1,122/140/140$ samples for training/validation/test. DC uses NYU Depth V2 (Silberman et al., 2012), with RGB images as context. The RGB-depth pairs are resized to 192×256, yielding $1,349$ training images and 100 test images from 35 unseen scenes. We report root mean squared error (RMSE) and mean absolute error (MAE) for both tasks, with peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) for RFR, and absolute relative error (AbsRel) and threshold accuracy δ_1 for DC.

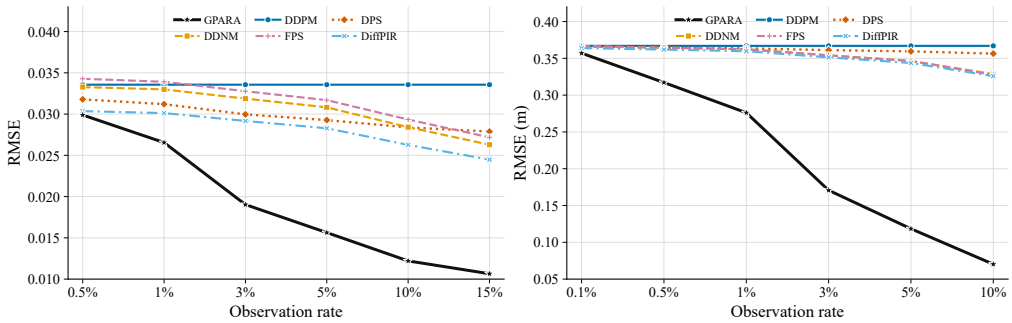


Figure 3: Posterior refinement across observation rates. RMSE as a function of the fraction of observed target locations for (a) RFR and (b) DC. All methods use the same frozen DDPM prior and identical randomly sampled observations at every observation rate.

Table 1: Posterior refinement under 491 matched random observations, approximately 3% of RFR and 1% of DC target entries. RMSE is reported as mean $\pm$ standard deviation across five independent seeds; other metrics show the mean values.

Method	Radio-field reconstruction				Depth completion			
	RMSE ↓	MAE ↓	SSIM ↑	PSNR ↑	RMSE ↓	MAE ↓	AbsRel ↓	δ_1 ↑
DDPM	0.035 ± 0.003	0.025	0.926	30.2	0.350 ± 0.005	0.273	0.113	0.878
DPS	0.031 ± 0.002	0.021	0.937	31.1	0.347 ± 0.005	0.271	0.112	0.881
DDNM	0.033 ± 0.003	0.023	0.919	30.6	0.344 ± 0.005	0.267	0.110	0.883
FPS	0.034 ± 0.003	0.024	0.921	30.4	0.345 ± 0.004	0.269	0.111	0.883
DiffPIR	0.030 ± 0.002	0.021	0.934	31.2	0.342 ± 0.004	0.266	0.110	0.885
GPARA (Ours)	$\mathbf{0.019 \pm 0.0007}$	**0.011**	**0.955**	**34.8**	$\mathbf{0.262 \pm 0.002}$	**0.196**	**0.081**	**0.951**

Baselines. Fixed-mask methods share the same pretrained DDPM within each task and receive identical random measurements. Baselines are DPS (Chung et al., 2023), DDNM (Wang et al., 2023), FPS (Dou & Song, 2024), and DiffPIR (Zhu et al., 2023). For active acquisition, we compare with ADS (Nolan et al., 2025) and AdaSense (Elata et al., 2024). AdaSense is paired with each posterior backend, including the proposed graph-Bayesian posterior.

4.2 Fixed-Observation Posterior Refinement

We first evaluate posterior refinement under identical randomly sampled observations.

Table 1 compares how each method uses the available evidence. On RFR, GPARA reduces RMSE and MAE by 36.7% and 47.6%, respectively, relative to DiffPIR, the strongest baseline here. On DC, GPARA reduces RMSE and AbsRel by 23.4% and 26.4%, respectively, while increasing δ_1 by 6.6 percentage points. Relative to DDPM prior, the RMSE reductions are 45.7% on RFR and 25.1% on DC. The consistent improvements indicate that the proposed graph-posterior refinement makes more effective use of the measurements. Note that all solvers are evaluated in an accelerated few-step regime, with additional results under larger reverse-step budgets reported in Appendix A.

Fig. 3 further evaluates RMSE across observation rates. It is observed that GPARA benefits more strongly from additional measurements than the competing baselines as the observation rate increases. On RFR, the gain over the strongest baseline grows from a modest margin in sparse regime to over 60% at higher rates; on DC, the gain similarly expands and reaches nearly 80%.

4.3 Posterior-Aligned Active Acquisition

We next allow each method to actively select measurements under the same sensing budget. "AdaSense + Graph" uses the proposed graph-based reconstruction backend. All methods are evaluated on the common prediction region shared across policies.

From Table 2, relative to the original DDPM prior, GPARA reduces RMSE by 60.0% on RFR and 32.6% on DC. Compared with random sensing, response-aligned acquisition provides a further 26.3% and 9.9% RMSE reduction, respectively. GPARA also outperforms the strongest baseline by

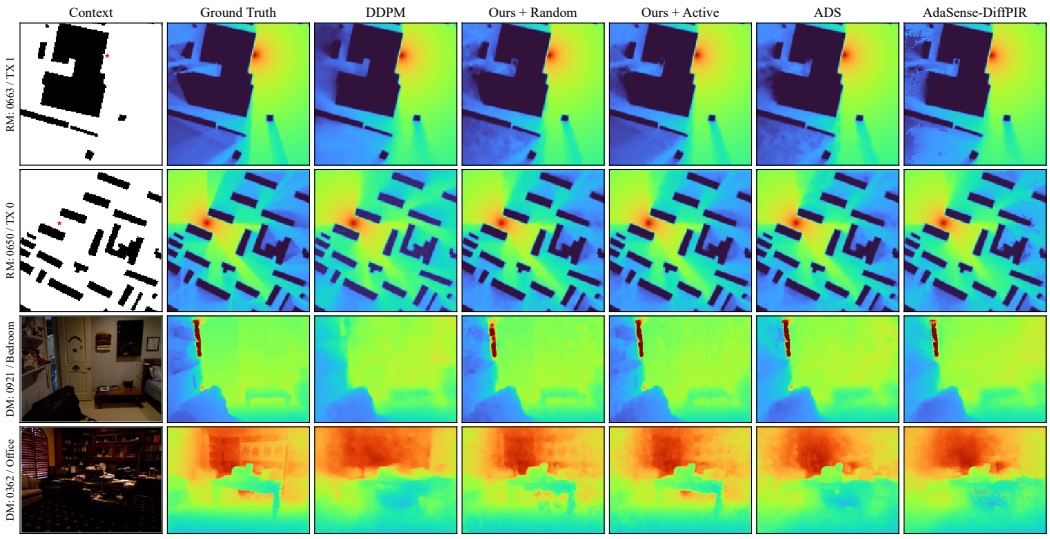


Figure 4: Qualitative comparison of active acquisition and reconstruction. The first two rows show RFR examples with 10% observations, and the last two rows are DC examples under 5%.

Table 2: Active acquisition with a matched 491-point sensing budget across five independent seeds.

Method	Radio-field reconstruction				Depth completion			
	RMSE ↓	MAE ↓	SSIM ↑	PSNR ↑	RMSE ↓	MAE ↓	AbsRel ↓	δ_1 ↑
ADS	0.026 ± 0.001	0.018	0.947	33.1	0.338 ± 0.002	0.267	0.111	0.882
AdaSense + DPS	0.026 ± 0.001	0.018	0.947	33.2	0.339 ± 0.002	0.268	0.111	0.881
AdaSense + DDNM	0.028 ± 0.002	0.020	0.937	32.7	0.336 ± 0.002	0.267	0.110	0.883
AdaSense + FPS	0.029 ± 0.002	0.021	0.935	32.5	0.336 ± 0.002	0.266	0.110	0.883
AdaSense + DiffPIR	0.026 ± 0.001	0.018	0.948	33.4	0.334 ± 0.002	0.264	0.110	0.885
AdaSense + Graph	0.018 ± 0.0006	0.013	0.956	36.4	0.268 ± 0.007	0.211	0.089	0.933
GPARA (Ours)	$\mathbf{0.014 \pm 0.0004}$	**0.009**	**0.962**	**38.1**	$\mathbf{0.236 \pm 0.006}$	**0.187**	**0.079**	**0.958**

46.2% and 29.3% on RMSE. Importantly, compared with AdaSense+Graph under the same backend, GPARA still achieves 22.2% and 11.9% lower RMSE, respectively, supporting the benefit of the proposed response-aligned acquisition strategy beyond the reconstruction backend itself.

The qualitative examples in Fig. 4 indicate that GPARA more faithfully recovers spatial variations and structural boundaries, with fewer reconstruction artifacts. The improvement of Ours + Active over Ours + Random further illustrates that the gain arises not only from posterior refinement, but also from acquiring more informative measurements. These examples should be interpreted as visual support for the quantitative trends rather than as independent evidence of general superiority.

