OPTIMIZING SAMPLING PATTERN CORRELATIONS FOR NEURAL FIELD TRAINING

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural fields are optimized from samples of a source signal whose locations can often be chosen freely. While existing approaches exploit this freedom by adapting individual sample locations to the signal, we investigate a complementary question: how should the samples be spatially correlated? We characterize these correlations through sampling-pattern power spectra and explore this space to identify patterns that optimize the outcome of neural field fitting. By evaluating sampling patterns through the resulting neural field, our approach accounts for the complete fitting process rather than integration accuracy alone. The optimized patterns outperform IID uniform sampling as well as correlated patterns designed for numerical integration, with particularly pronounced benefits under limited sampling budgets. Moreover, the patterns transfer across source signals, enabling inexpensive proxy signals to drive their optimization for more costly settings. We analyze the resulting sampling structures and characterize how they depend on the neural field fitting problem. Together, our results establish optimized sampling correlations as an effective tool for improving neural field fitting.

1 INTRODUCTION

Neural fields, also referred to as implicit neural representations or coordinate-based neural networks, are instance-specific neural representations of signals (Xie et al., 2022). They are particularly well suited for representing continuous signals, as their representational capacity is not tied to a predetermined spatial discretization. To obtain such a representation for a given source signal, a neural field is optimized by fitting it to samples queried from the signal at arbitrary continuous locations. Unlike conventional supervised learning with a fixed training set, *this makes the sampling locations themselves a degree of freedom*. Exploiting this freedom presents an opportunity for improving approximation quality under a limited sampling budget, particularly when querying the source signal is costly, e.g., when it requires expensive simulation or rendering.

Prior work has explored adaptive sampling strategies that analyze the signal and the current neural approximation to progressively refine sample locations during optimization (Zhang et al., 2024; Kheradmand et al., 2024; Zhang et al., 2026; Gao et al., 2023; Yifan et al., 2021). Complementing these signal-adaptive approaches, we consider a largely unexplored direction: *optimizing the spatial correlations that structure the sampling pattern* (Fig. 1). Concretely, we characterize these correlations through the sampling-pattern power spectrum (Ulichney, 1988; Lagae & Dutré, 2008) and explore this space to identify patterns that optimize approximation quality for a given neural field fitting problem and generalize across source signals.

In doing so, our approach establishes a connection between two largely distinct lines of research. The first is hyperparameter optimization, where an outer-loop optimizer searches over model and training configurations by repeatedly training and evaluating models in an inner loop (Bergstra et al., 2011; Snoek et al., 2012). We adopt this paradigm for instance-specific neural field fitting, but turn the outer-loop search toward a different object: the correlations that structure the training samples themselves. The second line of research concerns sampling pattern design for numerical integration, where spatial correlations between samples are carefully constructed to improve integration accuracy (Niederreiter, 1978; Yan et al., 2015; Singh et al., 2019). While these methods optimize sampling patterns for estimating a given integral, they do not account for the optimization dynamics induced by the samples. In neural field fitting, however, samples not only approximate the loss integral but also drive the

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.

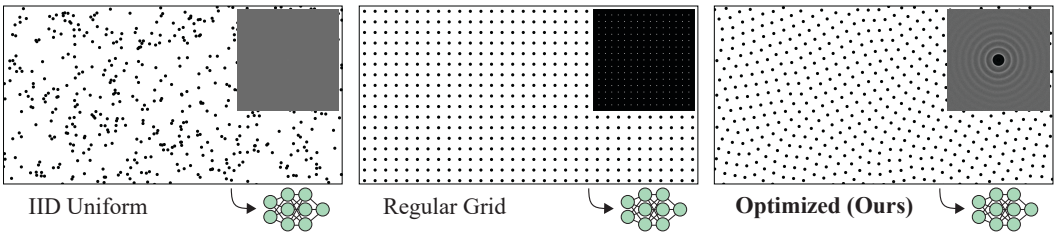

Figure 1: **Spatial correlations in sampling patterns.** IID uniform sampling (left) places samples independently, whereas a regular grid (center) exhibits strong spatial structure. These differences are reflected in their power spectra (insets). When used to fit a neural field, different sampling patterns induce different optimization dynamics and ultimately affect approximation quality. We use the power spectrum to parameterize the space of sampling patterns and identify correlation structures that optimize the outcome of neural field fitting and generalize across source signals (right).

optimization trajectory. Our approach therefore optimizes sampling correlations for the quality of the resulting neural field, with the full neural-field optimization process in the loop.

We show that our optimized correlated sampling patterns outperform independent and identically distributed (IID) uniform sampling, as well as correlated patterns specifically designed for numerical integration. The benefits are particularly pronounced under limited sampling budgets, e.g., when querying the source signal is computationally expensive. Importantly, we find that optimized sampling patterns transfer well across source signals, allowing inexpensive proxy signals to drive the pattern optimization for more costly settings. In summary, our contributions are:

- We introduce a framework for optimizing sampling correlations for neural field fitting, with the complete training process in the loop.
- We show improvements over uncorrelated and integration-optimized sampling, particularly under limited sampling budgets, and demonstrate transfer across source signals.
- We analyze the resulting sampling patterns and characterize how their structure depends on the neural field fitting problem.

All sampling patterns, data, and source code will be released upon publication.

2 Related Work

Correlations in Point Patterns Spatially correlated point patterns arise in a wide range of applications, including numerical integration, object placement, and artistic visual effects. Classical approaches to numerical integration exploit structure among sample locations. Quasi-Monte Carlo methods replace independent random samples with deterministic low-discrepancy sequences (Niederreiter, 1978), with prominent constructions including Sobol (1967) sequences. Related stratification strategies include N-rooks sampling (Shirley, 1991). Another prominent family is blue-noise sampling (Ulichney, 1988; Yan et al., 2015), for which numerous methods construct well-spaced point sets with favorable spatial and spectral properties (Lagae & Dutré, 2008; De Goes et al., 2012; Ahmed et al., 2022). We likewise exploit correlations between sample locations, but rather than adopting a predefined sampling structure, we optimize these correlations specifically for neural field fitting.

Analyses of point-pattern structure for numerical integration relate Monte Carlo variance to correlations between sample locations and, in the Fourier domain, to the interaction between the power spectra of the sampling pattern and integrand (Pilleboue et al., 2015; Singh et al., 2019). Related spectral formulations characterize point patterns through the distribution of pairwise sample differences (Wei & Wang, 2011). This connection has motivated the deliberate design of sampling correlations and spectra for desirable integration and aliasing properties (Heck et al., 2013; Kailkhura et al., 2016). More generally, point patterns can be directly optimized toward target spectral, integration, or uniformity objectives, either by optimizing sample locations or within structured sampling constructions (Zhou et al., 2012; Doignies et al., 2024). In contrast, we optimize point-pattern correlations with respect to the outcome of neural field optimization, thereby accounting for optimization dynamics rather than integration quality alone.

Beyond individual sampling constructions, more general approaches have been developed to synthesize point patterns from prescribed statistical properties (Leimkühler et al., 2019) or from example point sets (Doignies et al., 2023). Related to our approach, point correlations have also been represented in a compact parameter space from which correlation structures can subsequently be realized as point patterns (Huang et al., 2023). We similarly parameterize a space of correlations, but optimize within it according to neural field fitting performance.

Sampling for Model Optimization The choice of training samples can substantially affect model optimization. In the conventional supervised setting, importance sampling prioritizes informative examples from a given training set to reduce stochastic-gradient variance (Zhao & Zhang, 2015; Katharopoulos & Fleuret, 2018). Sampling distributions can also be adapted online to the evolving optimization state (Bouchard et al., 2015; Salaün et al., 2025). These approaches optimize the selection or weighting of existing training examples; in our setting, the continuous source signal instead allows us to choose the locations at which training samples are acquired.

This additional freedom has been exploited in neural field optimization in several ways. Existing approaches adapt sample locations based on properties of the source signal or current neural approximation, such as reconstruction error, frequency information, or sample influence on optimization (Yifan et al., 2021; Zhang et al., 2024; Kheradmand et al., 2024; Zhang et al., 2025; 2026). In a non-adaptive setting, the effect of point-set discrepancy on neural field training has been studied, demonstrating benefits of predefined farthest-point sampling structures (Xu et al., 2020). Our approach instead optimizes the correlations of non-adaptive sampling patterns specifically for the fitting problem, while retaining their ability to generalize across source signals.

More closely related to our setting, the spectral properties of training samples have been connected to neural-network generalization, demonstrating benefits of correlated patterns such as blue noise over random sampling (Kailkhura et al., 2020). This work transfers spectral tools for numerical integration to learning objectives, but does not account for the optimization dynamics induced by the samples. Sampling has also been optimized directly through downstream reconstruction models (Bahadir et al., 2019; Huijben et al., 2020). These methods jointly optimize acquisition patterns and reconstruction, whereas we consider continuous sampling locations for instance-specific neural field fitting.

Most closely related to our approach, sample locations have been meta-learned by differentiating through model fitting, directly optimizing their utility for previously unseen signals (Liu et al., 2023). This shares our goal of identifying transferable sampling patterns by evaluating them through the complete fitting process. However, explicitly meta-optimizing individual sample locations requires differentiating through the inner optimization and additionally requires derivatives of the supervision signal with respect to the sample coordinates, which can introduce substantial computational and implementation overhead for black-box source signals. Moreover, the dimensionality of the outer optimization grows with the number of samples. In contrast, our formulation relies solely on evaluations of the source signal and optimizes a compact representation of sampling correlations whose dimensionality is independent of the number of realized samples, allowing the same framework to operate across a wide range of sampling budgets.

3 Method

3.1 Problem Setting

We consider general signals $f : \Omega \to \mathbb{R}^{d_o}$, where $\Omega \subset \mathbb{R}^d$ is a compact domain. The variable $\mathbf{x} \in \Omega$ denotes an input coordinate, and $f(\mathbf{x}) \in \mathbb{R}^{d_o}$ the signal value. We assume that f can be queried at arbitrary coordinates $\mathbf{x} \in \Omega$, potentially at non-negligible acquisition cost, but is otherwise unknown. We seek to fit a neural field $f_\theta : \Omega \to \mathbb{R}^{d_o}$, parameterized by weights θ, to f. We measure the approximation quality of the neural field by its average loss over the continuous signal domain,

$$\mathcal{L}(\theta) = \frac{1}{|\Omega|} \int_\Omega \ell\left(f_\theta(\mathbf{x}), f(\mathbf{x})\right) \, \mathrm{d}\mathbf{x}, \tag{1}$$

where $\ell : \mathbb{R}^{d_o} \times \mathbb{R}^{d_o} \to \mathbb{R}_{\geq 0}$ is a pointwise error measure. Since f is only available through queries, the continuous loss $\mathcal{L}$ cannot generally be evaluated during fitting. We therefore consider a sampling budget of N queries and choose a corresponding sampling pattern $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N \subset \Omega$. Querying f at

