PROPAGATING GEOMETRY UNCERTAINTY FROM A SINGLE IMPLICIT NEURAL REPRESENTATION THROUGH A DIFFERENTIABLE PDE SOLVER

Anonymous authors
Paper under double-blind review

ABSTRACT

Implicit neural representations (INRs) such as neural signed distance functions are increasingly used as the geometry for physics simulations that solve partial differential equations (PDEs). However, in these simulations, the reconstructed geometry is usually treated as deterministic even though it is fitted to noisy data (such as point cloud). The effect of these geometric uncertainties on the solution or quantities of interest (QoIs) of a PDE is unclear. In this work, we propagate this uncertainty from the geometry (a trained INR) to PDE solutions. We freeze the INR's nonlinear features and treat only its final layer as free, which makes the boundary an explicit, low-dimensional function of a small set of deformation modes; both inference and propagation are then closed-form within this span. The noisy scan yields a Gaussian posterior over the mode coefficients, which a Shifted Boundary Method (SBM) solver propagates to PDE QoIs 300–350$\times$ faster than Monte Carlo. We validate against two ground truths, exact random-field geometry and a controlled ground truth drawn from the posterior itself, to separate the accuracy of the posterior estimate from the expressive limits of the frozen representation. The result is close to ideal on the controlled test, and against exact geometry the posterior still accounts for most of the variability.

1 INTRODUCTION

Geometric reconstruction is often an intermediate step of a physics simulation: a point cloud is turned into a surface, and the volume enclosed by the surface is the domain of the partial differential equation (PDE). When the observations are noisy, the deterministic reconstructed surface does not describe the geometric uncertainty that propagates to the solution. In particular, neural signed-distance and occupancy fields are now standard shape representations (Park et al., 2019; Sitzmann et al., 2020; Gropp et al., 2020) and are increasingly used as simulation geometry (Karki et al., 2025; 2026a; Daviet et al., 2025). A reference procedure is to reconstruct many independently perturbed scans and solve the PDE on each, which requires repeated training and repeated simulation. Probabilistic surface reconstruction methods and neural-field uncertainty methods offer alternatives (Sellán & Jacobson, 2022; Holalkere et al., 2025; Desai, 2025); however, we are not aware of a demonstrated route for uncertainty propagation using these methods.

Performing uncertainty propagation using INRs raises three challenges. First, the geometry is the zero set of a neural field, which depends nonlinearly on thousands of network parameters, so a distribution over parameters does not directly give a distribution over geometric boundaries, and many parameter changes do not change the geometry at all. Second, scanner noise is typically spatially correlated; treating the points as independent inflates the importance of the scan's constraint on the shape. Third, the distribution must be over geometries that a PDE solver can use, and propagating it by Monte Carlo multiplies the cost of the simulation by the number of samples.

In this paper, we infer geometric uncertainty from a single trained INR and its scan, and propagate it to PDE quantities of interest (QoIs) (Figure 1). We avoid retraining by freezing the INR's nonlinear features, leaving the field affine in its final layer. The zero set's motion is nonlinear, but its first-order displacement is an explicit linear map. The map's orthonormal deformation modes come from a surface-weighted singular value decomposition. An analytic Gaussian posterior over 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.

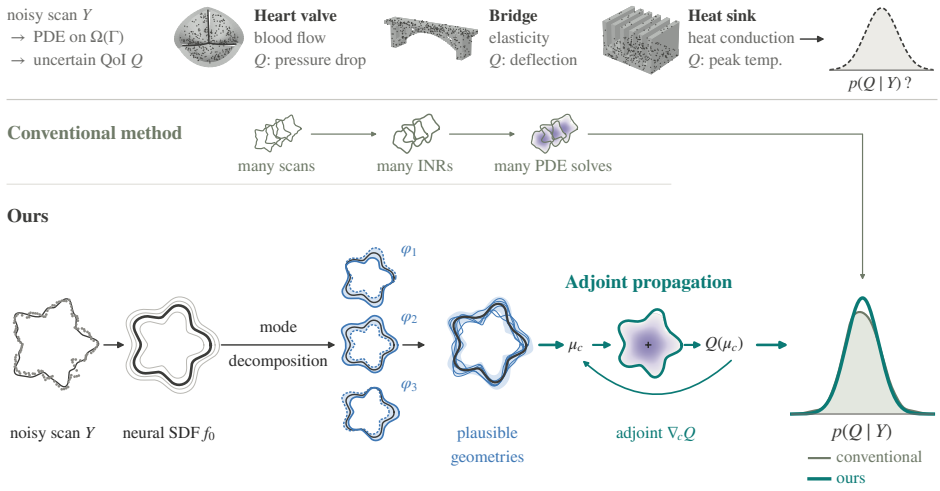


Figure 1: **Geometry uncertainty propagation from a single trained INR.** An INR fitted to a noisy point cloud provides boundary-deformation modes; the observations, their noise covariance and a coefficient prior define a posterior over mode coefficients, which a differentiable PDE solver propagates by adjoint-based linear propagation.

mode coefficients is obtained from the scan and its measurement-noise covariance, including spatial correlation. Samples are edited INRs, and a Shifted Boundary Method (SBM) solver (Karki et al., 2025; Atallah et al., 2020) evaluates them without remeshing; for a scalar QoI, linear propagation needs one forward and one adjoint solve. Because the modes are learned from the same scan, the construction is an empirical-Bayes approximation within a local family of surfaces.

Our contributions follow from a single construction: freezing a trained INR's nonlinear features makes its boundary an explicit, low-dimensional linear function of the final layer, giving an orthonormal span of deformation modes that everything else builds on.

1. **A geometric reduction of one trained INR:** the local map from final-layer perturbations to boundary motion, and surface-weighted deformation modes with an explicit map back to the network.
2. **Observation-conditioned inference and PDE propagation within that span:** a Gaussian coefficient posterior from a specified measurement model, propagated through an unfitted PDE solver with adjoint sensitivities.
3. **An evaluation that separates error sources:** comparisons with exact random-field geometry and with a controlled ground truth drawn from the posterior, which isolate the accuracy of inference and propagation within the span from what the span itself cannot express.

We evaluate this construction on three 2D shapes with correlated scan noise, using one network per shape, against two reference distributions detailed in Section 4: exact geometry, which involves no network, and a controlled ground truth drawn from the network's own posterior. The controlled test is close to ideal; against exact geometry the posterior captures most of the variability, with a shortfall consistent with what one network can represent.

2 Related Work

Neural implicit geometry and simulation. DeepSDF and Occupancy Networks established learned continuous shape representations (Park et al., 2019; Mescheder et al., 2019); periodic activations, Fourier features and geometric regularization control detail and derivatives (Sitzmann et al., 2020; Tancik et al., 2020; Gropp et al., 2020; Yüce et al., 2022). Neural implicit geometry has been coupled to flow and mechanics solvers (Karki et al., 2025; 2026a;b; Daviet et al., 2025); we build a distribution over such geometries and propagate it.

Probabilistic implicit surfaces. Gaussian-process implicit surfaces (Williams & Fitzgibbon, 2007), stochastic Poisson surface reconstruction (Sellán & Jacobson, 2022; Holalkere et al., 2025) and related formulations (Pujol & Chica, 2025; Genest & Coeurjolly, 2026) represent uncertainty in an implicit field. For neural fields, Bayes' Rays (Goli et al., 2024) and BayesSDF (Desai, 2025) estimate uncertainty of radiance and signed-distance fields, and latent-code or INR-based approaches address tomography and cardiac shapes (Vasconcelos et al., 2023; Verhülsdonk et al., 2026). We infer a distribution over boundary-deformation modes of one object-specific INR. This makes the boundary norm, parameter amplification, and truncation explicit, and connects the coefficients directly to a PDE solver.