Fig. 5 evaluates measurement efficiency across observation rates on RFR. GPARA achieves the lowest MAE at every observation rate. The gap over the conventional active baselines widens rapidly beyond the sparsest regime. AdaSense + Graph also improves substantially as the budget increases, confirming the benefit of the shared graph reconstruction backend.

Table 3 evaluates acquisition-stage computational cost on DC. Implementation configuration of GPARA is given in Appendix B.5. It requires only 20 sample-equivalent NFEs, substantially reducing dependence on repeated denoiser evaluations, but posterior computations remain a non-negligible part of its runtime. While slower than ADS in wall-clock time, it uses 54.6% less peak memory, remains comparable to the AdaSense variants, and is $1.71\times$ faster than AdaSense + Graph.

4.4 Mechanism and Ablation Studies

Posterior refinement mechanism. Table 4(a) separates the effects of graph coupling and correction placement under matched random observations. With step-wise injection, a uniform graph reduces RMSE by 40.4% on RFR and 22.8% on DC relative to pointwise conditioning. Learning the edge conductances provides a further 2.1% and 1.9% reduction under the same injection sched-

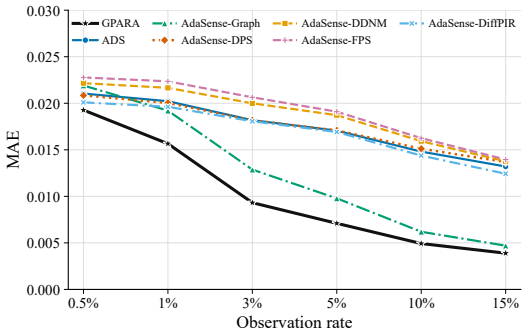


Figure 5: Active acquisition performance over the common prediction region across varying observation rates on RFR.

Table 3: Acquisition efficiency on DC. Time denotes acquisition-stage wall-clock time per image (mean ± standard deviation); NFE is the sample-equivalent number of denoiser evaluations, and Mem. is peak GPU memory usage.

Method	Time (s/img)	NFE	Mem. (MiB)
ADS	$\mathbf{3.14 \pm 0.24}$	252	1493.6
AdaS.+DPS	14.7 ± 0.21	1240	1497.4
AdaS.+DDNM	6.72 ± 0.15	1240	664.4
AdaS.+FPS	6.87 ± 0.15	1240	668.9
AdaS.+DiffPIR	6.88 ± 0.15	1240	663.6
AdaS.+Graph	11.5 ± 0.27	1240	671.7
GPARA	$\mathbf{6.73 \pm 0.18}$	**20**	677.9

Table 4: Mechanism ablations of GPARA. (a) Compares posterior mechanisms under identical random observations. (b) Compares acquisition criteria while keeping others fixed.

(a) Posterior refinement mechanism

Model	Injection	RFR	DC
Point Bayes	Step-wise	0.0327	0.346
Uniform graph	Final-only	0.0240	0.309
Learn Graph	Final-only	0.0233	0.307
Uniform graph	Step-wise	0.0195	0.267
Learn Graph	Step-wise	**0.0191**	**0.262**

(b) Acquisition utility

Acquisition score	RFR	DC
Local reducibility ℓ_i	0.0246	0.268
Weighted analytic U_i	0.0157	0.256
Unweighted U_i ($\boldsymbol{W}_k = \boldsymbol{I}$)	0.0163	0.263
Direct learned score U_i^{ln}	0.0174	0.268
Final calibrated $\widetilde{U}_i$	**0.0142**	**0.236**

ule. Moving from final-only to step-wise correction improves the learned-graph variant by 18.0% and 14.7%, respectively. These comparisons show that substantial improvements can be obtained through graph coupling and repeated measurement correction, while learned conductances yield smaller but consistent additional improvement gains.

Acquisition-score components. Table 4(b) evaluates the acquisition-score components with a shared reconstruction backend. Replacing local reducibility with the unweighted utility reduces RMSE by 33.7% on RFR and 1.9% on DC. Risk weighting provides a further 3.7% and 2.7% reduction, while adding residual calibration improves the weighted analytic policy by 9.6% and 7.8%, respectively. To assess the value of the explicit analytic anchor, we additionally train a direct learn scorer U_i^{ln} that predicts the candidate score without an additive analytic term, using the same loss functions as the calibrated scorer. Relative to this baseline, the weighted analytic policy reduces RMSE by 9.8% on RFR and 4.5% on DC, whereas the calibrated policy achieves reductions of 18.4% and 11.9%. These results support retaining the analytic posterior-response anchor and learning residual corrections, rather than learning the entire acquisition score directly.

5 Conclusion

We introduced GPARA for active grounding with frozen diffusion priors. The method learns a graph surrogate over prediction residuals and reuses its response structure for measurement refinement and candidate evaluation. Under a matched surrogate, the theoretical analysis relates weighted posterior-risk reduction to expected one-step refinement benefit. The implementation combines cached prior predictions, sparse conditioning, and bounded score calibration to avoid repeated diffusion sample generation during acquisition. Experiments on radio-field reconstruction and depth completion show improvements in both refinement and active sensing. Ablations further show that step-wise structured refinement outperforms terminal correction and that the weighted analytic acquisition criterion is stronger than a learned-only scorer, supporting the role of the proposed posterior-response mechanism beyond supervised policy capacity. The current analysis is limited to the graph-Gaussian surrogate and one-step acquisition, while nonlinear observations and longer-horizon sensing remain important directions for future work.

AI USE STATEMENT

We used OpenAI ChatGPT at several stages of this work. For writing and manuscript preparation, it assisted with outlining the paper, drafting and reorganizing sections, and paragraph-level language polishing. For retrieval and discovery, it was used to identify relevant literature, summarize related methods, and compare their assumptions and methodological differences with ours; all cited references and claims were checked by the authors against the original sources. For mathematical claims and proofs, it assisted with checking formulations, derivations, and proof steps; all mathematical claims and proofs were verified by the authors. For experiment execution, it assisted with experimental design, code generation and debugging, and interpretation of experimental results. We did not use generative AI to generate synthetic datasets or results. All AI-assisted code, mathematical claims, experimental procedures, citations, and manuscript text were developed under the authors' supervision rather than fully automated generation. The authors reviewed, revised, and verified all AI-assisted outputs and take full responsibility for the final content of this work.

ETHICS STATEMENT

This work uses the publicly available radio-field and RGB-depth datasets and involves neither new human-subject data collection nor collection of personally identifiable information.

REPRODUCIBILITY STATEMENT

We provide the methodological and experimental details required to reproduce GPARA. Section 3 specifies the graph-posterior formulation, active-acquisition criterion, and the complete acquisition-reconstruction pipeline summarized in Algorithm 1. Appendix B provides the numerical details underlying the method. The learned graph predictor and residual scorer are generic function approximators rather than architecture-specific components of the formulation; Appendix C specifies their functional interfaces and staged training objectives. Section 4 describes the datasets, sensing budgets, evaluation metrics, baseline configurations, and experimental protocols. All comparisons use the same frozen diffusion backbones. The anonymized source code of this work is also available at `https://anonymous.4open.science/r/29873-GPARA/`.

REFERENCES

Asad Aali, Giannis Daras, Brett Levac, Sidharth Kumar, Alexandros G. Dimakis, and Jonathan I. Tamir. Ambient diffusion posterior sampling: Solving inverse problems with diffusion models trained on corrupted data. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/2dd7f33ffbb59b4ff987be5442a13016-Abstract-Conference.html.

Luis Barba, Johannes Kirschner, Tomas Aidukas, Manuel Guizar-Sicairos, and Benjamín Béjar. Diffusion active learning: Towards data-driven experimental design in computed tomography, 2025. URL https://arxiv.org/abs/2504.03491.

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

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.

Zehao Dou and Yang Song. Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/333bbde5ed6397cd9bc5bac552b3b7a6-Abstract-Conference.html.

Noam Elata, Tomer Michaeli, and Michael Elad. Adaptive compressed sensing with diffusion-based posterior sampling. In *Computer Vision – ECCV 2024*, volume 15136 of *Lecture Notes in Computer Science*, pp. 290–308, 2024. doi: 10.1007/978-3-031-73229-4_17. URL https://eccv.ecva.net/virtual/2024/poster/315.

Ben Fei, Zhaoyang Lyu, Liang Pan, Junzhe Zhang, Weidong Yang, Tianyue Luo, Bo Zhang, and Bo Dai. Generative diffusion prior for unified image restoration and enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9935–9946, 2023. doi: 10.1109/CVPR52729.2023.00958.

Adam Foster, Martin Jankowiak, Eli Bingham, Paul Horsfall, Yee Whye Teh, Tom Rainforth, and Noah Goodman. Variational Bayesian optimal experimental design. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/hash/d55cbf210f175f4a37916eafe6c04f0d-Abstract.html.

Adam Foster, Desi R. Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential Bayesian experimental design. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3384–3395, 2021. URL https://proceedings.mlr.press/v139/foster21a.html.