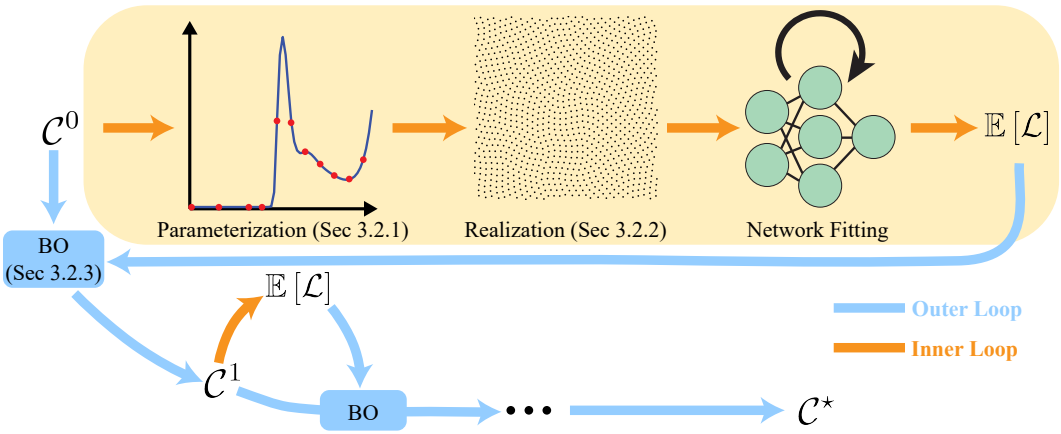

Figure 2: **Method overview**. Our method alternates between inner-loop evaluation and outer-loop optimization. In each inner loop, we parameterize the sampling correlations (Sec. 3.2.1), realize patterns with the corresponding radial spectra (Sec. 3.2.2), and use the resulting samples to train a neural field on the target signals, yielding the expected validation losses $\mathbb{E}[\mathcal{L}]$ in Eq. 6. In the outer loop, Bayesian optimization (Sec. 3.2.3) proposes the new correlation parameters. These two loops are repeated for a finite BO budget, then return the best feasible correlation parameters $\mathcal{C}^\star$ found during the search.

these coordinates yields the finite-sample training objective

$$\mathcal{L}_{\mathcal{X}}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \ell\left(f_\theta(\mathbf{x}_i), f(\mathbf{x}_i)\right). \tag{2}$$

Neural field fitting proceeds by minimizing $\mathcal{L}_{\mathcal{X}}$ using gradient-based optimization. Consequently, the chosen sampling pattern influences not only how well the continuous objective is approximated, but also the optimization trajectory itself. Our goal is to choose $\mathcal{X}$ such that the resulting neural field achieves low approximation error $\mathcal{L}$ over the continuous domain. Rather than adapting individual sample locations to a particular f, we seek correlations among the sample coordinates that yield high-quality solutions across source signals and stochastic realizations of neural field optimization.

3.2 Proposed Framework

Our approach searches for sampling correlations that maximize neural field fitting quality, and we visualize the method framework in Fig. 2. To enable this search, we first introduce two key ingredients. First, we parameterize the correlation structure of sampling patterns through a compact representation of their power spectrum (Sec. 3.2.1). Second, we realize point patterns with a given target correlation structure (Sec. 3.2.2), providing a mapping from correlation parameters to concrete sampling patterns. We then combine these ingredients in an outer-loop constrained Bayesian optimization (Sec. 3.2.3), where the candidate correlation parameters are optimized to improve the reconstruction quality of neural fields fitting. Additional implementation details for each stage are provided in Appendix A.

3.2.1 Parameterizing Sampling Correlations

Our goal is to characterize the spatial structure of $\mathcal{X}$ without considering its Nd coordinate degrees of freedom individually. Independent uniform samples exhibit no preferred spatial relationships, whereas structured patterns introduce characteristic correlations between sample locations, such as repulsion or regular spacing. We characterize these correlations through the power spectrum of the point pattern (Ulichney, 1988; Wei & Wang, 2011; Singh et al., 2019) (Fig. 3a, b). Treating $\mathcal{X}$ as a sum of Dirac impulses on the unit torus, its power spectrum is

$$P_{\mathcal{X}}(\mathbf{k}) = \frac{1}{N}\left|\sum_{i=1}^{N} \exp\left(-2\pi \mathrm{i}\, \mathbf{k}^{\mathsf{T}}\mathbf{x}_i\right)\right|^2, \tag{3}$$

where $\mathbf{k} \in \mathbb{Z}^d$ indexes the discrete spatial frequencies. Different spatial correlations give rise to different characteristic power spectra, as illustrated in the insets of Fig. 1. Representing the full

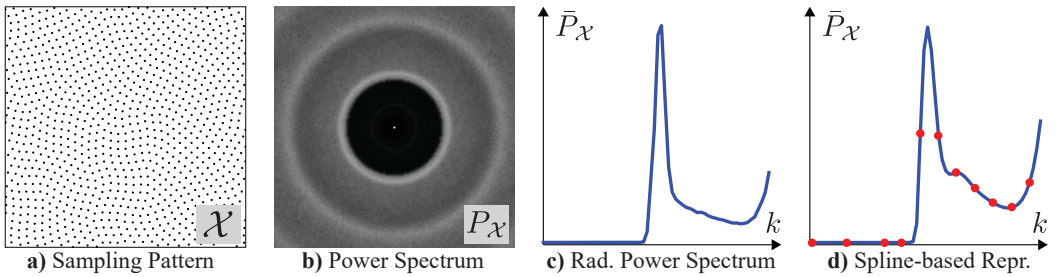

a) Sampling Pattern **b)** Power Spectrum **c)** Rad. Power Spectrum **d)** Spline-based Repr.

Figure 3: **Compact parameterization of point correlations.** From a sampling pattern *(a)*, we compute its power spectrum *(b)* and radially average it to obtain a one-dimensional spectral profile *(c)*. We compactly parameterize such profiles using a B-spline *(d)*, whose control values (heights of the red dots) define a low-dimensional search space over radial spectral statistics.

high-dimensional power spectrum requires a large number of degrees of freedom; therefore, we use its radial average power spectrum over frequency direction as a compact representation (Fig. 3c)

$$\bar{P}_{\mathcal{X}}(k) = \mathbb{E}_{\mathbf{k}\,:\,\|\mathbf{k}\|\approx k}\left[P_{\mathcal{X}}(\mathbf{k})\right], \tag{4}$$

where $k \in \mathbb{Z}_{\geq 0}$ denotes radial frequency and $\|\mathbf{k}\| \approx k$ denotes the discrete frequency vectors falling into the radial bin centered at k. This projection retains only radial spectral statistics of the power spectrum while discarding the angular information. For isotropically realizable 1D profiles, the realization procedure can produce patterns with isotropic full spectra. For profiles that are isotropically non-realizable, the same radial statistics can still be matched, but doing so introduces directional variation in the full spectra, resulting in anisotropic realizations.

Despite radial averaging, representing the spectrum at every discrete frequency k still results in a high-dimensional representation. At the same time, relevant radial spectra exhibit regularity across frequency, suggesting that they can be represented with substantially fewer degrees of freedom. We therefore parameterize the radial spectrum more compactly using a cubic B-spline through M control points $\{(k_j, c_j)\}_{j=1}^{M}$ (Fig. 3d). The radial locations $\{k_j\}_{j=1}^{M}$ are fixed at predefined locations, while the spectral values $\mathcal{C} = \{c_j\}_{j=1}^{M}$ constitute the free parameters of our representation. Details on the placement and range of the radial locations are provided in Appendix A. Since direct cubic B-spline interpolation can produce negative values between control points, we perform the interpolation in inverse-softplus space and apply softplus to the resulting curve, ensuring a nonnegative spectrum. The resulting parameterization $\mathcal{C}$ compactly describes the isotropic spatial correlations in $\mathcal{X}$ using only $M \ll Nd$ parameters, yielding a low-dimensional space that can be efficiently explored.

3.2.2 Realizing Sampling Patterns

The previous subsection provides a compact mapping from a sampling pattern $\mathcal{X}$ to its correlation parameters $\mathcal{C}$. While the outer optimization operates in this compact parameter space, neural field fitting ultimately requires a concrete set of sample coordinates. We therefore also require the inverse mapping: given a set of spectral parameters $\mathcal{C}$, we seek a sampling pattern $\mathcal{X}$ whose correlations realize the corresponding target spectrum. We obtain the target radial spectrum $\bar{P}_{\mathcal{C}}(k)$ by evaluating the spline defined by $\mathcal{C}$ at each radial frequency k. Since the mapping from a sampling pattern to its spectral representation is fully differentiable, we realize a target spectrum by directly optimizing all Nd sample coordinates using gradient-based optimization. Despite its dimensionality, this realization is comparatively inexpensive, as it operates solely on the point pattern and does not involve neural field fitting.

We obtain the realized sampling pattern by minimizing the realization objective, encouraging the radial power spectrum of $\mathcal{X}$ to match the target:

$$\mathcal{X}^{\star}(\mathcal{C}, \zeta) = \arg\min_{\mathcal{X}} \mathbb{E}_k \|\bar{P}_{\mathcal{X}}(k) - \bar{P}_{\mathcal{C}}(k)\|. \tag{5}$$

Here, ζ denotes the random initialization of the sample coordinates; in practice, we initialize them uniformly and independently over the domain. Different initializations therefore yield different point patterns while matching the same target radial spectrum.

The radial representation, however, does not uniquely determine the full-dimensional spectrum of a sampling pattern. Radial averaging discards directional information, such that many different

full-dimensional spectra can yield the same target $\bar{P}_{\mathcal{C}}$. Moreover, we find that some target radial spectra cannot be accurately matched by isotropic point-pattern realizations. In such cases, Eq. 5 may instead be minimized by an anisotropic spectrum whose radial average nevertheless matches the target. Despite these ambiguities, we observe that gradient-based realization exhibits a strong implicit bias. For a fixed $\mathcal{C}$, different random initializations tend to produce highly similar full-dimensional spectra, even when the resulting spectra are anisotropic. Consequently, while $\mathcal{C}$ explicitly describes only radial correlations, together with our realization procedure it consistently enumerates corresponding full-dimensional spectral structures, including anisotropic ones. We further analyze this behavior in Appendix G.

Overall, the realization procedure provides concrete sampling patterns $\mathcal{X}^\star(\mathcal{C}, \zeta)$ for candidate correlation parameters $\mathcal{C}$, enabling evaluation of their neural field fitting performance in the next section.

3.2.3 Optimizing Sampling Correlations

With the parameterization and realization procedure in place, we can optimize sampling correlations according to their neural field fitting performance. Our approach consists of an inner neural field optimization nested within an outer optimization of the correlation parameters $\mathcal{C}$ (Fig. 2).

For each inner-loop evaluation, a candidate $\mathcal{C}$ is first realized as a sampling pattern $\mathcal{X}^\star(\mathcal{C}, \zeta)$ following Sec. 3.2.2. We denote neural field fitting by $\mathcal{A}(f, \mathcal{X}, \xi)$, which returns the fitted neural field parameters for source signal f and sampling pattern $\mathcal{X}$. Here, ξ captures stochasticity in the fitting process. We evaluate the neural field using the continuous-domain loss $\mathcal{L}$ from Eq. 1, approximated using a large number of Monte Carlo samples drawn independently of the training samples. Before evaluating a candidate through neural field fitting, we additionally need to ensure that its target spectrum can be realized. Not every radial spectrum described by our spline parameterization necessarily corresponds to a realizable point pattern; in particular, nonnegativity alone is not sufficient for a function to constitute the power spectrum of a finite set of points. Since we do not have a tractable criterion for determining realizability directly from a spectral profile, we instead assess it through the realization procedure itself. Specifically, we consider a candidate $\mathcal{C}$ feasible if its average spectral realization error does not exceed a fixed threshold $\tau \in \mathbb{R}_{\geq 0}$.