Last-layer inference and geometric uncertainty in PDEs. Neural-linear and Bayesian last-layer models (MacKay, 1992; Snoek et al., 2015; Riquelme et al., 2018; Daxberger et al., 2021; Immer et al., 2021; Harrison et al., 2024) provide the Gaussian last-layer construction we build on; our geometric component is the observation operator linking scan residuals to motion of the zero set. PDEs on random domains (Xiu & Tartakovsky, 2006; Harbrecht et al., 2008; Castrillón-Candás et al., 2016; Badia et al., 2021), shape calculus (Sokołowski & Zolésio, 1992; Delfour & Zolésio, 2011), uncertain segmentations in simulation (Maher et al., 2021; Gralton et al., 2024) and Bayesian inverse problems (Stuart, 2010; Dunlop et al., 2017) study how a given geometric distribution affects solutions; we obtain that distribution from one trained network.

3 METHOD

A trained INR, with its nonlinear features frozen, defines a reduced family of nearby surfaces through perturbations of its final layer. We construct modes that are orthonormal in boundary displacement, infer a distribution over their coefficients from a measurement model, and propagate it to QoIs.

3.1 GEOMETRY AND OBSERVATION MODEL

Let $\Omega^\star \subset \mathbb{R}^d$, $d \in \{2, 3\}$, denote the unknown physical domain, with boundary $\Gamma^\star = \partial\Omega^\star$. We observe a point cloud

$$Y = \{\mathbf{y}_i\}_{i=1}^N, \qquad \mathbf{y}_i = \mathbf{x}_i^\star + \boldsymbol{\epsilon}_i, \qquad \mathbf{x}_i^\star \in \Gamma^\star. \tag{1}$$

The measurement errors may be anisotropic and spatially correlated. We write their covariance blocks as $\mathbf{C}_{ij} = \mathrm{Cov}(\boldsymbol{\epsilon}_i, \boldsymbol{\epsilon}_j)$.

A trained neural field is written as

$$f_0(\mathbf{x}) = \boldsymbol{\psi}(\mathbf{x})^\top \boldsymbol{\beta}_0, \qquad \Omega_0 = \{\mathbf{x} : f_0(\mathbf{x}) < 0\}, \qquad \Gamma_0 = \{\mathbf{x} : f_0(\mathbf{x}) = 0\}, \tag{2}$$

where $\boldsymbol{\psi} : \mathbb{R}^d \to \mathbb{R}^p$ contains the frozen penultimate features, including a constant feature for the final-layer bias. Throughout, signed distance is negative inside the domain. Consequently,

$$\mathbf{n}_0(\mathbf{x}) = \frac{\nabla f_0(\mathbf{x})}{\|\nabla f_0(\mathbf{x})\|} \tag{3}$$

is the outward unit normal on Γ_0.

We assume that Γ_0 is a compact, regular surface with a tubular neighborhood of unique closest-point projection, and that f_0 and the features are sufficiently smooth there, with $\|\nabla f_0\| \geq g_{\min} > 0$. These assumptions concern the learned reference surface; exact corners in a target geometry fall outside the smooth-surface analysis.

3.2 DEFORMATION MODES FROM A FROZEN INR

We now reduce the frozen INR to a small basis of boundary motions: an orthonormal set of deformation modes, together with an exact linear map back to the last-layer weights that produced them. Perturbing the final layer gives

$$f_{\delta\boldsymbol{\beta}}(\mathbf{x}) = f_0(\mathbf{x}) + \boldsymbol{\psi}(\mathbf{x})^\top \delta\boldsymbol{\beta}. \tag{4}$$

For sufficiently small perturbations, the nearby zero set can be written as a normal displacement of Γ_0: $\mathbf{x} \mapsto \mathbf{x} + h(\mathbf{x})\mathbf{n}_0(\mathbf{x})$. Expanding the zero-set condition yields

$$h(\mathbf{x}) = -\frac{\boldsymbol{\psi}(\mathbf{x})^\top \delta\boldsymbol{\beta}}{\|\nabla f_0(\mathbf{x})\|} + \mathcal{O}(\|\delta\boldsymbol{\beta}\|^2). \tag{5}$$

Thus, a negative field perturbation on the reference boundary produces positive outward motion. Appendix A establishes the local expansion and its remainder bound.

Let $\{\mathbf{q}_m\}_{m=1}^M \subset \Gamma_0$ be boundary samples, and let $\mathbf{W} = \mathrm{diag}(w_1, \ldots, w_M)$ contain positive, normalized surface quadrature weights, $\sum_m w_m = 1$. Define the boundary-displacement Jacobian and its weighted singular value decomposition (SVD):

$$A_{m,:} = -\frac{\boldsymbol{\psi}(\mathbf{q}_m)^\top}{\|\nabla f_0(\mathbf{q}_m)\|}, \qquad \mathbf{W}^{1/2}\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^\top. \tag{6}$$

Write the retained singular values as $s_1 \geq \cdots \geq s_r > 0$ and define

$$\mathbf{T} = \mathbf{V}_r\mathbf{S}_r^{-1}, \qquad \boldsymbol{\Phi} = \mathbf{A}\mathbf{T} = \mathbf{W}^{-1/2}\mathbf{U}_r. \tag{7}$$

The coefficient vector $\mathbf{c} \in \mathbb{R}^r$ specifies the last-layer update $\delta\boldsymbol{\beta} = \mathbf{T}\mathbf{c}$, whose first-order boundary displacement is $\boldsymbol{\Phi}\mathbf{c}$. By construction,

$$\boldsymbol{\Phi}^\top\mathbf{W}\boldsymbol{\Phi} = \mathbf{I}_r, \qquad \|\boldsymbol{\Phi}\mathbf{c}\|_{\mathbf{W}}^2 = \|\mathbf{c}\|_2^2, \tag{8}$$

where $\|\mathbf{v}\|_{\mathbf{W}}^2 = \mathbf{v}^\top\mathbf{W}\mathbf{v}$. The coefficients therefore measure amplitudes of boundary displacement in the chosen surface norm.

Radial construction used in the experiments. All scans in our experiments are taken along known rays from a center $\mathbf{o}$, and boundaries are extracted as radii $\rho(\theta)$. We therefore build the modes from the radial operator $A_{m,:} = -\boldsymbol{\psi}(\mathbf{q}_m)^\top/(\nabla f_0(\mathbf{q}_m)^\top\mathbf{e}_m)$, where $\mathbf{e}_m$ is the unit ray direction at $\mathbf{q}_m$, with arc-length weights $\mathbf{W}$. Then $\boldsymbol{\Phi}\mathbf{c}$ is the first-order radial displacement $\delta\rho(\theta)$, and the observation operator is $\boldsymbol{\Phi}$ evaluated at the scan angles. The normal and radial constructions coincide only for a circle centered at $\mathbf{o}$ (Appendix A).

For every coefficient vector, we evaluate the edited neural field

$$f_{\mathbf{c}}(\mathbf{x}) = f_0(\mathbf{x}) + \boldsymbol{\psi}(\mathbf{x})^\top\mathbf{T}\mathbf{c}, \qquad \Gamma(\mathbf{c}) = \{\mathbf{x} : f_{\mathbf{c}}(\mathbf{x}) = 0\}. \tag{9}$$

The field is exactly affine in $\mathbf{c}$, while its zero set depends nonlinearly on $\mathbf{c}$. In particular, $\boldsymbol{\Phi}\mathbf{c}$ predicts the first-order boundary motion; $\Gamma(\mathbf{c})$ is the realized geometry. We truncate poorly conditioned directions because

$$\|\mathbf{T}\mathbf{c}\|_2^2 = \sum_{k=1}^{r} \frac{c_k^2}{s_k^2}. \tag{10}$$

A small singular value can therefore require a large parameter update even for a modest displacement.