Alexandros Graikos, Esmeralda S Whitammer, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors. In *Advances in Neural Information Processing Systems*, volume 35, pp. 14715–14728. Curran Associates, Inc., 2022. doi: 10.52202/068431-1070. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/5e6cec2a9520708381fe520246018e8b-Paper-Conference.pdf.

Yuntong Gu, Xiangming Meng, Zhiyuan Lin, Sheng Wu, and Linling Kuang. Generative spectrum cartography: Unified reconstruction and active sensing via diffusion models, 2025. URL https://arxiv.org/abs/2512.20108.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html.

Jacopo Iollo, Christophe Heinkelé, Pierre Alliez, and Florence Forbes. Bayesian experimental design via contrastive diffusions. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/b5cc526f12164b2144bb2e06f2e84864-Abstract-Conference.html.

Desi R Ivanova, Adam Foster, Steven Kleinegesse, Michael Gutmann, and Thomas Rainforth. Implicit deep adaptive design: Policy-based experimental design without likelihoods. In *Advances in Neural Information Processing Systems*, volume 34, pp. 25785–25798. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/d811406316b669ad3d370d78b51b1d2e-Paper.pdf.

Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 23593–23606, 2022. doi: 10.52202/068431-1714. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/95504595b6169131b6ed6cd72eb05616-Abstract-Conference.html.

Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9492–9502, 2024. doi: 10.1109/CVPR52733.2024.00907.

Samar Khanna, Patrick Liu, Linqi Zhou, Chenlin Meng, Robin Rombach, Marshall Burke, David B. Lobell, and Stefano Ermon. DiffusionSat: A generative foundation model for satellite imagery. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=I5webNFDgQ.

Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008. URL https://www.jmlr.org/papers/v9/krause08a.html.

Zeyu Li, Hongkun Dou, Shen Fang, Wang Han, Yue Deng, and Lijun Yang. Physics-aligned field reconstruction with diffusion bridge. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/9421261e06f1a63a352b068f1ac90609-Abstract-Conference.html.

Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/7711026951f3dc3afb2be427bf98bc73-Abstract-Conference.html.

Amr Morssy, Marcus R. Frean, and Paul D. Teal. Informed adaptive sensing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3230–3241, 2024. doi: 10.1109/TPAMI.2023.3340990.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8162–8171. PMLR, 2021. URL https://proceedings.mlr.press/v139/nichol21a.html.

Oisín Nolan, Tristan S. W. Stevens, Wessel L. van Nierop, and Ruud J. G. van Sloun. Active diffusion subsampling. *Transactions on Machine Learning Research*, 2025. URL https://openreview.net/forum?id=OGifiton47.

Orazio Pinti, Jeremy M. Budd, Franca Hoffmann, and Assad A. Oberai. Graph Laplacian-based Bayesian multi-fidelity modeling. *Computer Methods in Applied Mechanics and Engineering*, 435:117647, 2025. doi: 10.1016/j.cma.2024.117647. URL https://doi.org/10.1016/j.cma.2024.117647.

Konstantinos D. Polyzos, Alireza Sadeghi, Wei Ye, Steven Sleder, Kodjo Houssou, Jeff Calder, Zhi-Li Zhang, and Georgios B. Giannakis. Bayesian active learning for sample efficient 5G radio map reconstruction. *IEEE Transactions on Wireless Communications*, 23(12):19382–19396, 2024. doi: 10.1109/TWC.2024.3483112. URL https://doi.org/10.1109/TWC.2024.3483112.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022. doi: 10.1109/CVPR52688.2022.01042. URL https://openaccess.thecvf.com/content/CVPR2022/html/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.html.

Litu Rout, Yujia Chen, Abhishek Kumar, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng Chu. Beyond first-order Tweedie: Solving inverse problems using latent diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9472–9481, June 2024. URL https://openaccess.thecvf.com/content/CVPR2024/html/Rout_Beyond_First-Order_Tweedie_Solving_Inverse_Problems_using_Latent_Diffusion_CVPR_2024_paper.html.

Håvard Rue and Leonhard Held. *Gaussian Markov Random Fields: Theory and Applications*. Chapman and Hall/CRC, 2005. doi: 10.1201/9780203492024.

David I Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013. doi: 10.1109/MSP.2012.2235192.

Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *Computer Vision – ECCV 2012*, pp. 746–760, 2012. doi: 10.1007/978-3-642-33715-4_54.

Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse problems with latent diffusion models via hard data consistency. In

International Conference on Learning Representations, volume 2024, pp. 7624–7654, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/1f09e1ee5035a4c3fe38a5681cae5815-Paper-Conference.pdf.

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

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*, 2021. URL https://openreview.net/forum?id=PxTIG12RRHS.

Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=mRieQgMtNTQ.

Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. SeeSR: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25456–25467, 2024a. URL https://openaccess.thecvf.com/content/CVPR2024/html/Wu_SeeSR_Towards_Semantics-Aware_Real-World_Image_Super-Resolution_CVPR_2024_paper.html.

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. In *Advances in Neural Information Processing Systems*, volume 37, pp. 118389–118427. Curran Associates, Inc., 2024b. doi: 10.52202/079017-3759. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/d65c4ce22241138c1784ff753d4c746c-Paper-Conference.pdf.

Çağkan Yapar, Ron Levie, Gitta Kutyniok, and Giuseppe Caire. Dataset of pathloss and ToA radio maps with localization application, 2024. URL https://arxiv.org/abs/2212.11777.

Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20895–20905, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Zhang_Improving_Diffusion_Inverse_Problem_Solving_with_Decoupled_Noise_Annealing_CVPR_2025_paper.html.

Zizhao Zhang, Adriana Romero, Matthew J. Muckley, Pascal Vincent, Lin Yang, and Michal Drozdzal. Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2049–2053, 2019. doi: 10.1109/CVPR.2019.00215.

Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1219–1229, 2023. doi: 10.1109/CVPRW59228.2023.00129. URL https://openaccess.thecvf.com/content/CVPR2023W/NTIRE/html/Zhu_Denoising_Diffusion_Models_for_Plug-and-Play_Image_Restoration_CVPRW_2023_paper.html.

Rayhan Zirvi, Bahareh Tolooshams, and Anima Anandkumar. Diffusion state-guided projected gradient for inverse problems. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/4d5b7f7a0603b75e39ffebbf95e3b7ee-Abstract-Conference.html.

A ADDITIONAL RESULTS

Table 5: Sensitivity to the number of reverse diffusion steps on RFR under 3% random observations. Results are full-region RMSE (mean $\pm$ standard deviation) over five paired seeds.

Method	5 steps	10 steps	20 steps	50 steps	100 steps	500 steps
DDPM	0.035 ± 0.0034	0.030 ± 0.0006	0.030 ± 0.0014	0.030 ± 0.0010	0.030 ± 0.0009	0.031 ± 0.0009
DPS	0.031 ± 0.0023	0.028 ± 0.0006	0.028 ± 0.0013	0.027 ± 0.0009	0.024 ± 0.0003	$\mathbf{0.018 \pm 0.0002}$
DDNM	0.033 ± 0.0031	0.028 ± 0.0005	0.028 ± 0.0013	0.027 ± 0.0009	0.027 ± 0.0007	0.027 ± 0.0009
FPS	0.034 ± 0.0033	0.027 ± 0.0005	0.027 ± 0.0009	0.026 ± 0.0002	0.025 ± 0.0003	0.022 ± 0.0004
DiffPIR	0.030 ± 0.0021	0.027 ± 0.0006	0.026 ± 0.0008	0.026 ± 0.0003	0.025 ± 0.0003	0.021 ± 0.0004
GPARA	$\mathbf{0.019 \pm 0.0007}$	$\mathbf{0.017 \pm 0.0002}$	$\mathbf{0.016 \pm 0.0003}$	$\mathbf{0.016 \pm 0.0002}$	$\mathbf{0.016 \pm 0.0001}$	$\mathbf{0.016 \pm 0.0002}$

Table 6: Alignment between acquisition scores and empirical one-step reconstruction benefit under shared random prefixes.

Score	RFR		DC	
	Spearman $\uparrow$	Regret ($\times 10^{-5}$) $\downarrow$	Spearman $\uparrow$	Regret ($\times 10^{-4}$) $\downarrow$
High-accuracy analytic	0.42 ± 0.26	1.26 ± 1.45	0.19 ± 0.27	9.23 ± 15.60
Deployed analytic	0.35 ± 0.27	1.14 ± 1.35	0.11 ± 0.25	9.51 ± 15.52
Calibrated utility	0.40 ± 0.25	1.06 ± 1.25	0.15 ± 0.24	$\mathbf{3.60 \pm 4.97}$
Direct learned scorer	$\mathbf{0.44 \pm 0.29}$	$\mathbf{0.94 \pm 1.28}$	$\mathbf{0.23 \pm 0.27}$	6.08 ± 14.75

Sensitivity to reverse diffusion steps. Table 5 evaluates sensitivity to the reverse-step budget. GPARA reaches its best performance within roughly 20 steps, whereas several baselines continue to improve markedly as more reverse steps are added; for example, DPS decreases from 0.031 RMSE at five steps to 0.018 at 500 steps. This suggests that conventional inverse solvers rely more heavily on repeated denoising and data-consistency updates to accumulate and propagate measurement information. In contrast, the explicit graph-posterior response in GPARA enables effective cross-location propagation with substantially fewer reverse steps, making the method more compatible with accelerated sampling. These results suggest that the proposed explicit posterior propagation reduces sensitivity to the denoising-step budget, although this ablation isolates reconstruction accuracy rather than end-to-end wall-clock efficiency.