The outer loop searches over these feasible correlation parameters $\mathcal{C}$ to minimize the expected fitting loss across source signals, sampling-pattern realizations, and stochastic training realizations. We formulate this as the constrained optimization problem

$$\begin{aligned} \mathcal{C}^\star = \arg\min_{\mathcal{C}} \mathbb{E}_{f,\zeta,\xi} \left[\mathcal{L} \left(\mathcal{A} \left(f, \mathcal{X}^\star(\mathcal{C}, \zeta), \xi \right) \right) \right] \\ \text{s.t.} \quad \mathbb{E}_{\zeta} \left[\| \bar{P}_{\mathcal{X}^\star(\mathcal{C},\zeta)}(k) - \bar{P}_{\mathcal{C}}(k) \| \right] \leq \tau. \end{aligned} \tag{6}$$

Candidates that violate this constraint are marked as infeasible and excluded from further consideration. Each evaluation of the outer objective requires multiple neural field fitting runs, motivating its treatment as an expensive black-box objective. We therefore solve it using constrained Bayesian Optimization (BO) (Snoek et al., 2012; Gelbart et al., 2014). Specifically, we use the Ax framework (Olson et al., 2025) with Log Expected Improvement (Ament et al., 2023) as the acquisition function. In practice, for each candidate $\mathcal{C}$, we generate K independent sampling-pattern realizations, whose realization losses estimate the feasibility constraint, and evaluate each on L source signals, resulting in KL neural field fitting runs whose losses are averaged to estimate the outer objective. The resulting sampling-pattern correlation parameters $\mathcal{C}^\star$ are optimized for neural field fitting performance and can be realized as concrete sampling patterns using the procedure from Sec. 3.2.2.

4 Results and Evaluation

In this section, we first describe our experimental protocol (Sec. 4.1). We then evaluate our framework on different 2D and 3D signals (Sec. 4.2).

4.1 Experimental Protocol

Target Signals We evaluate our framework on three 2D tasks and three 3D tasks. For computational tractability, we run BO using only two source signals, each with three random seeds. For 2D tasks,

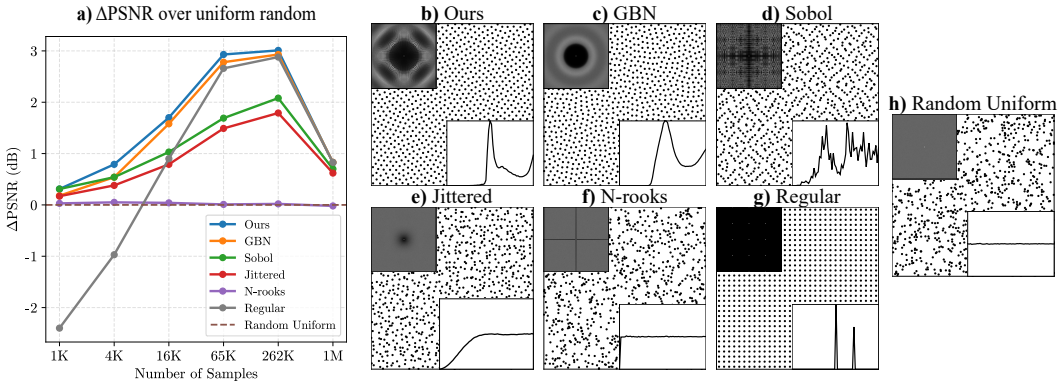

Figure 4: **LDR image fitting.** *(a)* Mean ΔPSNR relative to random uniform sampling across budgets; higher is better. *(b–h)* Representative point patterns with their full and radial power spectra.

we consider low-dynamic-range (LDR) and high-dynamic-range (HDR) image fitting. For LDR images, we use DIV2K (Agustsson & Timofte, 2017) and evaluate on fourteen held-out images. For HDR images, we use Poly Haven environment maps[1] and evaluate on fifteen held-out maps. Continuous images are constructed using bicubic interpolation. For 3D, we consider occupancy and SDF fitting, evaluating the optimized patterns on fourteen and five held-out meshes, respectively. To concentrate SDF samples near the surface, we partition the domain into a grid and place samples only in mesh-intersecting cells, yielding a mesh-dependent total budget (Appendix A). The evaluation meshes are listed in Appendix D.2 and Appendix D.3.

We additionally consider two expensive-query settings: 2D Monte Carlo rendering (MCR) and 3D Monte Carlo geometry processing (MCGP). In both settings, acquiring samples incurs substantial computational cost, making the choice of sample locations particularly consequential under a limited sampling budget. Due to this cost, we do not optimize the sampling patterns in these settings, but instead evaluate patterns optimized on cheaper proxy signals. For MCR, we evaluate on three scenes: *Country Kitchen*, *Grey and White Room*, and *White Room*[2] with each query using 2^{14} stochastic samples. For MCGP, target values are solutions of a Laplace equation over $[-1, 1]^3$, estimated using the Walk-on-Spheres solver (Sawhney & Crane, 2020). We impose a spatially varying function on the geometry surface and a piecewise constant function on the bounding cube; the exact boundary conditions and evaluated geometries are provided in Appendix D.4.

Neural Fields. For 2D signals, we utilize neural fields with Axis-Aligned Positional Encoding (AAPE), while for 3D signals, we use Random Fourier Features (RFF) (Tancik et al., 2020). Results for other neural fields, SIREN (Sitzmann et al., 2020) and FINER (Liu et al., 2024), are provided in Appendix E. We additionally investigate whether the optimized sampling patterns are specific to the chosen architectures and how well they perform across architectures and signals in Appendix B.1.

Baselines We compare our optimized sampling patterns against a range of established sampling strategies with different spatial properties. Random uniform sampling serves as the uncorrelated baseline, while regular grids, stratified (jittered) sampling, Latin-hypercube (N-rooks) sampling (Shirley et al., 1991), and scrambled Sobol sequences (Sobol, 1967; Owen, 1998) introduce different forms of stratification or low-discrepancy structure to improve spatial coverage and numerical integration. We further compare against Gaussian Blue Noise (GBN) (Ahmed et al., 2022), a correlated sampling pattern designed for favorable spectral properties and improved Monte Carlo integration. For larger budgets, the optimized and GBN patterns are obtained by tiling their base patterns of $1{,}024$ points, while the remaining baselines are generated directly at the target acquisition budgets.

Evaluation We evaluate each setting using six random seeds, which determine the neural-field initialization and sampling-pattern realization. All evaluations are performed on the fitted neural fields using 4×10^6 points sampled uniformly at random over the domain, except for SDF fitting, where we reconstruct the geometry using Marching Cubes (Lorensen & Cline, 1987). All metrics are reported as

[1] `https://polyhaven.com/hdris`
[2] `https://mitsuba.readthedocs.io/en/stable/src/gallery.html`

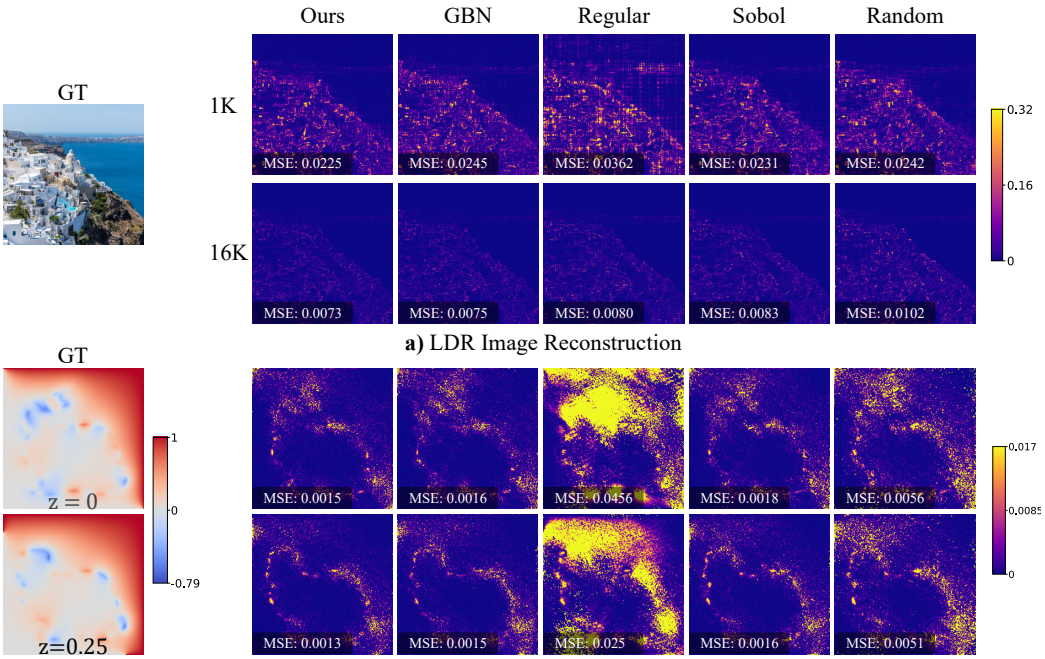

Figure 5: **Qualitative results.** *(a)* LDR reconstruction at two budgets: squared-error maps and MSE for different sampling patterns, evaluated on a $2{,}048 \times 2{,}048$ grid. *(b)* MCGP using a budget of 1M samples: two 256×256 axial slices at $z = 0$ and $z = 0.25$ of the WoS-estimated solution field; colorbar denotes function values. We report squared-error maps and MSE. Lower MSE is better.

improvements over random uniform patterns, denoted as Δ. Specifically, for all 2D signals, we report improvement in Peak Signal-to-Noise Ratio (Δ PSNR); for 3D signals, we report Δ Intersection over Union (IoU) for occupancy fields, Δ Chamfer-L_2 distance for SDFs computed using 10^5 surface samples, and Δ MSE in log domain for MCGP, formulated as $\Delta\text{MSE}_{\text{dB}} = 10 \log_{10} \left(\frac{\text{MSE}_{\text{random}}}{\text{MSE}_{\text{method}}} \right)$.

4.2 Results

LDR Images Fig. 4 shows the reconstruction improvements relative to random uniform across budgets, with different sampling patterns and their corresponding 2D and radial power spectra. Our optimized pattern consistently achieves the highest reconstruction quality across all sampling budgets, with the improvement over Sobol and random uniform increasing up to the budget of 262K samples. At the largest budget, $1M$, the gain decreases substantially, suggesting that the benefit of correlation diminishes as the signals are densely sampled. As shown in Fig. 4 (b), our resulting patterns exhibit angular anisotropy under fixed $\mathcal{C}$; we further study its effect on reconstruction quality in Appendix B.2. Qualitative results are provided in Fig. 5(a) and Appendix F.1.

HDR Images Fig. 6(a) illustrates the performance gain across different sampling budgets. Similar to the LDR reconstruction results, our optimized pattern is competitive with or outperforms all baselines across the evaluated budgets.

Monte Carlo Rendering We perform MCR evaluation, using the patterns optimized on the LDR and HDR reconstruction tasks. As demonstrated in Fig. 6(b), both transferred patterns consistently outperform the baselines at low and intermediate sampling budgets. This demonstrates that the optimized patterns from our framework can generalize across different source signals, allowing inexpensive proxy signals to drive pattern optimization for costly target signals.