Retained modes, rank and rank rule. The SVD in Eq. 6 orders the deformation modes by singular value. The first r of them, represented by the pair $(\boldsymbol{\Phi}, \mathbf{T})$ of Eq. 7, are the *retained modes*; r is the *rank*, and their span is the *retained span*. The rank is set by the *rank rule*: keep modes with $s_k/s_1 \geq 10^{-3}$, at most 32; for the SIRENs t, so the cap alone sets the rank (Section 4.1). Validity at this cap is checked only for the tanh ablations (Appendix C), where it fails to produce closed geometry in some references by rank 13 and in most by rank 15. Modes beyond r are not retained; when they are appended to the posterior (Section 3.3) they are called *discrepancy modes*, and the retained modes are unchanged. The singular values describe geometric sensitivity; the probability law of $\mathbf{c}$ is specified separately.

3.3 Observation-conditioned uncertainty in the coefficients

Let d_0 be the signed distance to Γ_0, and define

$$\mathbf{s}_i = \mathrm{cp}_{\Gamma_0}(\mathbf{y}_i), \qquad z_i = d_0(\mathbf{y}_i), \qquad B_{ik} = \phi_k(\mathbf{s}_i). \tag{11}$$

Moving the surface outward decreases the signed distance of a fixed observation. To first order,

$$d_{\Gamma(\mathbf{c})}(\mathbf{y}_i) \simeq z_i - \mathbf{B}_{i,:}\mathbf{c}. \tag{12}$$

Projecting the measurement errors onto the nominal normals gives the linear observation model

$$\mathbf{z} = \mathbf{B}\mathbf{c} + \boldsymbol{\eta}, \qquad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \qquad R_{ij} = \mathbf{n}_0(\mathbf{s}_i)^\top \mathbf{C}_{ij} \mathbf{n}_0(\mathbf{s}_j). \tag{13}$$

This approximation freezes the correspondences and normals at the reference surface. The covariance retains spatial correlations in the measurement model.

The likelihood in Eq. 13 uses normal-distance residuals. For scans parameterized by known rays, we instead use radial residuals and the corresponding radial observation operator, derived in Appendix A. This preserves the distinction between radial and normal displacement.

We place a Gaussian prior on the retained coefficients,

$$\mathbf{c} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Lambda}), \qquad \boldsymbol{\Lambda} \succ 0. \tag{14}$$

The choice $\boldsymbol{\Lambda} = \tau^2 \mathbf{I}_r$ assigns equal prior variance to the orthonormal retained modes. The prior scale controls unresolved geometric variation, while $\mathbf{R}$ describes measurement uncertainty.

Conditioned on the learned representation, the linear-Gaussian model has the analytic posterior

$$\mathbf{c} \mid Y \sim \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c), \tag{15}$$

$$\boldsymbol{\Sigma}_c = \left(\boldsymbol{\Lambda}^{-1} + \mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B}\right)^{-1}, \qquad \boldsymbol{\mu}_c = \boldsymbol{\Sigma}_c \mathbf{B}^\top \mathbf{R}^{-1} \mathbf{z}. \tag{16}$$

Its linearized boundary mean and covariance are

$$\mathbb{E}[\mathbf{h} \mid Y] = \boldsymbol{\Phi} \boldsymbol{\mu}_c, \qquad \mathrm{Cov}(\mathbf{h} \mid Y) = \boldsymbol{\Phi} \boldsymbol{\Sigma}_c \boldsymbol{\Phi}^\top. \tag{17}$$

Sampling $\mathbf{c}$ and evaluating Eq. 9 produces corresponding edited zero sets without retraining the nonlinear features.

Discrepancy modes. Directions outside the retained span are implicitly fixed at the network's own shape, with zero variance. To admit uncertainty just outside the span, we append K discrepancy modes: the deformation modes $r+1, \ldots, r+K$ of the same SVD, which are orthogonal to the retained modes by construction and chosen without the truth. The retained and discrepancy modes receive the same prior and are conditioned jointly on the scan; the discrepancy modes enter the boundary band and the adjoint propagation, and Monte Carlo draws are realized through the network as before. We use $K = 10$.

Because the reference INR and its modes are constructed from Y, this is a data-adaptive, empirical-Bayes approximation conditioned on the learned representation. It does not marginalize uncertainty in the backbone or account automatically for geometric variation outside the retained family.

3.4 Propagation through a differentiable PDE solver

For the domain $\Omega(\mathbf{c}) = \{\mathbf{x} : f_{\mathbf{c}}(\mathbf{x}) < 0\}$, let the discrete PDE solution satisfy

$$\mathbf{F}(\mathbf{u}(\mathbf{c}), \mathbf{c}) = \mathbf{0}, \qquad Q(\mathbf{c}) = \mathcal{J}(\mathbf{u}(\mathbf{c}), \mathbf{c}). \tag{18}$$

Here $\mathbf{F}$ is the assembled residual and $\mathcal{J}$ is a scalar QoI, including any explicit dependence on geometry. Our experiments use a linear Poisson problem. Even when the PDE is linear in $\mathbf{u}$, its solution and QoIs generally depend nonlinearly on $\mathbf{c}$.

We use the SBM on a background mesh to evaluate these geometries without body-fitted remeshing. At a surrogate-boundary point $\widetilde{\mathbf{x}}$, define

$$\mathbf{p}_{\mathbf{c}} = \mathrm{cp}_{\Gamma(\mathbf{c})}(\widetilde{\mathbf{x}}), \qquad \mathbf{d}_{\mathbf{c}} = \mathbf{p}_{\mathbf{c}} - \widetilde{\mathbf{x}}. \tag{19}$$

Thus $\mathbf{d}_{\mathbf{c}}$ points from the surrogate boundary to the physical boundary. A first-order shifted Dirichlet trace uses $u(\widetilde{\mathbf{x}}) + \nabla u(\widetilde{\mathbf{x}})^\top \mathbf{d}_{\mathbf{c}}$ to approximate $u(\mathbf{p}_{\mathbf{c}})$. Geometric queries use the realized zero set; the edited neural field need not remain an exact signed distance function.

For local differentiation, the background active set is held fixed at the expansion geometry while the boundary projections, shifts and other geometry-dependent assembly terms vary with $\mathbf{c}$. This avoids discrete changes in the residual caused by element reclassification.

Linear propagation. Let $\bar{\mathbf{c}} = \boldsymbol{\mu}_c$ and solve the PDE at this coefficient vector. A first-order expansion gives

$$Q(\mathbf{c}) \simeq Q(\bar{\mathbf{c}}) + \mathbf{g}^\top (\mathbf{c} - \bar{\mathbf{c}}), \qquad \mathbf{g} = \nabla_{\mathbf{c}} Q(\bar{\mathbf{c}}). \tag{20}$$

All derivatives below are evaluated at $(\mathbf{u}(\bar{\mathbf{c}}), \bar{\mathbf{c}})$. The adjoint and total coefficient gradient satisfy

$$\mathbf{F}_{\mathbf{u}}^{\top}\boldsymbol{\lambda} = \nabla_{\mathbf{u}}\mathcal{J}, \qquad \mathbf{g} = \nabla_{\mathbf{c}}\mathcal{J} - \mathbf{F}_{\mathbf{c}}^{\top}\boldsymbol{\lambda}. \tag{21}$$

Consequently,

$$Q \mid Y \;\approx\; \mathcal{N}\big(Q(\boldsymbol{\mu}_c), \mathbf{g}^{\top}\boldsymbol{\Sigma}_c\mathbf{g}\big)\,. \tag{22}$$

For a scalar QoI, this requires one forward and one adjoint solve, together with the evaluation of the residual derivatives.

Monte Carlo propagation. To assess the approximation in Eq. 20, posterior Monte Carlo (MC) propagation draws $\mathbf{c}^{(j)} \sim \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)$, constructs $\Gamma(\mathbf{c}^{(j)})$ and solves the PDE for each draw. Both propagation routes use the same Gaussian coefficient posterior. Their comparison isolates the effect of nonlinear geometry-to-QoI propagation.