One-step diagnostic protocol. We evaluate whether each acquisition score predicts the empirical benefit of the next measurement under shared random observation prefixes. For each task, we use 20 test instances and construct three prefixes per instance at approximately 10%, 50%, and 90% of the final sensing budget, yielding 60 matched prefix states. At each prefix $\mathcal{S}$, all scorers evaluate the same randomly sampled candidate set $\mathcal{C}^{\mathrm{rand}}$ with $|\mathcal{C}^{\mathrm{rand}}| = 16$. Each candidate is acquired independently and evaluated using four paired runs of the complete Stage-II reconstruction. Its empirical benefit is $g_i = E_{\mathrm{pre}} - E_i$, where E_{pre} and E_i denote reconstruction MSE before and after measuring candidate i, respectively, over a common evaluation region. For candidate scores s_i and the selected location $a = \arg\max_{i \in \mathcal{C}} s_i$, we compute

$$\rho_s = \mathrm{Corr}\left(\mathrm{rank}_{\mathcal{C}}(s_i), \mathrm{rank}_{\mathcal{C}}(g_i)\right), \quad \mathrm{Regret} = \max_{i \in \mathcal{C}} g_i - g_a = E_a - \min_{i \in \mathcal{C}} E_i.$$

Here, $\mathrm{Corr}(\cdot)$ denotes Pearson correlation applied to within-prefix ranks, with ties assigned average ranks. It measures candidate-ranking agreement, while regret quantifies the selected action's excess MSE over the empirically best candidate. Both metrics are computed per prefix and then aggregated.

The deployed analytic score uses the same numerical settings as the main experiments. The high-accuracy analytic diagnostic uses PCG with up to 96 iterations and tolerance 10^{-7}. Table 6 shows that calibration improves the deployed analytic score on average, increasing Spearman correlation from 0.35 to 0.40 on RFR and from 0.11 to 0.15 on DC, while reducing regret. The direct learned scorer achieves higher one-step rank on both tasks and the lowest RFR regret, whereas the calibrated achieves the lowest DC regret. These diagnostics assess only a single decision from a shared state; they do not directly predict full-budget closed-loop performance. Indeed, over complete acquisition, the calibrated policy achieves lower final RMSE than the directly learned scorer.

These diagnostics evaluate a single next-measurement decision from shared exogenous states, whereas Table 4 evaluates full-budget policy rollouts whose future states depend on previous selections. The random-prefix distribution is also closer to the randomly sampled observation states used during scorer training, while closed-loop policies induce structured, policy-dependent states. Moreover, each acquisition changes the posterior and therefore all subsequent decisions. The analytic anchor explicitly recomputes marginal value and redundancy from the updated posterior geometry at every refresh, whereas the direct scorer must learn these effects from supervised states. These differences provide plausible explanations for why the calibrated policy achieves lower final RMSE than the direct scorer despite not uniformly dominating it in the one-step diagnostics.

B Proofs and Numerical Details

B.1 Proof of Proposition 1

Positive definiteness of the posterior precision. For a symmetric graph with nonnegative conductances $c_{ij} = c_{ji} \geq 0$ and every $\boldsymbol{v} \in \mathbb{R}^M$, the graph Laplacian satisfies

$$\boldsymbol{v}^\top \boldsymbol{L}_g \boldsymbol{v} = \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} c_{ij} (v_i - v_j)^2 \geq 0. \tag{B.1}$$

Hence for $\tau > 0$ and $\lambda \geq 0$, $\boldsymbol{L}_g \succeq 0$ and $\boldsymbol{Q}_g \succeq \tau \boldsymbol{I} \succ 0$. Since $\boldsymbol{R}_{\mathcal{S}} \succ 0$, for any nonzero $\boldsymbol{v}$, we have

$$\boldsymbol{v}^\top \boldsymbol{A}_{\mathcal{S}} \boldsymbol{v} = \tau \|\boldsymbol{v}\|_2^2 + \lambda \boldsymbol{v}^\top \boldsymbol{L}_g \boldsymbol{v} + \|\boldsymbol{H}_{\mathcal{S}} \boldsymbol{v}\|_{\boldsymbol{R}_{\mathcal{S}}^{-1}}^2 \geq \tau \|\boldsymbol{v}\|_2^2 > 0. \tag{B.2}$$

Thus $\boldsymbol{A}_{\mathcal{S}} \succeq \tau \boldsymbol{I} \succ 0$. This does not require the graph to be connected.

Conditional Gaussian posterior. The residual $\boldsymbol{x}_0 = \widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{\delta}$ and the observation model imply

$$\boldsymbol{y}_{\mathcal{S}} \mid \boldsymbol{\delta}, \widehat{\boldsymbol{x}}_{0|t}, \boldsymbol{c} \sim \mathcal{N}\left(\boldsymbol{H}_{\mathcal{S}} \left(\widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{\delta}\right), \boldsymbol{R}_{\mathcal{S}}\right). \tag{B.3}$$

Combining this likelihood with $q_\Phi(\boldsymbol{\delta} \mid \boldsymbol{c}_g) = \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}_g^{-1})$ gives

$$q_\Phi\left(\boldsymbol{\delta} \mid \widehat{\boldsymbol{x}}_{0|t}, \boldsymbol{y}_{\mathcal{S}}, \boldsymbol{c}_g\right) \propto \exp\left(-\frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{Q}_g \boldsymbol{\delta} - \frac{1}{2} \left\|\boldsymbol{y}_{\mathcal{S}} - \boldsymbol{H}_{\mathcal{S}} \left(\widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{\delta}\right)\right\|_{\boldsymbol{R}_{\mathcal{S}}^{-1}}^2\right). \tag{B.4}$$

Expanding the exponent and omitting terms independent of $\boldsymbol{\delta}$, the negative log-density is

$$\frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{A}_{\mathcal{S}} \boldsymbol{\delta} - \boldsymbol{\delta}^\top \boldsymbol{H}_{\mathcal{S}}^\top \boldsymbol{R}_{\mathcal{S}}^{-1} \left(\boldsymbol{y}_{\mathcal{S}} - \boldsymbol{H}_{\mathcal{S}} \widehat{\boldsymbol{x}}_{0|t}\right) + \text{const}. \tag{B.5}$$

Define $\boldsymbol{m}_{\mathcal{S},t} = \boldsymbol{A}_{\mathcal{S}}^{-1} \boldsymbol{H}_{\mathcal{S}}^\top \boldsymbol{R}_{\mathcal{S}}^{-1} \left(\boldsymbol{y}_{\mathcal{S}} - \boldsymbol{H}_{\mathcal{S}} \widehat{\boldsymbol{x}}_{0|t}\right)$. Completing the square transforms (B.5) into

$$\frac{1}{2} \left(\boldsymbol{\delta} - \boldsymbol{m}_{\mathcal{S},t}\right)^\top \boldsymbol{A}_{\mathcal{S}} \left(\boldsymbol{\delta} - \boldsymbol{m}_{\mathcal{S},t}\right) + \text{const}. \tag{B.6}$$

Because $\boldsymbol{A}_{\mathcal{S}} \succ 0$, this expression defines a proper Gaussian distribution. Consequently,

$$q_\Phi\left(\boldsymbol{\delta} \mid \widehat{\boldsymbol{x}}_{0|t}, \boldsymbol{y}_{\mathcal{S}}, \boldsymbol{c}_g\right) = \mathcal{N}\left(\boldsymbol{m}_{\mathcal{S},t}, \boldsymbol{\Sigma}_{\mathcal{S}}\right), \quad \boldsymbol{\Sigma}_{\mathcal{S}} = \boldsymbol{A}_{\mathcal{S}}^{-1}. \tag{B.7}$$

This proves the posterior formulas and uniqueness in Proposition 1.

B.2 Proof of Proposition 2

Consider the current observation set $\mathcal{S}_k$ and a candidate i. The prospective measurement is

$$y_i = \boldsymbol{e}_i^\top \boldsymbol{x}_0 + \eta_i, \quad \eta_i \sim \mathcal{N}(0, \sigma_y^2), \tag{B.8}$$

where $\sigma_y^2 > 0$, $\boldsymbol{e}_i$ is the ith canonical basis vector, and η_i is independent of the residual and noise.

