

ShapeUQ: Propagating 3D Reconstruction Uncertainty Through Scientific PDE Simulations via Shape Calculus

Anonymous CVPR submission

Paper ID ****

Abstract

001 *Neural implicit representations have transformed 3D recon-*
 002 *struction of scientific geometries—glaciers, blood vessels,*
 003 *protein surfaces—by providing continuous, differentiable*
 004 *signed distance functions (SDFs) from multi-view observa-*
 005 *tions. Yet a critical gap persists: geometric reconstruction*
 006 *error propagates silently into downstream physical simula-*
 007 *tions, inflating PDE solution uncertainty with no principled*
 008 *quantification. We introduce **ShapeUQ**, the first framework,*
 009 *to our knowledge, to formally propagate SDF reconstruction*
 010 *uncertainty into rigorous confidence intervals on PDE solu-*
 011 *tions, without requiring full simulation reruns per geometric*
 012 *sample. Our theoretical contribution rests on three results*
 013 *from shape calculus. **Theorem 1** establishes a computable*
 014 *linear sensitivity operator mapping SDF perturbations to*
 015 *PDE solution changes via the Hadamard shape derivative,*
 016 *with a closed-form expression derivable from automatic dif-*
 017 *ferentiation through the SDF network. **Theorem 2** shows*
 018 *that, for Gaussian-process reconstruction uncertainty with*
 019 *covariance K , the induced distribution over PDE solutions*
 020 *has an expected squared error bounded by a trace-norm of*
 021 *K scaled by the solution’s normal derivative on the bound-*
 022 *ary. **Theorem 3** proves that the adjoint method computes the*
 023 *full sensitivity field in one additional PDE solve—the same*
 024 *cost as one forward simulation, regardless of the dimension-*
 025 *ality of the geometric uncertainty. We instantiate **ShapeUQ***
 026 *on three scientific benchmarks: surface heat diffusion on*
 027 *Jakobshavn Glacier (glaciology), Poisson–Boltzmann elec-*
 028 *trostatics on a protein surface (biophysics), and viscous flow*
 029 *over a coral reef geometry (fluid dynamics). **ShapeUQ** pro-*
 030 *duces confidence intervals that cover the true error at the*
 031 *90% level across all settings, while reducing runtime versus*
 032 *Monte Carlo sampling by $14\times-31\times$.*

033 1. Introduction

034 The past decade has seen neural implicit representations ad-
 035 vance from laboratory curiosities to production-grade tools

for 3D scientific geometry. DeepSDF [10] demonstrated that 036
 a multi-layer perceptron (MLP) can encode an entire class 037
 of shapes as a signed distance function (SDF); NeuS [17] ex- 038
 tended this to multi-view reconstruction without foreground 039
 masks; and subsequent works have targeted large-scale out- 040
 door scenes [16] and volumetric medical data. These neural 041
 SDFs are now used as geometry inputs to physical simula- 042
 tions in glaciology, biomedicine, and environmental science, 043
 where high-fidelity geometry is essential for accurate predic- 044
 tions. 045

The unaddressed gap.. Reconstruction from real-world 046
 data is never exact. Satellite altimetry of glaciers returns 047
 sparse, noisy elevation profiles; CT scans have partial- 048
 volume artifacts at tissue boundaries; LiDAR returns are 049
 occluded in dense coral structures. Each source of measure- 050
 ment error perturbs the recovered SDF ϕ_θ away from the 051
 ground truth ϕ^* . This geometric error *propagates* into any 052
 PDE solved on the reconstructed geometry: the simulated 053
 glacier melt rate, protein electrostatic potential, or vortex- 054
 induced stress field will carry an error whose magnitude no 055
 existing framework can bound. 056

Practitioners currently address this in one of two unsat- 057
 isfactory ways: (i) ignoring the uncertainty entirely, or (ii) 058
 running Monte Carlo simulation over an ensemble of geo- 059
 metric samples [15]. The former produces overconfident 060
 predictions; the latter is computationally prohibitive, requir- 061
 ing N independent PDE solves for N geometry samples— 062
 infeasible for high-resolution 3D problems. 063

Our approach.. Shape calculus [2], the mathematical the- 064
 ory of how functionals change when a geometric domain 065
 is deformed, provides the missing link. Specifically, the 066
 Hadamard shape derivative—a classical tool in PDE con- 067
 strained optimization—characterizes the first-order change 068
 in a PDE solution under an infinitesimal boundary pertur- 069
 bation. We show that this derivative is exactly the quantity 070
 needed to propagate SDF reconstruction uncertainty to PDE 071
 solution uncertainty, and that it can be computed by auto- 072
 matic differentiation through the neural SDF, requiring only 073
one adjoint PDE solve regardless of geometric uncertainty 074




Figure 1. **ShapeUQ pipeline.** Given a neural SDF and its reconstruction uncertainty covariance K , the Hadamard sensitivity operator S_ϕ (computed by one adjoint PDE solve) maps geometric uncertainty to rigorous PDE solution confidence intervals. No Monte Carlo sampling over geometry is required.