The method applies when the coefficient distribution is concentrated on regular, nearby geometries for which the local surface model and the numerical PDE remain valid; draws whose zero set is not a single closed curve are counted as invalid and reported.

4 Experiments

We evaluate four claims: that the modes span meaningful shape variation (Section 4.1); that the posterior captures the variability a network's own retained modes can express (Section 4.2 and Section 4.3); that linear propagation matches Monte Carlo and the ground truth (Section 4.4); and that it does so at a fraction of the propagation cost, using a circle, a multi-lobed flower and a sharp-cornered rectangle to test the smooth-surface assumption of Section 3.1 where it is most likely to fail.

Setup. We use a circle, a five-lobe flower and a rectangle, each observed through a scan of 500 points with correlated noise ($\sigma = 0.0025$; squared-exponential in arc length with correlation length $\ell = 0.2R_0$ for the circle and $0.1R_0$ otherwise). For each shape one INR, a SIREN (128×2 for the circle and flower, 128×3 for the rectangle), is trained on one scan. The posterior uses $r = 32$ retained modes, $K = 10$ discrepancy modes and $\tau = 5\sigma$. We compare with two ground truths. *Exact geometry*: 500 curves obtained by adding to the true boundary a Gaussian random field with the smooth part of the scan-noise covariance; no network is involved. *Controlled*: 500 curves drawn from the network's own posterior (retained modes) and extracted as explicit zero sets. The controlled test is an internal-consistency check of realization, linearization and propagation. Details are in Appendix A and Appendix B.

4.1 Deformation modes and spectral structure

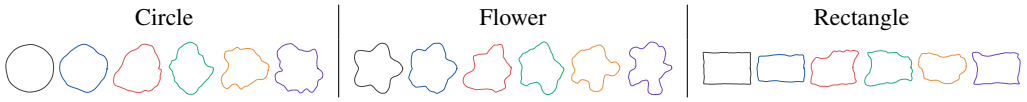


Figure 2: **Deformation modes of one trained INR.** Each strip shows the reference zero set (left) and five realized geometries, each obtained by varying a single mode coefficient (modes 1–5), $\delta\boldsymbol{\beta} = \mathbf{T}\mathbf{c}$; amplitudes are illustrative, not uncertainty levels.

Figure 2 shows realized geometries obtained by varying one coefficient at a time. The normalized singular values (Figure 3) decay slowly and without a gap, $s_4/s_1 \approx 0.44$–0.54 and $s_{32}/s_1 \approx 0.06$–0.10: 70–95 modes pass the relative threshold, so the cap of 32 sets the rank. The singular values measure geometric sensitivity to last-layer updates and order the modes from global motions to local detail; they do not by themselves set the coefficient variances.

4.2 Representational capacity against exact geometry

Because the spectrum has no gap, the cap rather than the threshold sets the rank, so the retained modes may not span the full variability of a given shape; we now measure what they can express. We fit the retained modes of the last layer to each ground-truth curve and compare the realized zero set with that curve (Figure 4). Against exact geometry, the fit recovers 0.90, 0.66 and 0.66 of each

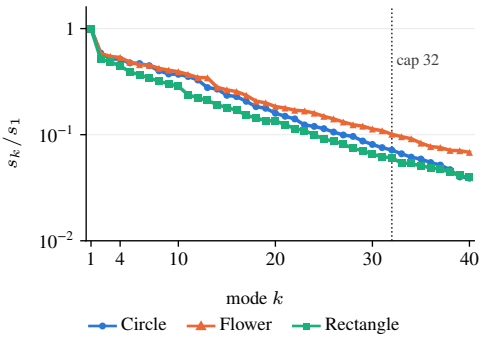


Figure 3: **Spectral structure.** Normalized singular values s_k/s_1 of each shape's deformation modes. The spectrum decays slowly and has no gap; the dotted line marks the rank cap 32, and all plotted values lie above the relative threshold 10^{-3}.

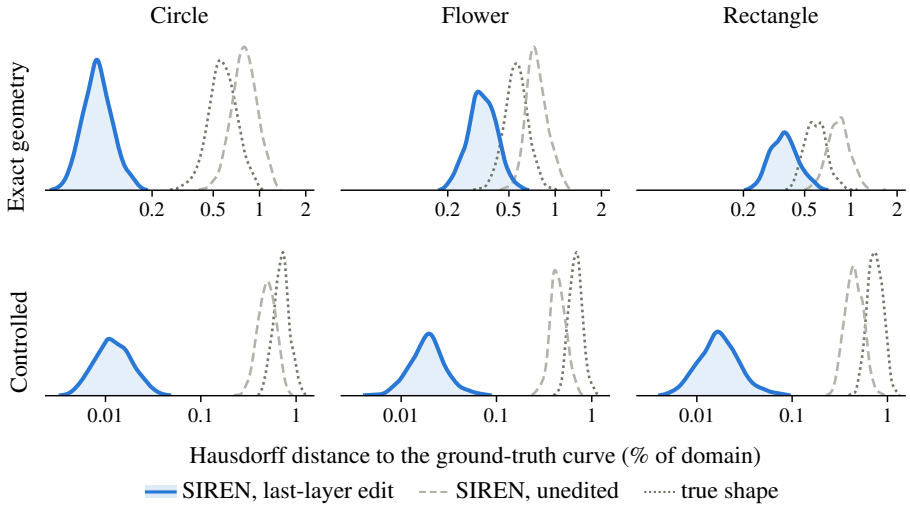


Figure 4: **Representational capacity of one network.** The retained modes of the last layer are fitted to each of 500 exact random-field curves (top) and to each of 500 curves drawn from the network's own posterior (bottom). The unedited network's own zero set ($\mathbf{c} = \mathbf{0}$) is shown for reference; it is not itself fitted to either curve set.

curve's displacement (defined in Appendix B)for the circle, flower and rectangle (median Hausdorff distance 0.09%, 0.33%, 0.37% of the domain). In the controlled test it recovers 0.97–0.98 (median distances 0.012–0.019%), the residual being the nonlinearity of the zero set in the coefficients. The gap between the rows, about 0.08 of the displacement for the circle and 0.3 for the flower and rectangle, measures what the retained modes of one network cannot express.

4.3 Posterior variability against exact geometry

Table 1 compares the posterior with each ground-truth ensemble through four metrics, each answering a different question (Figure 5; definitions in Appendix B): whether the posterior has the right total *amount* of boundary variation (variance ratio, ideal 1); whether it varies in the right *places* (covariance difference of trace-normalized covariances, compared with a split-half baseline of the ground truth); whether the variation of an *integrated* quantity such as the enclosed area is right (area standard-deviation ratio, ideal 1); and whether its linear 5–95% band contains the *truth* as often as the nominal 0.90 (coverage).

In the controlled test the metrics are close to their ideal values: variance ratios 0.97–1.03, covariance differences below the split-half baseline, and coverage 0.90–0.91, although the flower's area standard-deviation ratio is 0.92. Against exact geometry the same network's posterior captures 68–88% of the boundary variance and covers the true boundary at 0.83–0.90, with the flower lowest. Because the 1000 boundary rays lie on one correlated truth, not independent trials, the effective sample size is much smaller, so the flower's shortfall is suggestive rather than precisely estimated. The error of the posterior mean (0.9–1.0σ in arc-length RMS) exceeds the band (0.8–0.9σ) on the flower and rectangle. This shortfall is consistent with the representational limit of Section 4.2: the scan does not reveal to the posterior the directions that the retained modes cannot express.