Occupancy Fields We then extend our evaluation to 3D signals. Fig. 6(c) illustrates IoU gain over random uniform across budgets. As demonstrated, our sampling pattern outperforms weaker generic baselines and generally achieves higher mean IoU than all the baselines, with the largest improvements observed at smaller sampling budgets.

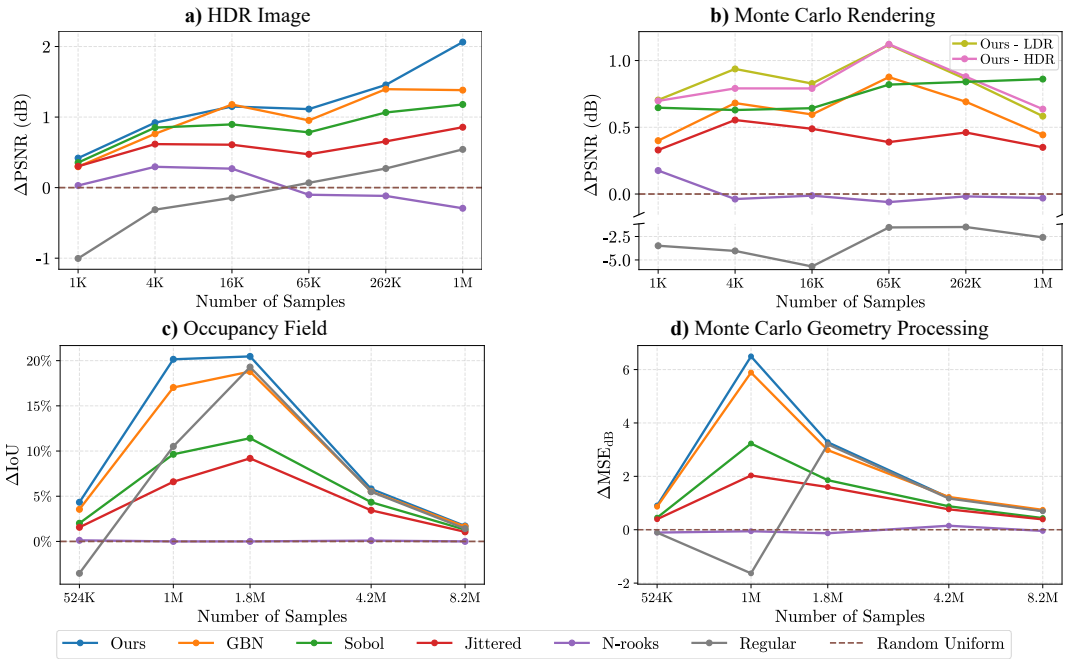

Figure 6: **Sampling pattern performance across tasks.** *(a)* HDR image fitting, reporting ΔPSNR. *(b)* Monte Carlo rendering, reporting ΔPSNR using the patterns optimized on LDR and HDR images. *(c)* 3D occupancy fitting, reporting ΔIoU. *(d)* Monte Carlo geometry processing, reporting $\Delta\text{MSE}_{\text{dB}}$ using the occupancy-optimized pattern. Higher is better in all panels.

Signed Distance Fields We evaluate SDF reconstruction at a single acquisition setting, using 1,024 samples per mesh-intersecting cell. Tab. 1 reports Chamfer-L_2 distance and IoU mean improvements for the RFF neural field, showing that our patterns remain competitive with GBN and generally achieve better reconstruction quality than other baselines.

Monte Carlo Geometry Processing As with 2D Monte Carlo rendering, directly optimizing sampling patterns on MCGP would require repeatedly querying the expensive solver. Therefore, we transfer the optimized pattern from the inexpensive occupancy field to MCGP and evaluate it on three meshes, with numerical comparisons visualized in Fig. 6(d). As demonstrated, despite being evaluated using patterns on different inexpensive 3D signals, our transferred pattern is competitive with or outperforms the baselines in expensive MCGP reconstruction, highlighting the generalization and efficiency of our framework. Qualitative results are provided in Fig. 5(b) and Appendix F.2.

Table 1: **SDF reconstruction.** We report ΔChamfer ($\times 10^{-7}$) and ΔIoU (%) ↑ over random uniform; higher is better. Best in bold, second-best underlined.

Sampling	ΔChamfer ↑	ΔIoU ↑
Regular	-0.226	0.16
N-rooks	0.343	0.00
Jittered	0.495	0.09
Sobol	0.469	0.13
GBN	<u>0.882</u>	**0.20**
Ours	**1.093**	<u>0.19</u>

5 Conclusion

We introduced a framework for optimizing sampling-pattern correlations specifically for neural field fitting. By parameterizing correlations through a compact spectral representation and evaluating candidate patterns with the complete fitting process in the loop, our approach identifies sampling structures that improve over uncorrelated and integration-oriented baselines, particularly under limited sampling budgets. The optimized patterns transfer across source signals and to settings with expensive target-signal queries, allowing inexpensive proxy tasks to guide sampling-pattern optimization. Our analysis further shows that the preferred correlations can exhibit nontrivial anisotropic structure. More broadly, our results suggest that correlations among training samples constitute an important degree of freedom for neural field optimization, motivating further study of their interaction with learning dynamics.

AI DISCLOSURE

In this work, we used LLMs to improve the writing, assist in creating experimental configurations, and help review code for potential issues. All suggested experimental configurations and code-related feedback were checked by the first author. We did not use generative AI to produce any figures; all figures were created by the authors. The authors take full responsibility for the final content of this work, including all text, claims, experimental configurations, code, and other artifacts produced with the assistance of generative AI or LLMs.

REFERENCES

Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.

Abdalla GM Ahmed, Jing Ren, and Peter Wonka. Gaussian blue noise. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022.

Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for bayesian optimization. *Advances in neural information processing systems*, 36:20577–20612, 2023.

Cagla Deniz Bahadir, Adrian V Dalca, and Mert R Sabuncu. Learning-based optimization of the under-sampling pattern in mri. In *International Conference on Information Processing in Medical Imaging*, pp. 780–792. Springer, 2019.

James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, volume 24, pp. 2546–2554, 2011.

Guillaume Bouchard, Théo Trouillon, Julien Perez, and Adrien Gaidon. Online learning to sample. *arXiv preprint arXiv:1506.09016*, 2015.

Fernando De Goes, Katherine Breeden, Victor Ostromoukhov, and Mathieu Desbrun. Blue noise through optimal transport. *ACM Transactions on Graphics (TOG)*, 31(6):1–11, 2012.

Bastien Doignies, Nicolas Bonneel, David Coeurjolly, Julie Digne, Loïs Paulin, Jean-Claude Iehl, and Victor Ostromoukhov. Example-based sampling with diffusion models. In *SIGGRAPH Asia 2023 conference papers*, pp. 1–11, 2023.

Bastien Doignies, David Coeurjolly, Nicolas Bonneel, Julie Digne, Jean-Claude Iehl, and Victor Ostromoukhov. Differentiable owen scrambling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 43(6), December 2024. doi: 10.1145/3687764.

Zhiwei Gao, Liang Yan, and Tao Zhou. Failure-informed adaptive sampling for pinns. *SIAM Journal on Scientific Computing*, 45(4):A1971–A1994, 2023.

Michael A. Gelbart, Jasper Snoek, and Ryan P. Adams. Bayesian optimization with unknown constraints. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pp. 250–259. AUAI Press, 2014.

Daniel Heck, Thomas Schlömer, and Oliver Deussen. Blue noise sampling with controlled aliasing. *ACM Transactions on Graphics (TOG)*, 32(3):1–12, 2013.

Xingchang Huang, Tobias Ritschel, Hans-Peter Seidel, Pooran Memari, and Gurprit Singh. Patternshop: editing point patterns by image manipulation. *ACM Transactions on Graphics (TOG)*, 42(4): 1–14, 2023.

Iris AM Huijben, Bastiaan S Veeling, and Ruud JG van Sloun. Deep probabilistic subsampling for task-adaptive compressed sensing. In *International Conference on Learning Representations*, 2020.

Bhavya Kailkhura, Jayaraman J Thiagarajan, Peer-Timo Bremer, and Pramod K Varshney. Stair blue noise sampling. *ACM Transactions on Graphics (TOG)*, 35(6):1–10, 2016.

Bhavya Kailkhura, Jayaraman Thiagarajan, Qunwei Li, Jize Zhang, Yi Zhou, and Timo Bremer. A statistical mechanics framework for task-agnostic sample design in machine learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 11925–11935, 2020.

Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. In *International conference on machine learning*, pp. 2525–2534. PMLR, 2018.

Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Accelerating neural field training via soft mining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. URL https://arxiv.org/abs/1412.6980.

Ares Lagae and Philip Dutré. A comparison of methods for generating poisson disk distributions. In *Computer Graphics Forum*, volume 27, pp. 114–129. Wiley Online Library, 2008.

Thomas Leimkühler, Gurprit Singh, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. Deep point correlation design. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2019)*, 38(6), 2019. doi: 10.1145/3355089.3356562.

Chen Liu, Michael Fischer, and Tobias Ritschel. Learning to learn and sample brdfs. *Computer Graphics Forum (Proceedings of Eurographics)*, 42(2):201–211, 2023. doi: 10.1111/cgf.14754.

Zhen Liu, Hao Zhu, Qi Zhang, Jingde Fu, Weibing Deng, Zhan Ma, Yanwen Guo, and Xun Cao. Finer: Flexible spectral-bias tuning in implicit neural representation by variable-periodic activation functions. In *Proceedings of the IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2024.

William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '87, pp. 163–169, New York, NY, USA, 1987. Association for Computing Machinery. ISBN 0897912276. doi: 10.1145/37401.37422. URL https://doi.org/10.1145/37401.37422.

Harald Niederreiter. Quasi-monte carlo methods and pseudo-random numbers. 1978.

Miles Olson, Elizabeth Santorella, Louis C. Tiao, Sait Cakmak, Mia Garrard, Samuel Daulton, Zhiyuan Jerry Lin, Sebastian Ament, Bernard Beckerman, Eric Onofrey, Paschal Igusti, Cristian Lara, Benjamin Letham, Cesar Cardoso, Shiyun Sunny Shen, Andy Chenyuan Lin, Matthew Grange, Elena Kashtelyan, David Eriksson, Maximilian Balandat, and Eytan. Bakshy. Ax: A Platform for Adaptive Experimentation. In *AutoML 2025 ABCD Track*, 2025.

Art B Owen. Scrambling sobol'and niederreiter–xing points. *Journal of complexity*, 14(4):466–489, 1998.

Adrien Pilleboue, Gurprit Singh, David Coeurjolly, Michael Kazhdan, and Victor Ostromoukhov. Variance analysis for monte carlo integration. *ACM Transactions on Graphics (TOG)*, 34(4):1–14, 2015.

Corentin Salaün, Xingchang Huang, Iliyan Georgiev, Niloy J. Mitra, and Gurprit Singh. Online importance sampling for stochastic gradient optimization. In *Proceedings of the 14th International Conference on Pattern Recognition Applications and Methods (ICPRAM)*, pp. 130–140, 2025.

Rohan Sawhney and Keenan Crane. Monte carlo geometry processing: A grid-free approach to pde-based methods on volumetric domains. *ACM Trans. Graph.*, 39(4), 2020.

Peter Shirley et al. Discrepancy as a quality measure for sample distributions. In *Eurographics*, volume 91, pp. 183–194, 1991.

Peter S. Shirley. *Physically Based Lighting Calculations for Computer Graphics*. PhD thesis, University of Illinois at Urbana-Champaign, 1991.