Posterior-mean response to a new measurement. The prospective innovation can be written as

$$\nu_i = y_i - \boldsymbol{e}_i^\top \left(\widehat{\boldsymbol{x}}_{0|t} + \boldsymbol{m}_{\mathcal{S}_k,t}\right) = \boldsymbol{e}_i^\top \left(\boldsymbol{\delta}_t - \boldsymbol{m}_{\mathcal{S}_k,t}\right) + \eta_i. \tag{B.9}$$

Its conditional moments are

$$\mathbb{E}_{q_\Phi}\left[\nu_i \mid \mathcal{F}_k\right] = 0, \quad \mathbb{E}_{q_\Phi}\left[\nu_i^2 \mid \mathcal{F}_k\right] = \sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}, \tag{B.10}$$

and

$$\mathbb{E}_{q_\Phi}\left[\left(\boldsymbol{\delta}_t - \boldsymbol{m}_{\mathcal{S}_k,t}\right)\nu_i \mid \mathcal{F}_k\right] = \boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i. \tag{B.11}$$

Since the residual and the prospective measurement are jointly Gaussian, conditioning on y_i gives

$$\boldsymbol{m}_{\mathcal{S}_k\cup\{i\},t} = \boldsymbol{m}_{\mathcal{S}_k,t} + \frac{\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i}{\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}}\nu_i. \tag{B.12}$$

Hence the target posterior-mean change equals the residual posterior-mean change:

$$\Delta\widehat{\boldsymbol{x}}_i = \boldsymbol{m}_{\mathcal{S}_k\cup\{i\},t} - \boldsymbol{m}_{\mathcal{S}_k,t} = \boldsymbol{u}_i\nu_i, \quad \boldsymbol{u}_i = \frac{\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i}{\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}}. \tag{B.13}$$

Covariance-risk reduction. The additional independent point measurement updates the precision

$$\boldsymbol{A}_{\mathcal{S}_k\cup\{i\}} = \boldsymbol{A}_{\mathcal{S}_k} + \sigma_y^{-2}\boldsymbol{e}_i\boldsymbol{e}_i^\top. \tag{B.14}$$

Applying the rank-one inverse identity yields

$$\boldsymbol{\Sigma}_{\mathcal{S}_k\cup\{i\}} = \boldsymbol{\Sigma}_{\mathcal{S}_k} - \frac{\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i\boldsymbol{e}_i^\top\boldsymbol{\Sigma}_{\mathcal{S}_k}}{\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}}. \tag{B.15}$$

Using the same current weighting matrix $\boldsymbol{W}_k$ on both sides of the hypothetical update,

$$U_i(\mathcal{S}_k) = \operatorname{tr}\left(\boldsymbol{W}_k\left(\boldsymbol{\Sigma}_{\mathcal{S}_k} - \boldsymbol{\Sigma}_{\mathcal{S}_k\cup\{i\}}\right)\right) = \frac{\boldsymbol{e}_i^\top\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{W}_k\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i}{\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}} = \left(\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}\right)\|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2. \tag{B.16}$$

The second equality follows from cyclic invariance of the trace. Since $\boldsymbol{W}_k \succeq 0$, the utility is nonnegative. Finally, $\boldsymbol{u}_i$ is fixed given the current state. Combining (B.10) and (B.13),

$$\mathbb{E}_{q_\Phi}\left[\|\Delta\widehat{\boldsymbol{x}}_i\|_{\boldsymbol{W}_k}^2 \mid \mathcal{F}_k\right] = \mathbb{E}_{q_\Phi}\left[\nu_i^2 \mid \mathcal{F}_k\right]\|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2 = \left(\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}\right)\|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2. \tag{B.17}$$

The identity establishes that the covariance reduction and the expected squared posterior-mean change describe the same one-step effect under the surrogate. This proves Proposition 2.

B.3 Proof of Corollary 1

For the unclipped comparison in Corollary 1, the two uniformly scaled estimates are

$$\widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k} = \widehat{\boldsymbol{x}}_{0|t} + \rho_g\boldsymbol{m}_{\mathcal{S}_k,t}, \quad \widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k\cup\{i\}} = \widehat{\boldsymbol{x}}_{0|t} + \rho_g\boldsymbol{m}_{\mathcal{S}_k\cup\{i\},t}. \tag{B.18}$$

By (B.12), $\widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k\cup\{i\}} = \widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k} + \rho_g\boldsymbol{u}_i\nu_i$. Since $\boldsymbol{x}_0 - \widetilde{\boldsymbol{x}}_{0|t,\mathcal{S}_k} = \boldsymbol{\delta}_t - \rho_g\boldsymbol{m}_{\mathcal{S}_k,t}$, expanding the weighted squared-error difference gives

$$\begin{aligned}\Delta_i(\rho_g) &= \|\boldsymbol{\delta}_t - \rho_g\boldsymbol{m}_{\mathcal{S}_k,t}\|_{\boldsymbol{W}_k}^2 - \|\boldsymbol{\delta}_t - \rho_g\boldsymbol{m}_{\mathcal{S}_k,t} - \rho_g\boldsymbol{u}_i\nu_i\|_{\boldsymbol{W}_k}^2 \\ &= 2\rho_g\nu_i\boldsymbol{u}_i^\top\boldsymbol{W}_k\left(\boldsymbol{\delta}_t - \rho_g\boldsymbol{m}_{\mathcal{S}_k,t}\right) - \rho_g^2\nu_i^2\|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2.\end{aligned} \tag{B.19}$$

The scaled estimate is generally not the surrogate posterior mean when $\rho_g \neq 1$. To account for this distinction, decompose its residual error as

$$\boldsymbol{\delta}_t - \rho_g\boldsymbol{m}_{\mathcal{S}_k,t} = \left(\boldsymbol{\delta}_t - \boldsymbol{m}_{\mathcal{S}_k,t}\right) + (1-\rho_g)\boldsymbol{m}_{\mathcal{S}_k,t}. \tag{B.20}$$

Using the zero innovation mean in (B.10) and the cross-moment in (B.11),

$$\begin{aligned}&\mathbb{E}_{q_\Phi}\left[\left(\boldsymbol{\delta}_t - \rho_g\boldsymbol{m}_{\mathcal{S}_k,t}\right)\nu_i \mid \mathcal{F}_k\right] \\ &\quad= \mathbb{E}_{q_\Phi}\left[\left(\boldsymbol{\delta}_t - \boldsymbol{m}_{\mathcal{S}_k,t}\right)\nu_i \mid \mathcal{F}_k\right] + (1-\rho_g)\boldsymbol{m}_{\mathcal{S}_k,t}\mathbb{E}_{q_\Phi}\left[\nu_i \mid \mathcal{F}_k\right] \\ &\quad= \boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i.\end{aligned} \tag{B.21}$$

Taking the conditional expectation of (B.19) therefore yields

$$\begin{aligned}\mathbb{E}_{q_\Phi}\left[\Delta_i(\rho_g) \mid \mathcal{F}_k\right] &= 2\rho_g\boldsymbol{u}_i^\top\boldsymbol{W}_k\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i - \rho_g^2\left(\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}\right)\|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2 \\ &= \left(2\rho_g - \rho_g^2\right)\left(\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}\right)\|\boldsymbol{u}_i\|_{\boldsymbol{W}_k}^2 = (2\rho_g - \rho_g^2)U_i(\mathcal{S}_k),\end{aligned} \tag{B.22}$$

where the second equality uses $\boldsymbol{\Sigma}_{\mathcal{S}_k}\boldsymbol{e}_i = \left(\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii}\right)\boldsymbol{u}_i$. For a common $0 < \rho_g < 2$, the factor $2\rho_g - \rho_g^2 = \rho_g(2-\rho_g)$ is positive and independent of the candidate. Hence

$$\arg\max_{i\in\mathcal{C}_k} U_i(\mathcal{S}_k) = \arg\max_{i\in\mathcal{C}_k}\mathbb{E}_{q_\Phi}\left[\Delta_i(\rho_g) \mid \mathcal{F}_k\right], \tag{B.23}$$

with equality understood as equality of maximizing sets when ties occur. This proves Corollary 1.

B.4 A Bounded-Score Consequence

Lemma 1. *(Bounded Departure from the Analytic Anchor). Let $\bar{U}_i$ denote the normalized analytic utility, and let its bounded calibrated score be $\widetilde{U}_i = \bar{U}_i + c_i$ with $|c_i| \le \beta_{\mathrm{acq}}$. Define*

$$i^\star \in \arg\max_i \bar{U}_i, \quad \hat{\imath} \in \arg\max_i \widetilde{U}_i.$$

Then $0 \le \bar{U}_{i^\star} - \bar{U}_{\hat{\imath}} \le 2\beta_{\mathrm{acq}}$. Consequently, if the analytic margin satisfies

$$\bar{U}_{i^\star} - \max_{j \neq i^\star} \bar{U}_j > 2\beta_{\mathrm{acq}}, \tag{B.24}$$

then the analytic maximizer $i^\star$ remains the unique maximizer after calibration.