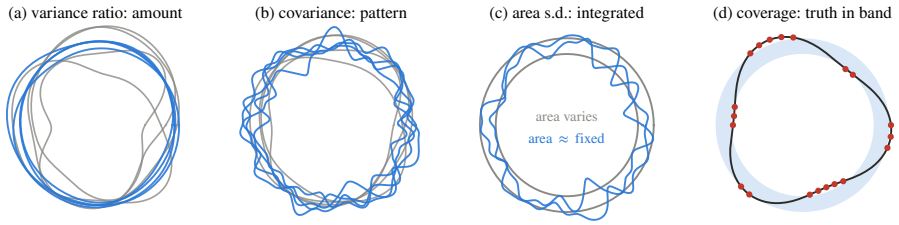


Figure 5: **What each metric of Table 1 measures** (schematic with synthetic curves; blue: posterior, gray: ground truth). (a) Variance ratio: total amount of boundary variation. (b) Covariance difference: where the variation occurs. (c) Area standard-deviation ratio: an integrated quantity. (d) Coverage: the share of the boundary within the posterior's 5–95% band (shaded; red if it falls outside).

Table 1: **Posterior variability against exact geometry and in the controlled test.** Exact: $r = 32$ retained and $K = 10$ discrepancy modes against 500 random-field curves; coverage of the true boundary. Controlled: the retained-mode posterior against 500 curves drawn from that same retained-mode posterior; coverage of those curves. Var.: variance ratio; cov. diff.: covariance difference with the split-half baseline in brackets; area s.d.: area standard-deviation ratio; cover.: share of 1000 boundary rays inside the linear 5–95% band (nominal 0.90).

	Exact geometry				Controlled			
Shape	var.	cov. diff.	area s.d.	cover.	var.	cov. diff.	area s.d.	cover.
Circle	0.88	0.27 (0.37)	0.93	0.90	0.97	0.21 (0.36)	1.01	0.90
Flower	0.81	0.40 (0.51)	0.97	0.83	0.97	0.27 (0.42)	0.92	0.90
Rectangle	0.68	0.45 (0.50)	0.91	0.87	1.03	0.27 (0.37)	1.03	0.91

4.4 PDE UNCERTAINTY PROPAGATION

We now test whether the boundary-level accuracy of Section 4.3 extends to a downstream PDE quantity, since the geometry-to-QoI map is nonlinear (Section 3.4). We solve $-\Delta u = 1$ in Ω, $u = 0$ on $\partial\Omega$, with the SBM on a level-8 quadtree (Appendix A), and evaluate u at 16 probes on the ray $\theta = 0$. We report $u - u^\star$, where $u^\star$ is the same discrete solve on the true shape. Linear propagation (Eq. 22) and posterior Monte Carlo (1000 draws) use the posterior of Section 4.3; the ground-truth distributions are the solutions on the 500 exact and 500 posterior-drawn curves.

In the controlled test, linear propagation, Monte Carlo and the ground truth agree closely (Figure 6, controlled-test panels): the linear standard deviation is 0.88–1.03 of the ground truth's (Monte Carlo 0.94–1.06), and the linear 5–95% interval contains 0.85–0.91 of the ground-truth solutions, with the flower's shortfall reflecting the linearization of the geometry-to-QoI map. Against exact geometry the linear spread is 0.90–0.93 of the ground truth's, and the 5–95% interval contains $u^\star$ at all 16 probes for each shape; without discrepancy modes the flower's interval misses $u^\star$ at 2 probes. Once the INR is trained (78–151 s), linear propagation takes 1.2–1.7 s per QoI against 6–10 min for 1000 Monte Carlo solves, about 300–350× faster; training an ensemble of 500 retrained INRs took 13–45 h (Table 2), a 605–1667× speedup end to end, including the training our method also requires.

4.5 ABLATIONS

Deformation space and rank. Tanh networks have a spectral drop of 5–7× after the third mode, keep 13–18 modes and recover only 0.66, 0.43 and 0.18 of the exact curves; the SIREN's flatter spectrum raises this to 0.66–0.90. Width, depth and skip connections matter little or hurt recovery. Increasing the rank does not remove the limit: a fit with the whole last layer would recover up to 0.99 linearly, but most such edits no longer give a closed zero set, and with tanh networks invalid draws appear in seven of nine references at rank 15 and in most draws at rank 17 (Appendix C). *Discrepancy modes.* They raise coverage against exact geometry from 0.81–0.87 to 0.83–0.90 and slightly over-cover the controlled test (0.92–0.93). *Scan model and likelihood.* Training on scans without the independent nugget (a small noise-floor term added to the covariance diagonal to keep it from being singular) and keeping a 5% nugget in the likelihood performs best; without the nugget the likelihood is nearly singular and the posterior collapses (coverage ≤ 0.15).

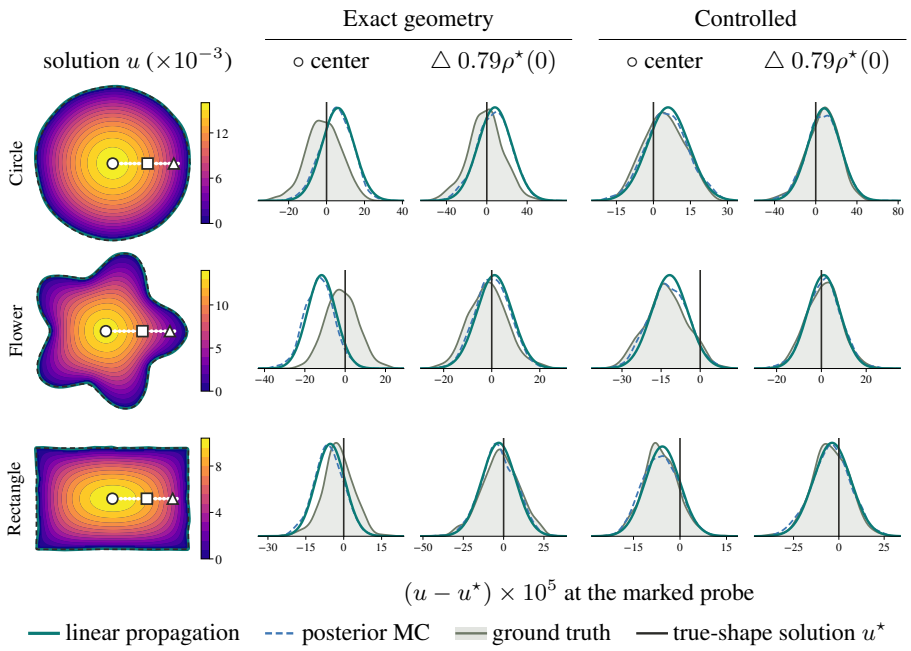


Figure 6: **Propagation of single-network geometry uncertainty through the PDE.** Leftmost panel: Poisson solution at the posterior-mean geometry, true boundary dashed, probes marked. Exact-geometry panels: linear propagation and posterior Monte Carlo against the solutions on 500 exact random-field curves; controlled-test panels: against 500 curves drawn from the posterior itself.

Table 2: **Cost of propagating geometry uncertainty to one PDE quantity of interest** (single core, Intel Xeon w7-3445, double precision; medians of three trials, which vary by up to about 2× on this shared machine). Linear propagation includes building the retained and discrepancy modes and posterior and one forward plus one adjoint solve. Speedup over ensemble is reported both for linear propagation alone and for the end-to-end cost, including the one-time INR training that the ensemble also pays per member. †Measured: training and one solve for each of the 500 retrained tanh INRs of the campaign.