Gurprit Singh, Cengiz Öztireli, Abdalla GM Ahmed, David Coeurjolly, Kartic Subr, Oliver Deussen, Victor Ostromoukhov, Ravi Ramamoorthi, and Wojciech Jarosz. Analysis of sample correlations for monte carlo rendering. *Computer Graphics Forum*, 38(2):473–491, 2019.

Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, volume 25, 2012.

I Sobol. The distribution of points in a cube and the accurate evaluation of integrals (in russian) zh. *Vychisl. Mat. i Mater. Phys*, 7:784–802, 1967.

Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.

R.A. Ulichney. Dithering with blue noise. *Proceedings of the IEEE*, 76(1):56–79, 1988.

Li-Yi Wei and Rui Wang. Differential domain analysis for non-uniform sampling. *ACM Transactions on Graphics (TOG)*, 30(4):1–10, 2011.

Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Computer Graphics Forum*, 2022. ISSN 1467-8659. doi: 10.1111/cgf.14505.

Yifan Xu, Tianqi Fan, Yi Yuan, and Gurprit Singh. Ladybird: Quasi-monte carlo sampling for deep implicit field based 3d reconstruction with symmetry. In *European Conference on Computer Vision*, pp. 248–263. Springer, 2020.

Dong-Ming Yan, Jian-Wei Guo, Bin Wang, Xiao-Peng Zhang, and Peter Wonka. A survey of blue-noise sampling and its applications. *Journal of Computer Science and Technology*, 30(3):439–452, 2015.

Wang Yifan, Shihao Wu, Cengiz Oztireli, and Olga Sorkine-Hornung. Iso-points: Optimizing neural implicit surfaces with hybrid representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 374–383, 2021.

Chen Zhang, Steven Luo, Jason Li, Yik-Chung Wu, and Ngai Wong. Nonparametric teaching of implicit neural representations. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.

Chen Zhang, Wei Zuo, Bingyang Cheng, Yikun Wang, Wei-Bin Kou, Yik-Chung Wu, and Ngai Wong. Ntk-guided implicit neural teaching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.

Weixiang Zhang, Shuzhao Xie, Chengwei Ren, Siyi Xie, Chen Tang, Shijia Ge, Mingzi Wang, and Zhi Wang. Evos: Efficient implicit neural training via evolutionary selector. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 30472–30482. IEEE, 2025.

Peilin Zhao and Tong Zhang. Stochastic optimization with importance sampling for regularized loss minimization. In *international conference on machine learning*, pp. 1–9. PMLR, 2015.

Yahan Zhou, Haibin Huang, Li-Yi Wei, and Rui Wang. Point sampling with general noise spectrum. *ACM Transactions on Graphics (TOG)*, 31(4):1–11, 2012.

In this Appendix, we provide the implementation details of our framework (Appendix A), ablation studies for the optimized sampling patterns (Appendix B), neural field training configuration (Appendix C), task-specific specification (Appendix D), additional quantitative results for other neural field architectures (Appendix E), qualitative results (Appendix F), and further analysis of angular ambiguity and the implicit bias of our realization procedure (Appendix G).

A Implementation Details

Bayesian Optimization. We initialize each search with 50 Sobol proposals and continue until 400 feasible candidates have been evaluated, with up to 8 proposals evaluated in parallel. We limit $K = 3$ and $L = 2$ to keep the search computationally tractable. For each candidate radial spectrum, we train neural fields using a fixed acquisition budget: $65{,}536$ samples for 2D tasks and $1{,}769{,}472$ samples for 3D tasks. We then use 4×10^6 uniformly random samples to evaluate $\mathcal{L}$ of the resulting neural fields.

Radial spectrum parameterization. We parameterize the radial spectrum compactly using a cubic B-spline through M control points $\{(k_j, c_j)\}_{j=1}^{M}$. For 2D point patterns, we use eleven radial control locations at radii $\{0, 10, 20, 25, 30, 35, 40, 45, 50, 55, 60\}$. In preliminary experiments, the first four spectral values consistently converged toward zero. We therefore fix $c_1, \ldots, c_4 = 0$ and optimize the remaining seven spectral values. For 3D, we use seven radial control locations at $\{0, 5, 10, 12, 14, 16, 19\}$, fix $c_1 = c_2 = 0$, and optimize $c_3, \ldots, c_7$. All optimized spectral values c_j are constrained in $[0, 4]$. We use a small number of control points to balance spectral flexibility with the sample efficiency of Bayesian optimization.

Realization. In practice, we approximate the radial expectation in Eq. 4 by Monte Carlo sampling frequency directions at each radius (Leimkühler et al., 2019). For each candidate radial spectrum, we optimize the coordinates of a point set containing $1{,}024$ samples. We use L_1 loss and the Adam optimizer (Kingma & Ba, 2015) with 5000 iterations, with an initial learning rate of 10^{-3}, multiplied by 0.99 every 200 iterations. After 1000 iterations, we terminate the realization if the spectrum is deemed infeasible with threshold $\tau = 7$. To construct larger acquisition budgets for neural field training, we simply tile the base pattern over the domain. All experiments are done on an NVIDIA Quadro RTX 8000 GPU. Jointly realizing six patterns takes around 2 minutes.

Sampling domains. Samples are drawn uniformly over the domain $[-1, 1]^d$, except for SDF fitting. For SDF fitting, we assume access to a coarse geometric prior indicating the surface location to guide the sampling. We simulate the prior by partitioning the domain into a $20 \times 20 \times 20$ grid and marking cells that intersect the mesh. Sampling is then restricted to these cells.

B Ablation Study and Discussion

In this section, we discuss the generalization of sampling patterns across different source signals and neural field architectures (Appendix B.1), followed by anisotropy analysis (Appendix B.2) and the reproducibility of Bayesian optimization (Appendix B.3).

B.1 Generalization of Sampling Patterns

First, we investigate the generalization of optimized patterns across different neural field architectures. Specifically, we independently optimize sampling patterns for AAPE, RFF, and SIREN on 2D sRGB signals at a fixed sampling budget of 65K, and cross-evaluate each pattern with all three architectures. We report mean PSNR in Tab. 2 and visualize their corresponding spectra in Fig. 7(a). Sampling patterns optimized for different architectures achieve nearly identical performance when cross-evaluated, while exhibiting different spectra. Beyond architectures, our MCR and MCGP experiments in Sec. 4.2 further demonstrate that optimized patterns can transfer across source-signal domains,

Table 2: **Cross-architecture evaluation.** Mean PSNR ↑.

Pattern optimized for	Eval. Architecture		
	AAPE	RFF	SIREN
AAPE	28.49	29.52	29.43
RFF	28.48	29.51	29.41
SIREN	28.48	29.51	29.42
GBN	28.34	29.43	29.34

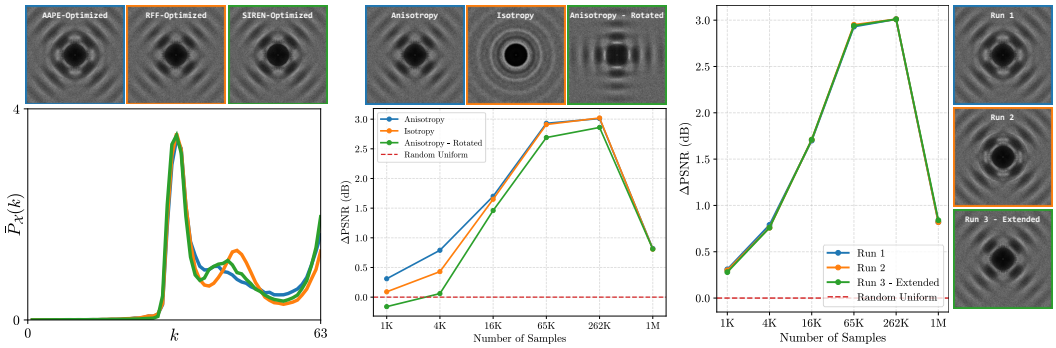

a) Full-dimensional and Radial Spectra **b)** Anisotropy Ablation **c)** Repeated BO runs

Figure 7: **Ablation study on optimized sampling patterns for LDR image fitting.** *(a)* Full-dimensional and radial spectra of patterns optimized independently for AAPE, RFF, and SIREN. *(b)* Anisotropy ablation comparing unregularized optimized pattern, isotropy-regularized pattern and $45°$ rotated anisotropic pattern. Performance is reported as ΔPSNR relative to random uniform. Higher is better. *(c)* Repeated BO runs, showing that BO can discover different spectra achieving similar reconstruction performance.

from inexpensive proxy signals to costly target signals. Together, these results suggest the optimized sampling patterns can generalize across source-signal domains and perform well on other architectures.

B.2 Anisotropy

We further investigate the effect of anisotropy on the reconstruction quality. We first enforce isotropy in the search space by introducing a penalty on directional spectral variation to the objective:

$$\lambda_{\text{iso}} \mathbb{E}_k \left[\text{Var}_{\mathbf{k}\,:\,\|\mathbf{k}\| \approx k} \left[P_{\mathcal{X}}(\mathbf{k}) \right] \right].$$

We set $\lambda_{\text{iso}} = 0.1$ and rerun Bayesian optimization on the LDR 2D image reconstruction task under this modified objective.

We additionally consider rotating the optimized anisotropic patterns by $45°$, which changes the orientation of the full-dimensional spectrum. We report the mean PSNR gain over random uniform of these patterns and their corresponding full-dimensional spectra in Fig. 7(b). Both isotropic and rotated anisotropic patterns perform worse than the optimized anisotropic patterns, especially in the low-budget regime; this suggests that in this setting, neural field training favors a specific anisotropic structure in the sampling pattern.

B.3 Bayesian Optimization Reproducibility

We investigate the reproducibility of Bayesian optimization by running the search independently multiple times. Run 1 and Run 2 are set to a BO budget of 400, while the budget of Run 3 is extended to 600 to explore the search space further. As demonstrated in Fig. 7(c) on the left, independent runs can discover high-performing patterns with very similar performance, but distinct spectra. This suggests the search space contains multiple high-performing spectra.

C Neural Field Training Configurations

For a fair comparison, we use the same neural-field training configuration for all acquisition patterns. In particular, the network architecture, initialization seed, training objective, optimizer, learning-rate schedule, and training budget are held fixed, such that only the acquisition pattern changes. All models are trained with Adam (Kingma & Ba, 2015) and cosine learning-rate decay. We use mean squared error as the training loss function, except for occupancy fitting, where we use binary cross-entropy. Please refer to Tab. 3 and Tab. 4 for detailed training configurations.

Table 3: Training configurations for 2D neural fields.

	AAPE	RFF	SIREN
Input dimension	2	2	2
Output dimension	3	3	3
Depth	4	4	4
Hidden width	256	256	256
Activation	ReLU	ReLU	–
AAPE frequencies	7	–	–
Training iterations	50,000	50,000	50,000
Batch size	2^{10}	2^{10}	2^{10}
Initial learning rate	10^{-3}	10^{-3}	10^{-4}
Minimum learning rate	10^{-5}	10^{-5}	10^{-6}
LR schedule	Cosine	Cosine	Cosine
Optimizer	Adam	Adam	Adam

Table 4: Training configurations for 3D neural fields.