Proof. By the optimality of $\widehat{i}$,

$$\bar{U}_{\widehat{i}} + c_{\widehat{i}} \ge \bar{U}_{i^\star} + c_{i^\star}.$$

Therefore,

$$\bar{U}_{i^\star} - \bar{U}_{\widehat{i}} \le c_{\widehat{i}} - c_{i^\star} \le |c_{\widehat{i}}| + |c_{i^\star}| \le 2\beta_{\mathrm{acq}}.$$

If the analytic margin exceeds $2\beta_{\mathrm{acq}}$, no other candidate can overcome this gap under an admissible residual correction, proving the second statement. □

Lemma 1 formalizes the role of bounded residual calibration. The learned scorer may refine the ordering of candidates whose analytic utilities are close, but it cannot arbitrarily overturn a sufficiently strong analytic preference. The final acquisition rule therefore remains anchored to the theoretically motivated posterior-risk criterion while retaining limited flexibility to compensate for surrogate mismatch.

B.5 Matrix-Free Graph Solves and Covariance Probing

GPARA does not explicitly form the dense posterior covariance $\boldsymbol{\Sigma}_{\mathcal{S}} = \boldsymbol{A}_{\mathcal{S}}^{-1}$. Instead, posterior refinement and covariance-dependent acquisition statistics are evaluated through sparse linear solves with the posterior precision $\boldsymbol{A}_{\mathcal{S}}$. For independent noise observations, write $\mathbf{A}_{\mathcal{S}} = \mathbf{D}_A - \lambda \mathbf{C}_g$, where

$$[\mathbf{D}_A]_{ii} = \tau + \lambda d_i + \sigma_y^{-2} \mathbf{1}_{\{i \in \mathcal{S}\}}, \quad d_i = \sum_j c_{ij}. \tag{B.25}$$

For a reference prediction $\boldsymbol{\mu}$, define $\mathbf{f}_{\mathcal{S}} = \boldsymbol{H}_{\mathcal{S}}^{\mathrm{T}} \boldsymbol{R}_{\mathcal{S}}^{-1} (\boldsymbol{y}_{\mathcal{S}} - \boldsymbol{H}_{\mathcal{S}} \boldsymbol{\mu})$. The posterior-mean correction $\boldsymbol{m}_{\mathcal{S}} = \boldsymbol{A}_{\mathcal{S}}^{-1} \mathbf{f}_{\mathcal{S}}$ is approximated by fixed-depth Jacobi iterations:

$$\boldsymbol{m}^{(j+1)} = \mathbf{D}_A^{-1} \left(\mathbf{f}_{\mathcal{S}} + \lambda \mathbf{C}_g \boldsymbol{m}^{(j)} \right). \tag{B.26}$$

Because $\tau > 0$ and $c_{ij} \ge 0$ the Jacobi iteration is contractive under the point-sensing model, with

$$q = \max_i \frac{\lambda d_i}{\tau + \lambda d_i + \sigma_y^{-2} \mathbf{1}_{\{i \in \mathcal{S}\}}} < 1. \tag{B.27}$$

Thus, Jacobi converges to the exact mean as its iteration depth increases. However, we use fixed iterations and treat the resulting correction as a finite-depth approximation rather than an exact solve.

For active acquisition, the utility requires only $\Sigma_{k,ii}$ and $(\boldsymbol{\Sigma}_k \boldsymbol{W}_k \boldsymbol{\Sigma}_k)_{ii}$. We estimate these quantities using P independent Rademacher probes as

$$\mathbf{z}^{(p)} \in \{-1, +1\}^M, \quad \mathbb{E}[\mathbf{z}^{(p)} \mathbf{z}^{(p)\mathsf{T}}] = \boldsymbol{I}. \tag{B.28}$$

For each probe, we solve the linear systems by the preconditioned conjugate gradient (PCG) method

$$\boldsymbol{A}_k \mathbf{v}^{(p)} = \mathbf{z}^{(p)}, \quad \boldsymbol{A}_k \mathbf{s}^{(p)} = \boldsymbol{W}_k^{1/2} \mathbf{z}^{(p)} \tag{B.29}$$

using preconditioned conjugate gradients with the diagonal Jacobi preconditioner $\boldsymbol{P}_k = \mathrm{diag}(\boldsymbol{A}_k)$.

With exact solves,

$$\mathbb{E}\left[z_i^{(p)} v_i^{(p)} \right] = \Sigma_{k,ii}, \quad \mathbb{E}\left[\left(s_i^{(p)} \right)^2 \right] = (\boldsymbol{\Sigma}_k \boldsymbol{W}_k \boldsymbol{\Sigma}_k)_{ii}. \tag{B.30}$$

Accordingly, we evaluate the numerical acquisition utility as

$$\widehat{U}_i = \frac{\widehat{n}_i}{\sigma_y^2 + \max(\widehat{d}_i, [\mathbf{D}_A]_{ii}^{-1}, \varepsilon_d)}, \quad \widehat{d}_i = \frac{1}{P}\sum_{p=1}^{P} z_i^{(p)} v_i^{(p)}, \quad \widehat{n}_i = \frac{1}{P}\sum_{p=1}^{P}\left(s_i^{(p)}\right)^2. \tag{B.31}$$

The two raw moment estimators are unbiased under exact solves, but their stabilized ratio is generally biased at finite P because of probe variance, correlation between numerator and denominator, the nonlinear ratio, and truncated iterative solves. We therefore treat $\widehat{U}_i$ as a numerical approximation of the analytic utility rather than an unbiased estimator.

For both tasks, we use prior samples $L = 4$, graph precision $\tau = 0.0625$, $\varepsilon_d = 10^{-12}$, 24 Jacobi iterations, and PCG with at most 24 iterations and a relative tolerance of 10^{-5}. RFR uses $P = 16$, $\lambda = 0.8$, and $\sigma_y = 10^{-4}$, whereas DC uses $P = 4$, $\lambda = 1.0$, and $\sigma_y = 2 \times 10^{-3}$.

B.6 Construction of the Observation-Set-Adaptive Risk Matrix

The graph posterior and the risk matrix capture two complementary aspects of active acquisition. The graph posterior $\mathbf{\Sigma}_k$ determines the spatial response induced by a prospective measurement: its entry $\Sigma_{k,ji}$ controls how strongly an innovation observed at candidate i influences variable j. In contrast, $\boldsymbol{W}_k$ emphasizes locations that remain less resolved under the current observation set. Locations exhibiting larger disagreement among the graph-conditioned samples receive larger weights. Consequently, a candidate is valuable not merely when it influences many variables, but when its response reaches regions with relatively high remaining disagreement.

Let $\{\widehat{\boldsymbol{x}}_0^{(l)}\}_{l=1}^{L}$, $\widehat{\boldsymbol{x}}_0^{(l)} \sim p_\theta(\boldsymbol{x}_0 \mid \boldsymbol{c})$ denote the L cached samples produced by the frozen diffusion prior. Such an ensemble is generated once before active acquisition and cached thereafter.

At refresh k, the observations $(\mathcal{S}_k, \boldsymbol{y}_{\mathcal{S}_k})$ are applied to $\widehat{\boldsymbol{x}}_0^{(l)}$ by Eq. (13). Let $\widetilde{\boldsymbol{x}}_{0|\mathcal{S}_k}^{(l)}$ be the resulting graph-conditioned sample after correction. The ensemble mean and empirical variance are

$$\widehat{\boldsymbol{\mu}}_{\mathcal{S}_k} = \frac{1}{L}\sum_{l=1}^{L} \widetilde{\boldsymbol{x}}_{0|\mathcal{S}_k}^{(l)}, \quad \widehat{\boldsymbol{v}}_{\mathcal{S}_k} = \frac{1}{L-1}\sum_{l=1}^{L}\left(\widetilde{\boldsymbol{x}}_{0|\mathcal{S}_k}^{(l)} - \widehat{\boldsymbol{\mu}}_{\mathcal{S}_k}\right)^2. \tag{B.32}$$

The average empirical variance over unobserved admissible locations is given by

$$\overline{v}_{\mathcal{S}_k} = \frac{\sum_{j\in\mathcal{C}_k} \widehat{v}_{\mathcal{S}_k,j}}{\max(|\mathcal{C}_k|, 1)}, \tag{B.33}$$

where $\widehat{v}_{\mathcal{S}_k,j}$ is the value at location j. We then construct $\boldsymbol{W}_k = \mathrm{diag}(w_{k,1}, ..., w_{k,M}) \succeq 0$ with

$$w_{k,i} = \mathbf{1}_{\{i\in\mathcal{C}_k\}} \times \left(\varsigma + \frac{\Delta w}{r}\Pi_{[0,r]}\left(\frac{\widehat{v}_{\mathcal{S}_k,i}}{\max(\overline{v}_{\mathcal{S}_k}, \varepsilon_v)}\right)\right). \tag{B.34}$$

Here, the hyperparameters $\varsigma > 0$, $\Delta w \geq 0$, $r > 0$, and $\varepsilon_v > 0$ specify the baseline weight, maximum additional weight, relative-variance cap, and numerical floor, respectively. In this work, we set $\varsigma = 0.25$, $\Delta w = 0.75$, $r = 4$, $\varepsilon_v = 10^{-8}$ and $L = 4$.