	Circle	Flower	Rectangle
Training the INR (once)	78 s	96 s	151 s
Linear propagation	1.7 s	1.2 s	1.7 s
Posterior MC	9.6 min	6.5 min	8.6 min
Speedup over MC	**343**×	**317**×	**305**×
Retrained ensemble, 500 INRs†	13.4 h	45.0 h	42.0 h
Speedup over ensemble (linear only)	$\mathbf{2.9 \times 10^4}$	$\mathbf{1.3 \times 10^5}$	$\mathbf{9.0 \times 10^4}$
Speedup over ensemble (end to end)	**605**×	**1667**×	**990**×

5 Discussion and Limitations

Our method quantifies uncertainty *within* a representation learned from one scan. The posterior places mass only on the retained and discrepancy modes, so it cannot express variability such as the rectangle's corners; a shared corner bias, confirmed with tanh ensembles, places the true PDE solution outside the posterior near the rectangle's wall (Appendix B). By contrast, the controlled test confirms realization and propagation are accurate within the posterior's own law, but does not establish calibration against external geometry, where the flower's coverage remains below nominal. More fundamentally, the construction is local: large perturbations can invalidate the normal-graph description or change the surrogate domain, requiring Monte Carlo or nonlinear propagation instead. In addition, results depend on the noise model, including the likelihood nugget, and on the network architecture. Finally, our experiments are two-dimensional with a linear elliptic PDE; three-dimensional surfaces and nonlinear PDEs may need more modes and less local propagation.

AI USE STATEMENT

Generative AI tools were used to help draft and edit the manuscript text and to write plotting code for the figures from saved experiment outputs. Every number was checked against the run records, and the authors take full responsibility for the content.

REPRODUCIBILITY STATEMENT

All computations are in float64 with fixed seeds; INR training and PDE solves run on the CPU, and batched zero-set extraction and the exact-geometry PDE solves optionally run on a GPU (identical to the CPU path to 10^{-13}). Appendix A lists the observation model, prior, networks, training and PDE settings.

REFERENCES

Nabil M. Atallah, Claudio Canuto, and Guglielmo Scovazzi. The second-generation Shifted Boundary Method and its numerical analysis. *Computer Methods in Applied Mechanics and Engineering*, 372:113341, 2020. ISSN 0045-7825. doi: 10.1016/j.cma.2020.113341. URL https://www.sciencedirect.com/science/article/pii/S0045782520305260.

Santiago Badia, Jerrad Hampton, and Javier Principe. Embedded multilevel Monte Carlo for uncertainty quantification in random domains. *International Journal for Uncertainty Quantification*, 11(1):119–142, 2021. doi: 10.1615/Int.J.UncertaintyQuantification.2021032984.

Julio E. Castrillón-Candás, Fabio Nobile, and Raúl F. Tempone. Analytic regularity and collocation approximation for elliptic PDEs with random domain deformations. *Computers and Mathematics with Applications*, 71(6):1173–1197, 2016. ISSN 0898-1221. doi: 10.1016/j.camwa.2016.01.005. URL https://www.sciencedirect.com/science/article/pii/S0898122116000080.

Gilles Daviet, Tianchang Shen, Nicholas Sharp, and David I.W. Levin. Neurally integrated finite elements for differentiable elasticity on evolving domains. *ACM Transactions on Graphics*, 44(2), April 2025. ISSN 0730-0301. doi: 10.1145/3727874. URL https://doi.org/10.1145/3727874.

Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux – effortless Bayesian deep learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 20089–20103. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/a7c9585703d275249f30a088cebba0ad-Paper.pdf.

M. C. Delfour and J.-P. Zolésio. *Shapes and Geometries: Metrics, Analysis, Differential Calculus, and Optimization*. Society for Industrial and Applied Mathematics, second edition, 2011. doi: 10.1137/1.9780898719826.

Rushil Desai. BayesSDF: Surface-based Laplacian uncertainty estimation for 3D geometry with neural signed distance fields, 2025. ICCV 2025 Workshops (per arXiv comment).

Matthew M. Dunlop, Marco A. Iglesias, and Andrew M. Stuart. Hierarchical Bayesian level set inversion. *Statistics and Computing*, 27(6):1555–1584, 2017. doi: 10.1007/s11222-016-9704-8.

Baptiste Genest and David Coeurjolly. Uncertainty-aware geometry processing on Gaussian process implicit surfaces. *ACM Transactions on Graphics*, 45(4), July 2026. ISSN 0730-0301. doi: 10.1145/3811280. URL https://doi.org/10.1145/3811280.

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. Bayes' Rays: Uncertainty quantification for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20061–20070, 2024. doi: 10.1109/CVPR52733.2024.01896.

Stephen G. Gralton, Farah Alkhatib, Benjamin Zwick, George Bourantas, Adam Wittek, and Karol Miller. Random Boundaries: Quantifying segmentation uncertainty in solutions to boundary-value problems. In Adam Wittek, Magdalena Kobielarz, Anju R. Babu, Martyn P. Nash, Poul M. F. Nielsen, and Karol Miller (eds.), *Computational Biomechanics for Medicine*, pp. 17–32, Cham, 2024. Springer Nature Switzerland. ISBN 978-3-031-64632-4.

Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 3789–3799, 2020.

Helmut Harbrecht, Reinhold Schneider, and Christoph Schwab. Sparse second moment analysis for elliptic problems in stochastic domains. *Numerische Mathematik*, 109(3):385–414, 2008. doi: 10.1007/s00211-008-0147-9.

James Harrison, John Willes, and Jasper Snoek. Variational Bayesian last layers. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Sx7BIiPzys.

Sidhanth Holalkere, David Bindel, Silvia Sellán, and Alexander Terenin. Stochastic Poisson surface reconstruction with one solve using geometric Gaussian processes. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 23382–23396. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/holalkere25a.html.

Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of Bayesian neural nets via local linearization. In Arindam Banerjee and Kenji Fukumizu (eds.), *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 703–711. PMLR, 13–15 Apr 2021. URL https://proceedings.mlr.press/v130/immer21a.html.

Samundra Karki, Mehdi Shadkhah, Cheng-Hau Yang, Aditya Balu, Guglielmo Scovazzi, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Direct flow simulations with implicit neural representation of complex geometry. *Computer Methods in Applied Mechanics and Engineering*, 446:118248, 2025.

Samundra Karki, Ming-Chen Hsu, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Mechanics simulation with implicit neural representations of complex geometries. *Computer-Aided Design*, 190:103978, 2026a. ISSN 0010-4485. doi: 10.1016/j.cad.2025.103978. URL https://www.sciencedirect.com/science/article/pii/S0010448525001393.

Samundra Karki, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Neural geometry for pdes: Regularity, stability, and convergence guarantees, 2026b. URL https://arxiv.org/abs/2602.02271.

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

Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2796–2804. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/kuleshov18a.html.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 6402–6413. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf.

David J. C. MacKay. A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3):448–472, 05 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.448. URL https://doi.org/10.1162/neco.1992.4.3.448.

Gabriel D. Maher, Casey M. Fleeter, Daniele E. Schiavazzi, and Alison L. Marsden. Geometric uncertainty in patient-specific cardiovascular modeling with convolutional dropout networks. *Computer Methods in Applied Mechanics and Engineering*, 386:114038, 2021. ISSN 0045-7825. doi: 10.1016/j.cma.2021.114038. URL https://www.sciencedirect.com/science/article/pii/S0045782521003698.

Alex Main and Guglielmo Scovazzi. The shifted boundary method for embedded domain computations. Part I: Poisson and Stokes problems. *Journal of Computational Physics*, 372:972–995, 2018.

Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3D reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4460–4470, 2019.

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 165–174, 2019. doi: 10.1109/CVPR.2019.00025.

E. Pujol and A. Chica. Bayesian 3D shape reconstruction from noisy points and normals. *Computer Graphics Forum*, 44(5):e70201, 2025. doi: 10.1111/cgf.70201. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.70201.

Carlos Riquelme, George Tucker, and Jasper Snoek. Deep Bayesian bandits showdown: An empirical comparison of Bayesian deep networks for Thompson sampling. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=SyYe6k-CW.

Silvia Sellán and Alec Jacobson. Stochastic Poisson surface reconstruction. *ACM Transactions on Graphics*, 41(6), November 2022. ISSN 0730-0301. doi: 10.1145/3550454.3555441. URL https://doi.org/10.1145/3550454.3555441.

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit Neural Representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7462–7473, 2020.

Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Prabhat, and Ryan Adams. Scalable Bayesian optimization using deep neural networks. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2171–2180, Lille, France, 07–09 Jul 2015. PMLR. URL https://proceedings.mlr.press/v37/snoek15.html.

Jan Sokołowski and Jean-Paul Zolésio. *Introduction to Shape Optimization*. Springer Series in Computational Mathematics. Springer, Berlin, Heidelberg, 1992. doi: 10.1007/978-3-642-58106-9.

A. M. Stuart. Inverse problems: A Bayesian perspective. *Acta Numerica*, 19:451–559, 2010. doi: 10.1017/S0962492910000061.

Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 7537–7547. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/55053683268957697aa39fba6f231c68-Paper.pdf.

Francisca Vasconcelos, Bobby He, Nalini M Singh, and Yee Whye Teh. UncertaINR: Uncertainty quantification of end-to-end implicit neural representations for computed tomography. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=jdGMBgYvfX.

Jan Verhülsdonk, Thomas Grandits, Francisco Sahli Costabal, Thomas Beiert, Simone Pezzuto, and Alexander Effland. Uncertainty quantification for cardiac shape reconstruction with deep signed distance functions via MCMC methods, 2026. URL https://arxiv.org/abs/2605.07987.

Oliver Williams and Andrew Fitzgibbon. Gaussian process implicit surfaces. In *Gaussian Processes in Practice*, April 2007. URL https://www.microsoft.com/en-us/research/publication/gaussian-process-implicit-surfaces-2/.

Dongbin Xiu and Daniel M. Tartakovsky. Numerical methods for differential equations in random domains. *SIAM Journal on Scientific Computing*, 28(3):1167–1185, 2006. doi: 10.1137/040613160. URL https://doi.org/10.1137/040613160.

Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19206–19216, 2022.

A MATHEMATICAL AND IMPLEMENTATION DETAILS

Deformation map. Assume $f_0, \psi \in C^2$ near Γ_0 and $\|\nabla f_0\| \geq g_{\min} > 0$ on a tube of half-width ρ around it. For $\mathbf{x} \in \Gamma_0$, $F_{\mathbf{x}}(t, \delta\boldsymbol{\beta}) = f_0(\mathbf{x} + t\mathbf{n}_0) + \boldsymbol{\psi}(\mathbf{x} + t\mathbf{n}_0)^\top \delta\boldsymbol{\beta}$ has $F_{\mathbf{x}}(0, \mathbf{0}) = 0$ and $\partial_t F_{\mathbf{x}}(0, \mathbf{0}) = \|\nabla f_0(\mathbf{x})\|$. For $\|\delta\boldsymbol{\beta}\|$ small enough that $\partial_t F_{\mathbf{x}} \geq g_{\min}/2$ on the tube, the implicit function theorem gives a unique root $t = h(\mathbf{x}; \delta\boldsymbol{\beta})$ in $(-\rho, \rho)$, uniformly in $\mathbf{x}$, and differentiating at $\delta\boldsymbol{\beta} = \mathbf{0}$ gives Eq. 5 with a remainder bounded by the second derivatives of f_0 and $\boldsymbol{\psi}$ on the tube. Along a ray $\mathbf{e}$ with $\mathbf{n}_0^\top \mathbf{e} > 0$ the same argument gives the radial displacement $-\boldsymbol{\psi}^\top \delta\boldsymbol{\beta} / (\nabla f_0^\top \mathbf{e})$, used for the radial modes.

Observation model, prior and posterior. Scan point i lies on ray θ_i with radius $\rho^\star(\theta_i) + \epsilon_i$. The radial noise covariance is

$$R_{ij} = \sigma^2 \big[(1 - \nu) \exp\big(-\Delta_{ij}^2 / 2\ell^2\big) + \nu\, \delta_{ij}\big], \qquad \Delta_{ij} = R_0 \min(|\theta_i - \theta_j|, 2\pi - |\theta_i - \theta_j|),$$

with $\sigma = 0.0025$ and nugget $\nu = 0.05$, and $\mathbf{B}$ holds the rows of $\boldsymbol{\Phi}$ at the 500 scan angles (a sub-grid of the 1000 boundary rays). For normal-distance observations, $\mathbf{B}$ instead evaluates the normal modes at the closest points and $\mathbf{R}$ projects $\mathbf{C}_{ij}$ on the normals. The prior is $\tau^2 \mathbf{I}$ with $\tau = 5\sigma$; Eq. 16 is evaluated with Cholesky factorizations of $\mathbf{R}$ and of the posterior precision. The rank rule and the discrepancy modes (Section 3) use the same SVD, computed on 1000 uniform rays with arc-length weights; the joint posterior of the $r + K$ coefficients is Eq. 16 with $\mathbf{B}$ extended by the discrepancy columns. The pointwise band is $\boldsymbol{\Phi}\boldsymbol{\mu}_c \pm z_{0.95}\, [\mathrm{diag}(\boldsymbol{\Phi}\boldsymbol{\Sigma}_c\boldsymbol{\Phi}^\top)]^{1/2}$ with $z_{0.95} = 1.645$ for the 5–95% band.

Networks and training. Every INR is a Cartesian coordinate MLP on the unit square with inputs $2\mathbf{x} - 1$, L hidden layers of width W and a linear head with bias, in float64. The SIRENs ($\sin(30\,\cdot)$ activations, initialization of Sitzmann et al. (2020)) are 128×2 (circle, flower; 17,025 parameters) and 128×3 (rectangle; 33,537); the tanh networks of the ablations are 64×3 (circles) and 128×3. Each INR is trained on its own scan: the target is the signed distance to the polygon through the noisy points at query points drawn uniformly and near the polygon (circle: 4096 queries, 50% near, band 0.025; flower and rectangle: 16,384, 70%, 0.012), plus the squared field at the scan points. Adam (Kingma & Ba, 2015) runs 20,000 steps with batch 256 (circle) or 512, learning rate 10^{-4} for SIREN and 10^{-3} for tanh, cosine-decayed to 5×10^{-5}; the final weights are used. The SIREN scans are drawn from the smooth field alone (noise standard deviation $\sigma\sqrt{1 - \nu}$).

PDE setup. The domain is the unit square with shapes centered at $(0.5, 0.5)$: a circle of radius 0.25, a flower $\rho^\star(\theta) = 0.25(1 + 0.18\cos 5\theta)$ and a 0.48×0.32 rectangle. The SBM in Nitsche form (Main & Scovazzi, 2018) uses a uniform level-8 quadtree ($h = 1/256$), bilinear elements with 2×2 Gauss quadrature, only fully interior cells (cell-inclusion parameter $\lambda = 0$), penalty $\alpha = 10$ and closest points by Newton projection on the implicit field. The nonsymmetric system is solved by sparse LU (relative residual 10^{-10}), and the factorization is reused for the adjoint. Probes are 16 equispaced points on $\theta = 0$ from the center to $0.85\rho^\star(0)$; $u^\star$ uses a 4096-vertex polygon of the true shape, and each ground-truth curve is solved as a 4000-vertex polygon.

B VALIDATION DETAILS

Reference distributions. *Exact geometry*: $\rho_k(\theta) = \rho^\star(\theta) + \varepsilon_k(\theta)$, $k = 1, \ldots, 500$, with ε_k a Gaussian random field of covariance $(1 - \nu)\sigma^2 \exp(-s^2/2\ell^2)$ drawn exactly on 4000 rays with fixed seeds. *Controlled*: 500 draws of the retained-mode posterior (seeds independent of the Monte Carlo draws), pushed into the last layer and extracted on 4000 rays; all 1500 are single closed curves. *Retrained ensembles* (ablations only), following the deep-ensemble approach (Lakshminarayanan et al., 2017): 500 tanh INRs per campaign, each trained on its own scan; they describe variability of the reconstruction procedure, not samples of the posterior.