	RFF	SIREN	FINER
Input dimension	3	3	3
Output dimension	1	1	1
Depth	9	9	9
Hidden width	256	256	256
Activation	ReLU	–	–
RFF feature dimension	256	–	–
RFF σ	12	–	–
Training iterations	5,000	5,000	5,000
Batch size	2^{15}	2^{15}	2^{15}
Initial learning rate	5×10^{-4}	10^{-4}	10^{-4}
Minimum learning rate	5×10^{-6}	10^{-6}	10^{-6}
LR schedule	Cosine	Cosine	Cosine
Optimizer	Adam	Adam	Adam

D Task-specific Specifications

D.1 Monte Carlo Rendering

For MCR, we evaluate on three scenes: Country Kitchen, The White Room, and Grey & White Room [3].

D.2 Occupancy Meshes

For occupancy field reconstruction, we evaluate on fourteen assets: Dragon, Happy Buddha, Lucy, Thai Statue, Bunny, Cheburashka, Cow, Fandisk, Homer, Horse, Igea, Indonesian Statue, Nefertiti, and Rocker Arm.

D.3 Signed Distance Field

For SDF reconstruction, we evaluate on five Stanford assets: Bunny, Dragon, Buddha, Lucy, and Thai statue.

[3] https://mitsuba.readthedocs.io/en/stable/src/gallery.html

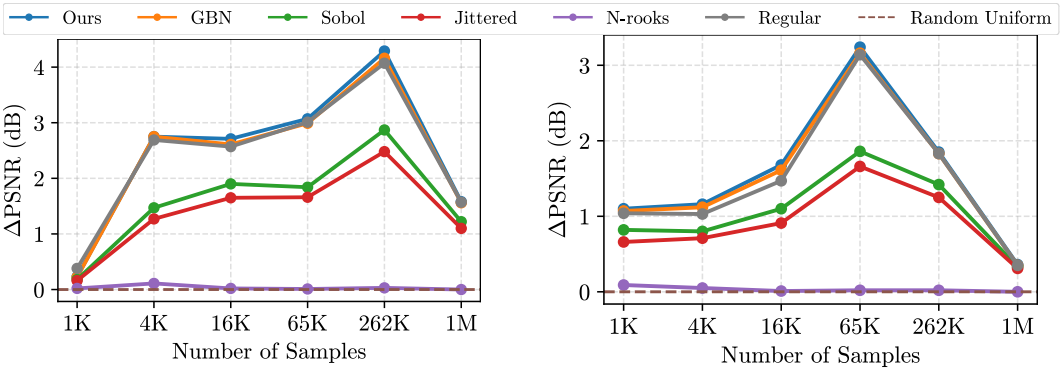

a) Random Fourier Features **b)** SIREN

Figure 8: **LDR image fitting across neural field architectures.** Mean ΔPSNR relative to random uniform sampling across budgets for *(a)* RFF and *(b)* SIREN; higher is better.

D.4 Monte Carlo Geometry Processing

We solve a Laplace equation defined on a geometry:

$$\Delta u(\mathbf{x}) = 0 \tag{7}$$

$$u = g \quad \text{if } \mathbf{x} \in \partial G \tag{8}$$

$$u = b \quad \text{if } \mathbf{x} \in \partial B \tag{9}$$

where G denotes the input geometry and $B = [-1, 1]^3$ denotes the bounding cube. We center and scale G to fit inside B.

On the geometry surface ∂G we define a spatially varying condition:

$$g(\mathbf{x}) = \sum_{i=1}^{10} a_i \cos(\omega_i^\top \mathbf{x} + \phi_i), \tag{10}$$

where ω_i, ϕ_i and a_i denote frequency, phase, and amplitude of the i-th cosine term, respectively. The amplitudes are normalized such that $\sum_i |a_i| = 1$. The parameters ω_i, ϕ_i and a_i are randomly initialized once and kept fixed for all experiments.

On the bounding surface of the cube ∂B, we define a piecewise-constant boundary condition:

$$b(\mathbf{x}) = \begin{cases} 1 & x + y + z > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

We evaluate on three geometries: Armadillo, Bunny, and Dragon.

E Additional Quantitative Results

E.1 LDR Images

We optimize the sampling patterns for RFF (Tancik et al., 2020) and SIREN (Sitzmann et al., 2020) on the LDR image-fitting task. As shown in Fig. 8, the optimized patterns remain competitive with GBN and outperform weaker baselines.

E.2 Occupancy Fields

We optimize the sampling patterns for SIREN (Sitzmann et al., 2020) and FINER (Liu et al., 2024) on the occupancy-fitting task. As shown in Fig. 9, the optimized patterns remain highly competitive with GBN and outperform all other baselines, with the highest gain at the low-budget regime. This demonstrates that our framework can discover high-performing patterns for other architectures.

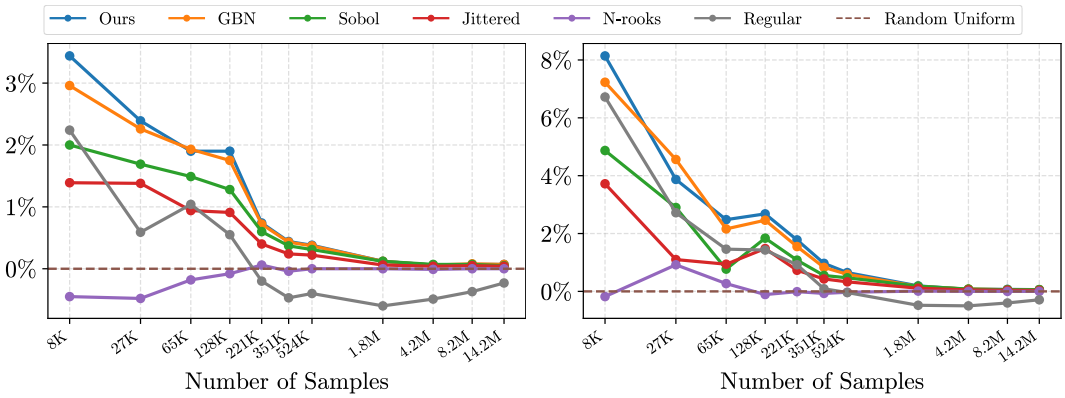

a) SIREN **b)** FINER

Figure 9: **Occupancy field fitting across neural field architectures.** Mean ΔIoU relative to random uniform sampling across budgets for *(a)* SIREN and *(b)* FINER; higher is better.

Table 5: **SDF reconstruction.** We report ΔChamfer ($\times 10^{-7}$) and ΔIoU (%) relative to random uniform sampling; higher is better. Best and second-best results are shown in bold and underlined, respectively.

	SIREN		FINER	
Sampling	ΔChamfer ($\times 10^{-7}$) ↑	ΔIoU (%) ↑	ΔChamfer ($\times 10^{-7}$) ↑	ΔIoU (%) ↑
Random	0.00	0.00	0.000	0.00
Regular	0.25	<u>0.04</u>	-0.401	<u>0.05</u>
N-rooks	0.99	0.01	-0.273	-0.01
Jittered	0.37	0.03	-0.386	0.03
Sobol	0.56	**0.05**	0.302	0.04
GBN	**1.23**	**0.05**	**0.842**	<u>0.05</u>
Ours	<u>1.09</u>	**0.05**	<u>0.623</u>	**0.06**

E.3 Signed Distance Fields

We optimize the sampling patterns for SIREN (Sitzmann et al., 2020) and FINER (Liu et al., 2024) on the SDF-fitting task. Tab. 5 summarizes the Chamfer-L_2 distance and IoU improvement for the two neural fields, demonstrating that our patterns remain competitive with GBN and generally achieve better reconstruction quality than other baselines.

F Additional Qualitative Results

F.1 2D Signal Reconstruction

LDR Images. We provide qualitative comparisons on held-out DIV2K images. Reconstructions are rendered on a $2{,}048 \times 2{,}048$ grid, with acquisition budgets shown by rows and sampling patterns by columns. We report the MSE and corresponding squared-error maps in Fig. 10, Fig. 11 and Fig. 12. These examples show the same overall behavior as the main-paper result: the optimized sampling pattern is competitive with GBN and yields lower reconstruction error than the generic baselines.

Monte Carlo Rendering We provide qualitative comparisons on fitting neural fields to signals whose queries are obtained from expensive Monte Carlo rendering. Reconstructions are rendered on a $2{,}048 \times 1{,}152$ grid, with acquisition budgets shown by rows and sampling patterns by columns. We report the MSE and corresponding squared-error maps in Fig. 13 and Fig. 14. These examples show the same overall behavior as the main-paper result: the optimized sampling patterns from inexpensive proxy tasks can transfer to this setting, yielding lower reconstruction error than the generic baselines.

F.2 3D Field Reconstruction

Monte Carlo Geometry Processing We provide qualitative comparisons of the RFF neural field fitted to solutions of Laplace equations defined on different geometries. We visualize three axial

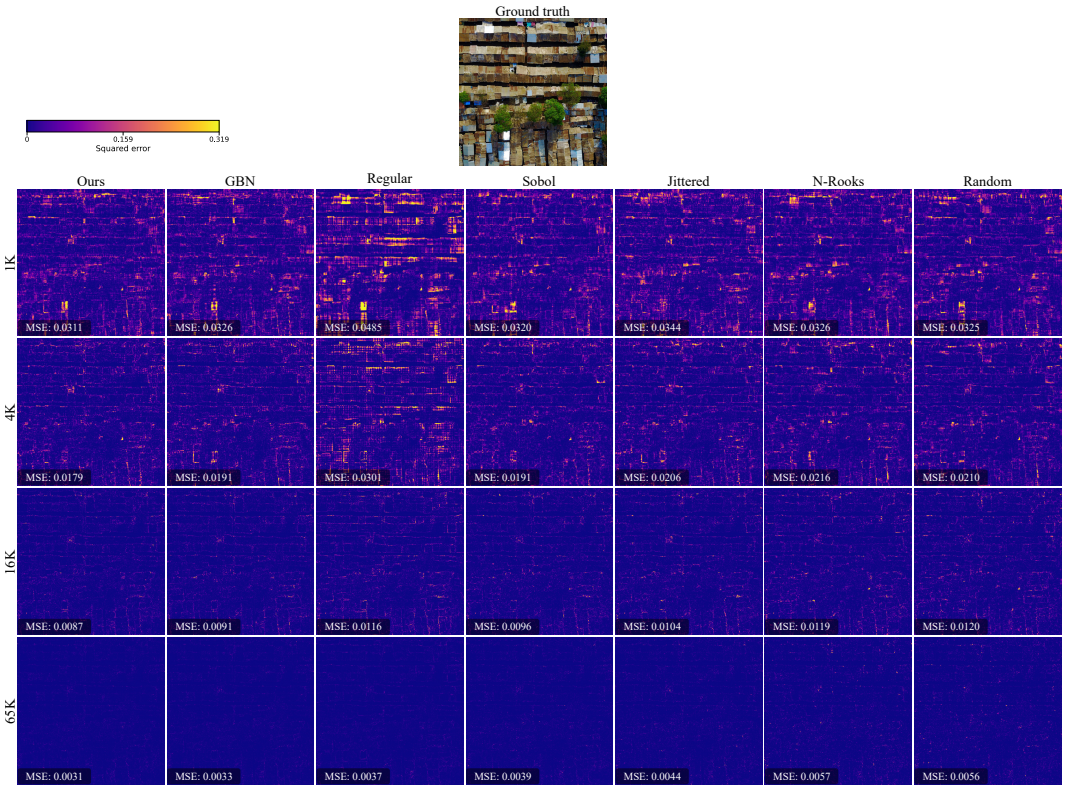

Figure 10: Visual comparison across sampling patterns and budgets. We report the corresponding MSE and error maps.