Under the unclipped linear graph update, the measurement-dependent correction contains a common additive shift across all cached samples, which cancels when computing the centered ensemble variance. Therefore, $\boldsymbol{W}_k$ is primarily adaptive to the observation set and graph-posterior response rather than directly to the realized measurement values. The small diffusion ensemble is generated only once and cached. The cached ensemble is reconditioned after each acquisition refresh, allowing the risk emphasis to adapt to the evolving observation set without regenerating diffusion samples.

B.7 Batched Acquisition Between Posterior Refreshes

To reduce posterior-refresh cost, the implementation can use batched acquisition as a computational approximation to the single-point rule in Section 3. At refresh k, we select at most $B = 8$ measurements, i.e., $B_k = \min\{B, \Gamma - |\mathcal{S}_k|\}$ from the current candidate set $\mathcal{C}_k$. The graph posterior, risk matrix, and calibrated scores are computed once and remain fixed while constructing the batch.

To discourage redundant selections, we reuse the weighted probe solutions $\mathbf{s}^{(p)}$ from (B.29). For candidates i and j, we define the response-overlap proxy

$$\widehat{\kappa}_{ij} = \Pi_{[0,1]}\left(\frac{\left(\sum_{p=1}^{P} s_i^{(p)} s_j^{(p)}\right)^2}{\max\left(\max(P\widehat{n}_i, 10^{-12})\max(P\widehat{n}_j, 10^{-12}), 10^{-12}\right)}\right), \tag{B.35}$$

where $\widehat{n}_i$ is defined in (B.31). Hence, no additional linear solves are required beyond those already used for acquisition-utility estimation. Let

$$D_k = \max\left(\max_{i\in\mathcal{C}_k} \widetilde{U}_i(\mathcal{S}_k) - \min_{i\in\mathcal{C}_k} \widetilde{U}_i(\mathcal{S}_k), 1\right). \tag{B.36}$$

Starting with $\mathcal{B}_k^{(0)} = \varnothing$, the batch is constructed greedily as

$$\begin{aligned} i_{k,\ell} &\in \arg\max_{i\in\mathcal{C}_k\setminus\mathcal{B}_k^{(\ell-1)}} \left(\widetilde{U}_i(\mathcal{S}_k) - 0.5D_k \sum_{j\in\mathcal{B}_k^{(\ell-1)}} \widehat{\kappa}_{ij}\right), \\ \mathcal{B}_k^{(\ell)} &= \mathcal{B}_k^{(\ell-1)} \cup \{i_{k,\ell}\}, \quad \ell = 1, ..., B_k. \end{aligned} \tag{B.37}$$

Thus, the first point maximizes the calibrated utility, while subsequent points are penalized according to their cumulative response overlap with already selected batch members. After the batch $\mathcal{B}_k = \mathcal{B}_k^{(B_k)}$ is determined, measurements are acquired and the observation set is updated as

$$\mathcal{S}_{k+1} = \mathcal{S}_k \cup \mathcal{B}_k.$$

The graph posterior, risk matrix, and acquisition scores are then recomputed for the next refresh. No newly acquired value is used during construction of the current batch. Note that the single-point utility remains the underlying acquisition criterion of GPARA, while batching serves only as a computational approximation rather than an exact joint batch-risk optimization.

C Network Architectures and Training

The graph predictor G_ϕ and residual scorer r_ψ are learned function approximators whose architectures are not intrinsic to the GPARA formulation. GPARA requires only that G_ϕ produce valid nonnegative graph conductances from the available context and that r_ψ output candidate-wise bounded score corrections. We therefore specify these interfaces and mappings below, while using standard lightweight architectures for their implementation. The two components are trained sequentially. With the diffusion backbone frozen, G_ϕ is first optimized through differentiable graph conditioning. G_ϕ is then frozen and r_ψ is trained using privileged supervision constructed from training targets. All networks are trained offline and remain fixed during deployment.

C.1 Network interfaces.

The graph predictor uses a lightweight UNet to generate four directional edge-logit maps from $\boldsymbol{c}_g$:

$$G_\phi(\boldsymbol{c}_g) \longrightarrow \mathbf{Z}_\phi = \{z_{i,\xi}\}_{i\in\mathcal{V},\xi\in\mathcal{D}} \in \mathbb{R}^{M\times 4}, \quad \mathcal{D} = \{(0,1),(1,0),(1,1),(1,-1)\}, \tag{C.1}$$

corresponding to right, down, down-right, and down-left canonical edges. Each undirected local edge is predicted once, and its reverse direction is introduced deterministically when assembling the symmetric graph. For $j = i + \xi$, let $m_{i,\xi} \in \{0,1\}$ be the valid-edge mask. The final conductance is

$$c_{ij} = c_{ji} = m_{i,\xi}\frac{2\mathrm{sigmoid}(z_{i,\xi})}{\|\xi\|_2}, \quad \xi \in \mathcal{D}. \tag{C.2}$$

Thus, horizontal and vertical conductances lie in $[0,2]$, whereas diagonal conductances are additionally scaled by $1/\sqrt{2}$. Invalid or out-of-domain edges and self-connections are assigned zero weight. The resulting adjacency matrix $\boldsymbol{C}_g = [c_{ij}]$ is sparse, symmetric, and nonnegative. Four canonical edge maps therefore represent an undirected graph with up to eight neighbors per interior node.

The acquisition scorer independently outputs one residual score per target location,

$$R_\psi(\mathbf{z}_{\mathcal{S}_k}) = [r_\psi(1;\mathbf{z}_{\mathcal{S}_k}), ..., r_\psi(M;\mathbf{z}_{\mathcal{S}_k})]^{\mathrm{T}} \in \mathbb{R}^M. \tag{C.3}$$

The scorer combines a convolutional encoder-decoder, four graph-message-passing blocks, and a budget-conditioning MLP. The input $\mathbf{z}_{\mathcal{S}_k}$ captures the available acquisition state, combining task context, accumulated observation information, graph-posterior statistics, the analytic acquisition feature, and the remaining sensing budget.

C.2 Learning the graph through posterior refinement.

For each training example, we randomly sample a reverse step t, obtain the clean estimate $\widehat{\boldsymbol{x}}_{0|t}$ from the frozen diffusion model, and sample a sparse observation set $\mathcal{S}$. The graph predictor constructs $\boldsymbol{C}_g$ and obtains $\boldsymbol{A}_{\mathcal{S}}$. We then apply the same differentiable Jacobi conditioner used at deployment to obtain the residual correction and refined prediction $\widetilde{\boldsymbol{x}}_{0|t}$ through Eq. (13). Gradients are propagated through both the graph construction and the unrolled conditioner to G_ϕ.

Let $\mathcal{V}_{\mathrm{val}} \subseteq \{1,\dots,M\}$ denote the valid spatial region, and $\mathcal{U} = \mathcal{V}_{\mathrm{val}} \setminus \mathcal{S}$ is the unobserved region. For any subset $\mathcal{R} \subseteq \mathcal{V}_{\mathrm{val}}$, define the masked reconstruction error as

$$E_{\mathcal{R}}(\hat{\boldsymbol{x}}_0, \boldsymbol{x}_0) = \frac{\sum_{i\in\mathcal{R}} |\hat{\boldsymbol{x}}_{0,i} - \boldsymbol{x}_{0,i}|^2}{\max(|\mathcal{R}|, 1)}. \tag{C.4}$$

The graph-training objective is formulated as

$$\mathcal{L}_{\mathrm{graph}} = E_{\mathcal{V}_{\mathrm{val}}}(\widetilde{\boldsymbol{x}}_{0|t}, \boldsymbol{x}_0) + \lambda_u E_{\mathcal{U}}(\widetilde{\boldsymbol{x}}_{0|t}, \boldsymbol{x}_0) + \lambda_s E_{\mathcal{S}}(\widetilde{\boldsymbol{x}}_{0|t}, \boldsymbol{y}_{\mathcal{S}}) + \lambda_h \mathcal{L}_{\mathrm{hinge}} + \lambda_a \mathcal{L}_{\mathrm{aff}}. \tag{C.5}$$

The first three terms supervise full-region reconstruction, propagation to unobserved locations, and observation consistency, respectively. To discourage graph refinement from degrading the frozen prediction, we use

$$\mathcal{L}_{\mathrm{hinge}} = \max\left(E_{\mathcal{U}}\left(\widetilde{\boldsymbol{x}}_{0|t}, \boldsymbol{x}_0\right) - E_{\mathcal{U}}\left(\widehat{\boldsymbol{x}}_{0|t}, \boldsymbol{x}_0\right), 0\right). \tag{C.6}$$

We additionally supervise local residual compatibility. Let $\boldsymbol{\delta}_t = \boldsymbol{x}_0 - \widehat{\boldsymbol{x}}_{0|t}$. For each valid canonical edge $(i,j) \in \mathcal{E}$,

$$a_{ij}^{\mathrm{res}} = \exp\left(-\frac{|\boldsymbol{\delta}_{t,i} - \boldsymbol{\delta}_{t,j}|}{s_{\mathrm{res}}} - \frac{|\boldsymbol{x}_{0,i} - \boldsymbol{x}_{0,j}|}{s_x}\right), \tag{C.7}$$

where $s_{\mathrm{res}} > 0$ and $s_x > 0$ are fixed scales. Let $p_{ij}^{\phi} = \mathrm{sigmoid}(z_{i,\xi})$ for $j = i + \xi$. The affinity loss acts on these pre-rescaling edge probabilities:

$$\mathcal{L}_{\mathrm{aff}} = -\frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} \left[a_{ij}^{\mathrm{res}} \log p_{ij}^{\phi} + (1 - a_{ij}^{\mathrm{res}}) \log(1 - p_{ij}^{\phi})\right]. \tag{C.8}$$

Each undirected edge appears once in $\mathcal{E}$. The affinity loss therefore supervises the pre-rescaling edge compatibility, while the reconstruction losses optimize the final conductances end-to-end through the differentiable graph-Bayesian conditioner. Importantly, these conductances couple prediction-residual corrections rather than directly enforcing similarity between reconstructed target values.