Metrics. Boundaries are radii $\boldsymbol{\rho} \in \mathbb{R}^{1000}$ on common rays. For a population with covariance C, the weighted covariance is $\mathbf{W}^{1/2} C\, \mathbf{W}^{1/2}$ (arc-length weights). The variance ratio is $\mathrm{tr}(\mathbf{W}C_{\mathrm{post}}) / \mathrm{tr}(\mathbf{W}C_{\mathrm{gt}})$; the covariance difference is the Frobenius norm of the difference of the trace-normalized weighted covariances; the area ratio is $\mathrm{std}(A_{\mathrm{post}}) / \mathrm{std}(A_{\mathrm{gt}})$ (unweighted). The baseline is the same difference between two random halves of the ground truth. Geometry metrics of the posterior with discrepancy modes use 1000 linear draws $\rho_0 + \boldsymbol{\Phi}\mathbf{c}$; the others use realized draws.

Recovered displacement. For each ground-truth curve $\boldsymbol{\rho}_k$ (exact or controlled), the target displacement from the unedited boundary $\boldsymbol{\rho}_0$ is $\mathbf{d}_k = \boldsymbol{\rho}_k - \boldsymbol{\rho}_0$, and the retained modes are fitted by weighted least squares, $\mathbf{c}_k = \arg\min_{\mathbf{c}} \|\boldsymbol{\Phi}\mathbf{c} - \mathbf{d}_k\|_{\mathbf{W}}$. The fitted head $\boldsymbol{\beta}_0 + \mathbf{T}\mathbf{c}_k$ is pushed into the last layer and its zero set extracted on the same rays, giving the realized radii $\tilde{\boldsymbol{\rho}}_k$. The recovered fraction is

$$\eta_k = 1 - \frac{\|\tilde{\boldsymbol{\rho}}_k - \boldsymbol{\rho}_k\|_{\mathbf{W}}}{\|\boldsymbol{\rho}_k - \boldsymbol{\rho}_0\|_{\mathbf{W}}},$$

so $\eta_k = 1$ for an exact fit and $\eta_k = 0$ when the edit leaves the error of the unedited network unchanged. We report the median of η_k over curves whose fitted zero set is a single closed curve. The *linear* recovery replaces $\tilde{\boldsymbol{\rho}}_k$ by $\boldsymbol{\rho}_0 + \boldsymbol{\Phi}\mathbf{c}_k$.

Coverage. Each coverage number states its nominal level, in the sense of calibration for probabilistic forecasts (Gneiting & Raftery, 2007; Kuleshov et al., 2018). *Boundary coverage* (Table 1): share of the 1000 rays at which the true boundary (exact) or each posterior-drawn curve, averaged over curves (controlled), lies inside the linear 5–95% band; nominal 0.90. *PDE, exact*: number of the 16 probes whose linear 5–95% interval contains $u^\star$. *PDE, controlled*: share of the 500×16 ground-truth solutions inside the linear 5–95% interval; nominal 0.90. *Repeated scans* (tanh, Table 3): each of 500 INRs builds its own posterior; boundary coverage as above (nominal 0.90), and PDE coverage as the share of probe–scan pairs whose linear 95% interval contains $u^\star$ (nominal 0.95).

Retrained tanh ensembles. Table 3 summarizes the comparison of a single tanh INR's posterior with 500 retrained INRs (medians over nine *reference networks*, individually trained INRs selected by validation loss from among the retrained INRs), and repeated-scan coverage. Agreement is close for the circle with $\ell/R_0 = 0.2$ and the flower, weaker for the circle with iid (independent, uncorrelated) noise and poor for the rectangle, where the retained modes capture about 70% of the ensemble variance and the posterior 85% of that. A truth-free widening by $\sqrt{\max(1, \chi^2/\mathrm{dof})}$ of the scan residuals raises the rectangle's boundary coverage from 0.54 to 0.85. Linear propagation matches posterior Monte Carlo (standard-deviation ratios 0.94–0.99), but near the rectangle's wall a corner bias shared by the retrained INRs places $u^\star$ outside the posterior.

Table 3: **Tanh networks: posterior against retrained ensembles, and repeated-scan coverage.** Ratios are posterior over ensemble (ideal 1); baseline in brackets. Boundary coverage: linear 5–95% band, nominal 0.90; PDE coverage: linear 95% interval over probe–scan pairs, nominal 0.95 (not run for the iid circle).

	Circle, iid	Circle, $\ell/R_0 = 0.2$	Flower	Rectangle
Variance ratio	0.61	1.07	0.82	0.58
Covariance difference	0.65 (0.28)	0.28 (0.34)	0.41 (0.37)	0.85 (0.31)
Area s.d. ratio	0.45	1.01	0.90	0.88
Boundary coverage (nominal 0.90)	0.87	0.87	0.83	0.54
PDE coverage (nominal 0.95)	–	0.945	0.915	0.821

C FOCUSED ABLATIONS

Table 4 varies the number of discrepancy modes and the scan and likelihood model for the SIREN of each shape; Table 5 shows the trade-off between rank and valid realized geometry for the tanh rectangle. The variance ratios in Table 4 use the analytic linear band, and differ from the sampled values of Table 1 by at most 0.01. Replacing the discrepancy modes by smooth Fourier modes orthogonalized against the retained modes gives similar coverage (for $K = 20$: 0.91, 0.86, 0.86) with more over-coverage in the controlled test (0.93–0.95).

Table 4: **Discrepancy modes and scan model** (SIREN). Top: K discrepancy modes; boundary coverage of the truth and of posterior-drawn curves (nominal 0.90), variance ratio against exact geometry and PDE probes (of 16) whose 5–95% interval contains $u^\star$. Bottom ($K = 0$): training scan with or without the independent nugget, likelihood with or without the 5% nugget; PDE spread is the linear standard deviation over the exact-geometry ensemble's (median over probes).

Shape	K	coverage	controlled	var. ratio	$u^\star$ in 5–95%
Circle	0 / 5 / 10	0.87 / 0.90 / 0.90	0.90 / 0.91 / 0.93	0.72 / 0.79 / 0.88	16 / 16 / 16
Flower	0 / 5 / 10	0.81 / 0.83 / 0.83	0.90 / 0.91 / 0.92	0.72 / 0.76 / 0.81	14 / 15 / 16
Rectangle	0 / 5 / 10	0.84 / 0.83 / 0.87	0.91 / 0.92 / 0.93	0.58 / 0.63 / 0.68	16 / 16 / 16

Scan nugget / likelihood nugget ($K = 0$)	coverage circle / flower / rect.	$u^\star$ in 5–95% circle / flower / rect.	PDE spread circle / flower / rect.
yes / 5%	0.87 / 0.78 / 0.75	16 / 8 / 16	0.81 / 0.84 / 0.90
no / 5%	0.87 / 0.81 / 0.84	16 / 14 / 16	0.90 / 0.88 / 0.92
no / none	0.14 / 0.15 / 0.00	4 / 3 / 0	0.13 / 0.34 / 0.10

Table 5: **Rank against valid realized geometry** (tanh rectangle, medians over nine references, 1000 draws per rank). Span: share of the retrained-ensemble variance captured by the retained modes; linear: linear posterior variance over ensemble variance.

Rank r	9	11	13	15	17
Retained span	0.49	0.60	0.69	0.75	0.87
Linear variance ratio	0.44	0.52	0.59	0.66	0.72
Invalid draws, median	0%	0%	0%	0.4%	98%
References with invalid draws	0/9	0/9	2/9	7/9	9/9

Architecture. Across tanh and SIREN networks of widths 32–256 and depths 2–5 retrained on the same scan, network size changes the recovered share of the exact curves little; the SIREN's usable rank (32 against 13–18) raises it. Width 256 at depths 4–5 fails to give a single closed zero set, and SIRENs of depth 4–8 with a skip connection recover less and cover less than the reported networks.