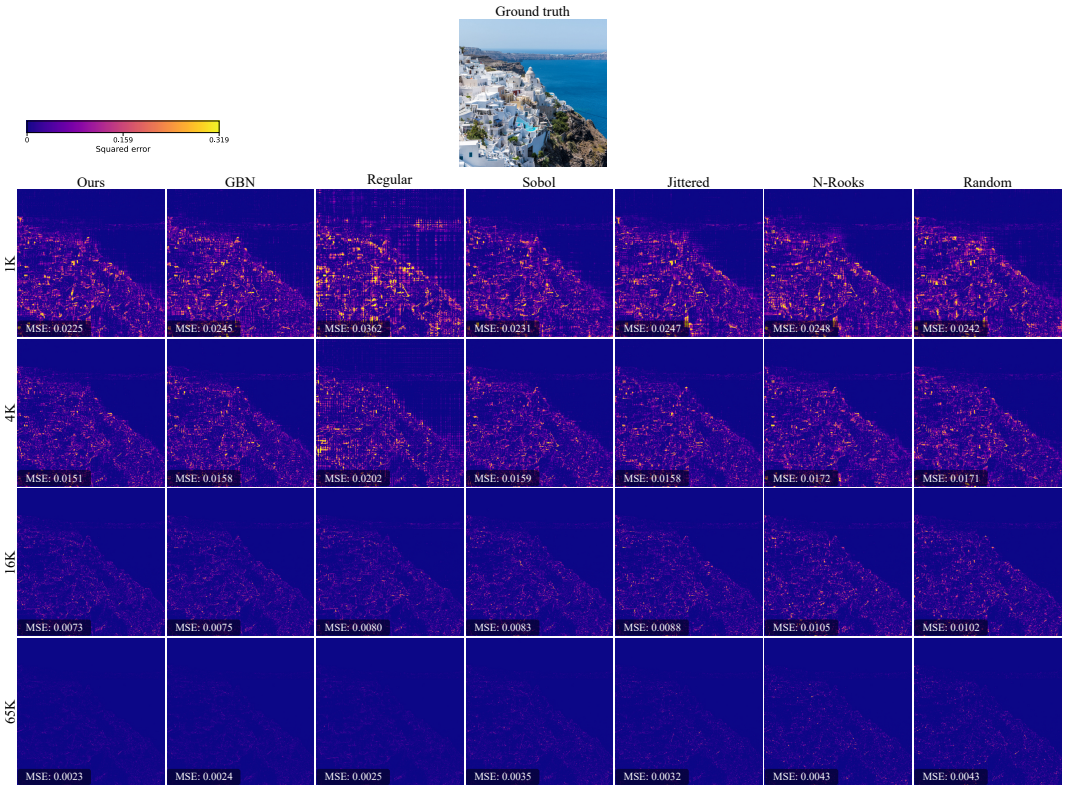

Figure 11: Visual comparison across sampling patterns and budgets. We report the corresponding MSE and error maps.

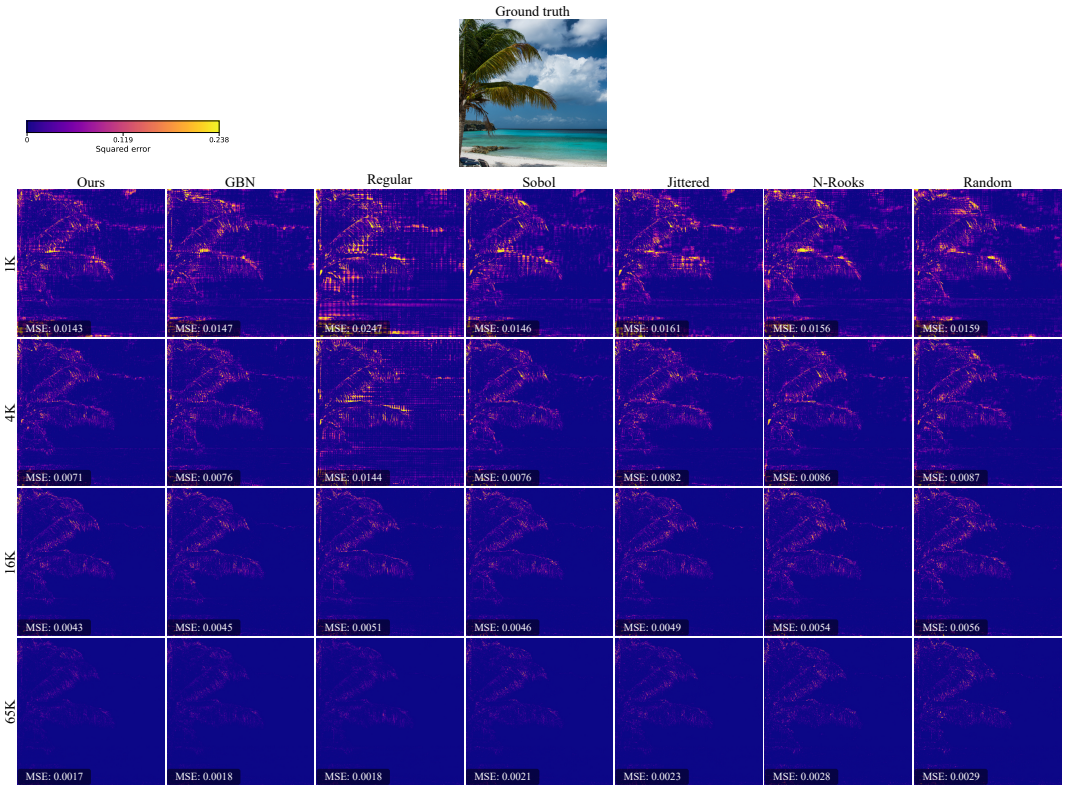

Figure 12: Visual comparison across sampling patterns and budgets. We report the corresponding MSE and error maps.

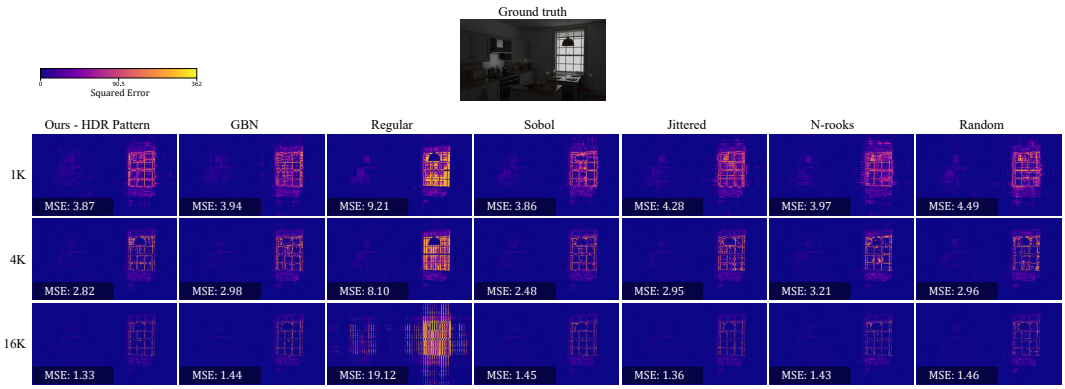

Figure 13: Visual comparison across sampling patterns and budgets. We report the corresponding MSE and error maps.

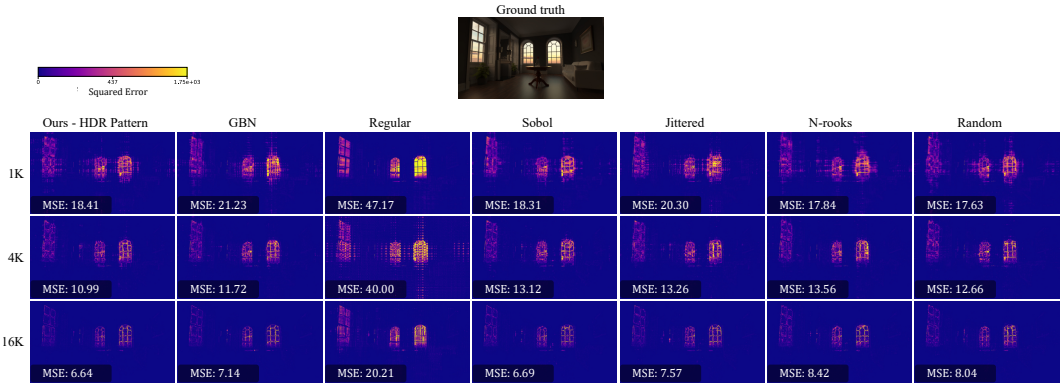

Figure 14: Visual comparison across sampling patterns and budgets. We report the corresponding MSE and error maps.

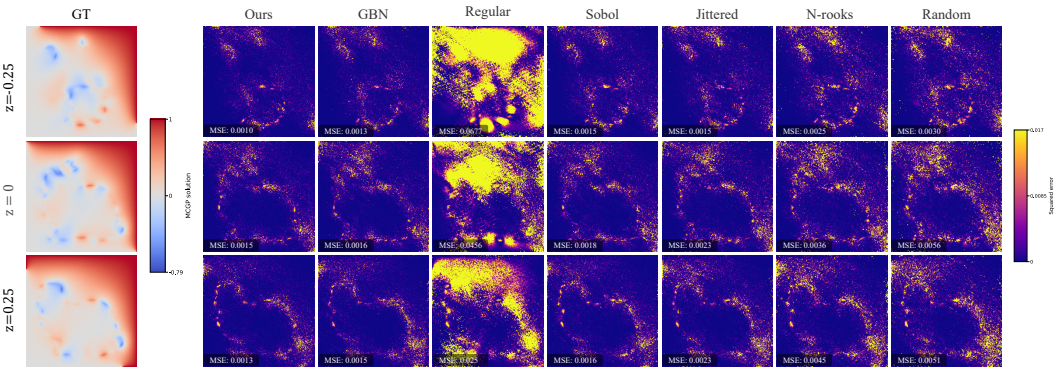

Figure 15: **RFF fitting to the solution of Laplace's equation on Bunny.** Ground-truth slices at $z \in \{-0.25, 0, 0.25\}$ and corresponding MSE and error maps for different sampling patterns at an acquisition budget of 1,024,000 samples. Error maps are rendered on a 256×256 grid for visualization.

slices of target fields obtained from the WoS Solver, and the corresponding MSE and error maps of neural fields trained using different sampling patterns in Fig. 15 and Fig. 16.