In this work, we set $\lambda_u = 2$, $\lambda_s = 0.1$, $\lambda_h = 0.1$, $\lambda_a = 0.01$, $s_{\mathrm{res}} = 0.05$ and $s_x = 0.08$.

C.3 Learning the Acquisition Residual Scorer

For each training example, an observation set is sampled by $\mathcal{S}_k \sim \mathrm{Uniform}\left(\{\mathcal{A} \subseteq \mathcal{V}_{\mathrm{val}} : |\mathcal{A}| < \Gamma\}\right)$. The graph conditioner assimilates $(\mathcal{S}_k, \mathbf{y}_{\mathcal{S}_k})$, after which we compute the analytic utilities U_i for every $i \in \mathcal{C}_k$ and assemble $\mathbf{z}_{\mathcal{S}_k}$. Consequently, revisiting the same training sample can produce different examples through different observation sets.

Let $\widehat{\boldsymbol{\mu}}_{\mathcal{S}_k}$ denote the current graph-conditioned estimate and the error $\boldsymbol{e} = \boldsymbol{x}_0 - \widehat{\boldsymbol{\mu}}_{\mathcal{S}_k}$. Using the graph response $\boldsymbol{u}_i = \boldsymbol{\Sigma}_{\mathcal{S}_k} \boldsymbol{e}_i / (\sigma_y^2 + \Sigma_{\mathcal{S}_k,ii})$, the training-only linear-response teacher evaluates

$$\widehat{\Delta}_i = 2\rho_g \nu_i^{\mathrm{tr}} \boldsymbol{u}_i^{\top} \mathbf{W}_{\mathrm{tr}} \boldsymbol{e} - \rho_g^2 (\nu_i^{\mathrm{tr}})^2 \boldsymbol{u}_i^{\top} \mathbf{W}_{\mathrm{tr}} \boldsymbol{u}_i. \tag{C.9}$$

Here, $\mathbf{W}_{\mathrm{tr}} = \mathrm{diag}(\mathbf{1}_{\{1\in\mathcal{V}_{\mathrm{val}}\}}, ..., \mathbf{1}_{\{M\in\mathcal{V}_{\mathrm{val}}\}})$, and $\nu_i^{\mathrm{tr}} = \boldsymbol{x}_{0,i} - (\overline{\boldsymbol{x}}_0 + \overline{\boldsymbol{m}}_{\mathcal{S}_k})_i$, where $\overline{\boldsymbol{x}}_0$ is the original cached prior mean and $\overline{\boldsymbol{m}}_{\mathcal{S}_k}$ is the unscaled correction. For any candidate quantity $\boldsymbol{x}_i$, define

$$\mathcal{Z}_{\mathcal{C}_k}(\boldsymbol{x}_i) = \frac{\boldsymbol{x}_i - \boldsymbol{\mu}_{\mathcal{C}_k}}{\max\{\boldsymbol{\sigma}_{\mathcal{C}_k}, 10^{-6}\}}, \tag{C.10}$$

where $\boldsymbol{\mu}_{\mathcal{C}_k}$ and $\boldsymbol{\sigma}_{\mathcal{C}_k}$ are the mean and standard deviation over $i \in \mathcal{C}_k$. The same normalization of the analytic anchor is used during training and deployment. We have the normalized analytic utility:

$$\bar{U}_i = \mathcal{Z}_{\mathcal{C}_k}\left(\log\max\{U_i, 10^{-12}\}\right), \quad b_i = \mathcal{Z}_{\mathcal{C}_k}\left(\widehat{\Delta}_i\right). \tag{C.11}$$

The calibrated score and training teacher are given by $\widetilde{U}_i = \bar{U}_i + \beta_{\mathrm{acq}} \tanh\left(r_\psi(i; \mathbf{z}_{\mathcal{S}_k})\right)$ and $h_i = \lambda_m \bar{U}_i + (1-\lambda_m) b_i$. The corresponding candidate distributions are

$$p_i^\tau = \frac{\exp(h_i/T_\tau)}{\sum_{j\in\mathcal{C}_k}\exp(h_j/T_\tau)}, \quad p_i^\psi = \frac{\exp(\widetilde{U}_i/T_\psi)}{\sum_{j\in\mathcal{C}_k}\exp(\widetilde{U}_j/T_\psi)}, \quad p_i^{\mathrm{a}} = \frac{\exp(\bar{U}_i/T_\psi)}{\sum_{j\in\mathcal{C}_k}\exp(\bar{U}_j/T_\psi)}, \tag{C.12}$$

where $T_\tau, T_\psi > 0$ are temperatures. These distributions are used for training, and deployment selects locations from scores. Let $i^\star = \arg\max_i h_i$, $\widehat{i} = \arg\max_i \widetilde{U}_i$, $h_{\max} = \max_i h_i$, and $h_{\min} = \min_i h_i$. The distillation and selection losses are

$$\mathcal{L}_{\mathrm{dist}} = -\sum_i p_i^\tau \log p_i^\psi, \ \mathcal{L}_{\mathrm{top}} = -\log\frac{\exp(\widetilde{U}_{i^\star})}{\sum_j \exp(\widetilde{U}_j)}, \ \mathcal{L}_{\mathrm{reg}} = \frac{\sum_i p_i^\psi (h_{\max} - h_i)}{\max(h_{\max} - h_{\min}, 10^{-8})}. \tag{C.13}$$

They respectively transfer the teacher preference distribution, promote its highest-scoring candidate, and penalize expected teacher regret. We further supervise the learned departure from the normalized analytic anchor using

$$\mathcal{L}_{\mathrm{cal}} = \frac{1}{2}\sum_i (p_i^\tau + p_i^{\mathrm{a}})\mathcal{H}\left(\widetilde{U}_i - \bar{U}_i - \Pi_{[-1,1]}(h_i - \bar{U}_i)\right), \tag{C.14}$$

where $\mathcal{H}(d) = d^2/2$ for $|d| < 1$, otherwise $|d| - 1/2$. We additionally introduce

$$\mathcal{L}_{\mathrm{mar}} = \mathbf{1}_{\{\widehat{i}\neq i^\star\}}\max\left(0.5 + \widetilde{U}_{\widehat{i}} - \widetilde{U}_{i^\star}, 0\right), \quad \mathcal{L}_{\mathrm{anc}} = \max\left(\sum_i p_i^{\mathrm{a}} h_i - \sum_i p_i^\psi h_i, 0\right). \tag{C.15}$$

The margin term promotes separation from an incorrect greedy winner, while the anchor-relative term prevents the learned policy from having lower teacher-expected value than the analytic-anchor distribution. All sums and extrema above are over $\mathcal{C}_k$.

The complete objective is $\mathcal{L}_{\mathrm{scorer}} = \sum_{j\in\mathcal{J}} \lambda_j \mathcal{L}_j$, where $\mathcal{J} = (\mathrm{dist}, \mathrm{top}, \mathrm{reg}, \mathrm{cal}, \mathrm{mar}, \mathrm{anc})$ specifies the ordered set of loss terms. Only the scorer ψ is optimized; prior generation, graph conditioning, analytic scoring, and teacher construction are detached. Following the order in $\mathcal{J}$, we use $\lambda_{\mathcal{J}} = (1, 0.1, 2, 0.1, 0.1, 0)$, $(T_\tau, T_\psi) = (0.15, 0.30)$ and $\lambda_m = 0$ for RFR, and $\lambda_{\mathcal{J}} = (1, 0.1, 0.5, 0.25, 0, 0.05)$, $(T_\tau, T_\psi) = (0.25, 0.25)$ and $\lambda_m = 0.2$ for DC. We set the refinement strength to $\rho_g = 1.4$ for RFR and $\rho_g = 1.0$ for DC, with a base residual-score scale of $\beta_{\mathrm{acq}} = 2.0$ for both tasks.