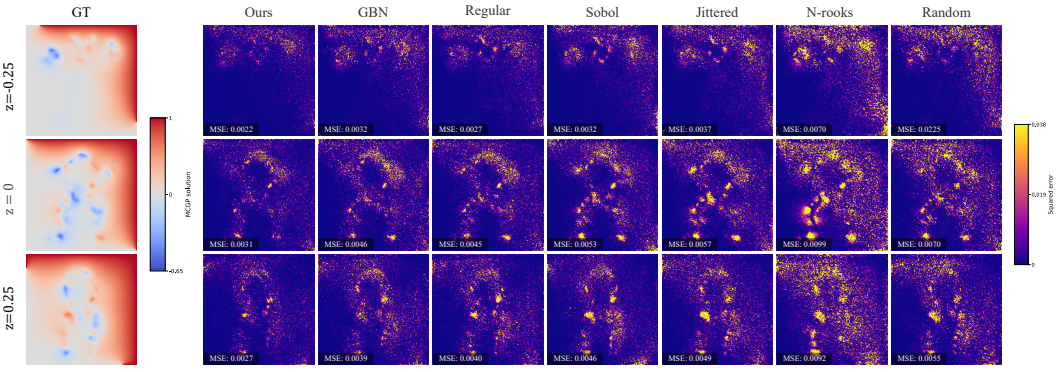

Figure 16: **RFF fitting to the solution of Laplace's equation on Armadillo.** Ground-truth slices at $z \in \{-0.25, 0, 0.25\}$ and corresponding MSE and error maps for different sampling patterns at an acquisition budget of 1,769,472 samples. Error maps are rendered on a 256×256 grid for visualization.

G Angular Ambiguity and Implicit Bias of Realization

Given the target radial spectrum $\bar{P}_{\mathcal{C}}$, our realization procedure seeks a point pattern $\mathcal{X}$ such that its radial power spectrum matches the target, by minimizing the objective:

$$\mathcal{X}^{\star}(\mathcal{C}, \zeta) = \arg\min_{\mathcal{X}} \mathbb{E}_k \|\bar{P}_{\mathcal{X}}(k) - \bar{P}_{\mathcal{C}}(k)\|. \tag{12}$$

For a point pattern on the unit torus, its power spectrum is evaluated at integer frequencies $\mathbf{k} \in \mathbb{Z}^d$,

$$P_{\mathcal{X}}(\mathbf{k}) = \frac{1}{N} \left| \sum_{i=1}^{N} \exp\left(-2\pi \mathrm{i}\, \mathbf{k}^{\mathsf{T}} \mathbf{x}_i\right) \right|^2, \tag{13}$$

and the radial spectrum is computed by averaging spectral power over frequency directions,

$$\bar{P}_{\mathcal{X}}(k) = \mathbb{E}_{\mathbf{k}\,:\,\|\mathbf{k}\| \approx k} \left[P_{\mathcal{X}}(\mathbf{k}) \right]. \tag{14}$$

The radial expectation can be estimated in different ways. One implementation is to average integer frequencies belonging to a radial bin. Another is to randomly sample unit directions uniformly, scale them by radius k, and then round the resulting frequencies to integer bins before averaging. In both cases, the radial averaging is a linear operator.

However, the optimization-based realization remains non-convex with respect to the point coordinates. In order to study how angular ambiguity arises in the optimized spectrum on average across realizations as well as the implicit inductive bias of the realization procedure, we instead consider a simplified problem in which the full-dimensional power spectrum is optimized directly.

G.1 Power Spectrum Optimization

Let

$$\mathcal{K} = \{\mathbf{k}_1, \cdots, \mathbf{k}_n\} \subset \mathbb{Z}^d$$

be a finite set of Fourier frequencies, and let

$$P = [P(\mathbf{k}_1), \cdots, P(\mathbf{k}_n)]^{\top} \in \mathbb{R}^n$$

be the vectorized power spectrum. The radial averaging can then be represented by a matrix $R \in \mathbb{R}^{m \times n}$, mapping a full-dimensional spectrum P to an m-dimensional radial spectrum. The simplified optimization problem then is:

$$\arg\min_{P \in \mathbb{R}^n} \|RP - \bar{P}_{\mathcal{C}}\|, \tag{15}$$

where $\|.\|$ denotes the chosen norm.

The structure of the radial averaging operator R depends on how the radial spectrum computation is implemented. One possible design is to group integer frequencies into radial bins. For bins $\{\mathcal{B}_k\}_{k=1}^{m}$, we define:

$$\mathcal{B}_k = \{\mathbf{k}_j \mid \mathrm{round}(\|\mathbf{k}_j\|) = k\} \tag{16}$$

and the corresponding averaging operator R is constructed as:

$$R_{kj} = \begin{cases} 1/|\mathcal{B}_k|, & \mathbf{k}_j \in \mathcal{B}_k, \\ 0, & \text{otherwise.} \end{cases} \tag{17}$$

An alternative is to uniformly sample the unit directions $\mathbf{u}$, scaling them to radius k, before mapping back to integer frequencies. Then, the entries of the operator R are:

$$R_{kj} = \Pr_{\mathbf{u}} \left[\mathrm{round}(k\mathbf{u}) = \mathbf{k}_j \right]. \tag{18}$$

Angular ambiguity. Since $m \ll n$, R has a large null space $\ker R$ with $\dim \ker R = n - \mathrm{rank}(R) \geq n - m$.

For any $v \in \ker R$, we have

$$R(P + v) = RP,$$

so all elements in $\ker R$ are invisible to the radial objective. This leaves the angular ambiguity and anisotropic structure in the spectrum unconstrained.

G.2 Implicit bias of gradient descent

Despite the existing ambiguity, gradient-based optimization on Eq. 15 does not freely explore the null space $\ker R$. In fact, consider an update with learning rate η:

$$g_t = R^\top \partial \|RP - \bar{P}_\mathcal{C}\| \tag{19}$$

$$P_{t+1} = P_t - \eta \cdot g_t, \tag{20}$$

we can observe that $g_t \in \mathrm{range}(R^\top) = (\ker R)^\perp$. This suggests that the optimization only traverses within the image space of $R^\top$.

Decompose the initialization P_0 into:

$$P_0 = P_0^{\|} + P_0^{\perp}, \qquad P_0^{\|} \in \mathrm{range}(R^\top), \quad P_0^{\perp} \in \ker R, \tag{21}$$

and we have

$$RP_0 = R(P_0^{\|} + P_0^{\perp}) = RP_0^{\|}. \tag{22}$$

Hence, every subsequent iterate can be rewritten as:

$$P_{t+1} = P_t^{\|} - \eta \cdot R^\top \partial \|RP_t^{\|} - \bar{P}_\mathcal{C}\| + P_0^{\perp}. \tag{23}$$

Thus, only the component $P_t^{||}$ moves within the image space $\mathrm{range}(R^\top)$ during the optimization, while $P_0^{\perp}$ remains fixed. Although the null space $\ker R$ gives rise to multiple spectra with the same target radial spectrum, gradient descent cannot move freely among these solutions. Instead, the final solution still inherits the null-space component of the initialization, inducing an initialization-dependent bias in the resulting spectrum.

Least Squares Objective. When the objective is least squares:

$$\arg\min_{P \in \mathbb{R}^n} \|RP - \bar{P}_\mathcal{C}\|_2^2, \tag{24}$$

the limiting solution of gradient descent admits a closed-form expression:

$$P_\infty = R^\dagger \bar{P}_\mathcal{C} + (I - R^\dagger R) P_0, \tag{25}$$

where $R^\dagger = R^\top (RR^\top)^{-1}$.

Thus, the resulting full-dimensional spectrum is determined by the target radial spectrum and the null-space component of the initialization. Consequently, when the initialization spectra are drawn from the same distribution, their realizations share the same limiting spectrum in expectation.

Random uniform initialization. Our coordinate-based realization initializes points independently and uniformly over the domain. In expectation, its power spectrum $\mathbb{E}[P_0] = \mathbf{1}$ for any non-zero integer frequency $\mathbf{k}$. The limiting solution of gradient descent is simply:

$$\mathbb{E}[P_\infty] = \mathbf{1} + R^\dagger(\bar{P}_\mathcal{C} - R\mathbf{1}). \tag{26}$$

G.3 Empirical Evidence

While the earlier analysis focuses on direct spectrum optimization, our realization procedure instead operates on point coordinates. However, empirically, we observe similar implicit bias. To demonstrate this, we fix a target radial spectrum and perform multiple realization runs from independent random 2D point set initializations using Adam (Kingma & Ba, 2015). We visualize the target and resulting radial spectra in Fig. 17 (a), and the corresponding full-dimensional spectra in Fig. 17 (b).

We consider three different realization objectives: L_2, L_1, and weighted-L_1. The weighted-L_1 loss is defined as:

$$\mathcal{L}_{\mathrm{w}L_1} = \mathbb{E}_k \left[w(k) \left| \bar{P}_\mathcal{X}(k) - \bar{P}_\mathcal{C}(k) \right| \right], \tag{27}$$

where

$$w(k) \propto (1 - k/k_{\max})^2 + 0.1, \tag{28}$$

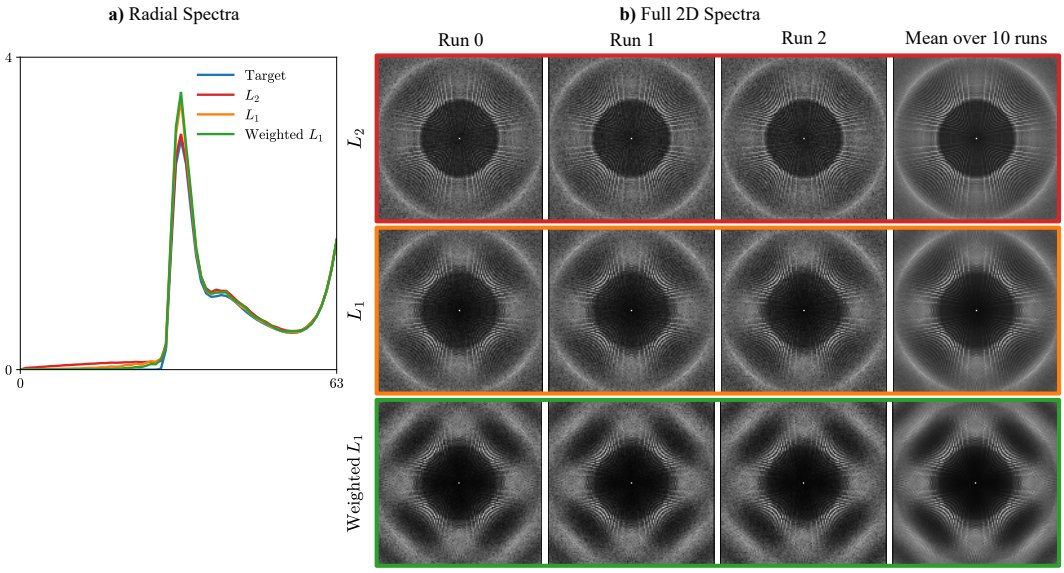

Figure 17: We realize the same target radial spectrum using L_2, L_1 and weighted-L_1 objectives. Initial point sets are randomly and uniformly initialized. *(a)* Target radial spectrum and mean realized radial spectra for three objectives. *(b)* Full two-dimensional spectra from three representative runs and the average spectra across 10 runs. Each run realizes five point patterns, and we report the mean spectrum of these five patterns.

$k_{\max}$ denotes the largest radial frequency (63 in our experiments), and the weights are normalized to have unit mean.

As noticed in Fig. 17 (a), the resulting radial spectra obtained from different loss functions closely match the target; however, their corresponding full-dimensional spectra differ substantially, as shown in Fig. 17 (b). For each objective, the resulting full-dimensional spectra exhibit a consistent anisotropic structure across multiple runs regardless of random initializations. The structure of the full-dimensional spectrum becomes particularly clear as we average over multiple runs. This indicates that the optimization resolves the angular ambiguity systematically rather than arbitrarily. In addition, changing the objective results in different, but again consistent anisotropic structure. This shows that the preferred anisotropy also depends on the implicit bias of the chosen realization objective.