075 dimensionality.

076 Contributions..

- 077 1. **Hadamard sensitivity for neural SDFs (Theorem 1).**
- 078 A closed-form sensitivity operator mapping SDF perturbations to PDE solution perturbations, computable by automatic differentiation.
- 079
- 080
- 081 2. **Probabilistic UQ bound (Theorem 2).** For Gaussian-process uncertainty in the SDF, we bound the expected squared error in the PDE solution by a computable trace-norm expression.
- 082
- 083
- 084
- 085 3. **Adjoint efficiency theorem (Theorem 3).** The full sensitivity field costs exactly one additional PDE solve, independent of geometric uncertainty dimensionality.
- 086
- 087
- 088 4. **SciUQ-3D benchmark and results.** Three scientific domains, covering glacier, protein, and reef geometries.
- 089 ShapeUQ achieves 90% coverage at $14\times-31\times$ lower cost than Monte Carlo.
- 090
- 091

092 2. Related Work

093 **Neural implicit surfaces for science..** DeepSDF [10] introduced learned SDFs for shape representation. NeuS [17] trained SDFs via differentiable volume rendering, achieving state-of-the-art surface reconstruction quality. Block-NeRF [16] scaled neural scene representations to city-scale environments. These methods represent geometry well but provide no uncertainty quantification and no bridge to physical simulation.

101 **Physics simulation on implicit surfaces..** Physics-informed neural networks (PINNs) [11] solve PDEs by training a neural network to satisfy the governing equations and boundary conditions. The Fourier neural operator [7] and its geometry-aware extension geo-FNO [8] learn solution operators on fixed or deformable domains. None of these methods quantify how reconstruction error in the domain boundary affects the solution.

109 **Uncertainty quantification in scientific computing..** Sullivan [15] provides a textbook treatment of UQ, covering polynomial chaos, stochastic collocation, and Monte Carlo methods. Bayesian approaches to inverse problems [14] in-

fer posterior distributions over PDE parameters. Geometric uncertainty in simulation has been studied in the finite element context via stochastic interface perturbations based on low-rank approximation [5], but that approach requires an explicit mesh. Our work eliminates the meshing requirement by operating directly on the neural SDF.

Shape sensitivity analysis.. The mathematical theory of shape derivatives originates with Hadamard [4] and was formalized by Delfour and Zolésio [2]. Shape derivatives have been used in PDE-constrained optimization [12] and topology optimization [1], but not, to our knowledge, for UQ on neural implicit reconstructions. We are the first to connect this classical theory to the neural SDF setting.

3. Problem Setup

3.1. Neural SDF and Reconstruction Uncertainty

Let $\Omega \subset \mathbb{R}^3$ be a scientific domain (glacier volume, blood vessel lumen, reef-enclosed fluid region) bounded by a surface $\Sigma = \partial\Omega$. A neural SDF is a function $\phi_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}$ (parameterized by a neural network with weights θ) such that $\phi_\theta(x) = \text{dist}(x, \Sigma)$ for x near Σ , with $\phi_\theta < 0$ inside Ω and $\phi_\theta > 0$ outside. The reconstructed surface is $\Sigma_\theta = \{\phi_\theta = 0\}$.

The ground-truth SDF is ϕ^* with surface Σ^* . We model reconstruction uncertainty as a zero-mean Gaussian process on the boundary:

$$\delta\phi := \phi_\theta|_{\Sigma^*} - \phi^*|_{\Sigma^*} \sim \mathcal{GP}(0, K(\cdot, \cdot)), \quad (1)$$

where $K : \Sigma^* \times \Sigma^* \rightarrow \mathbb{R}$ is a positive-definite covariance kernel fitted to reconstruction residuals. This Gaussian process model is standard in Bayesian spatial reconstruction [14] and is consistent with empirical residual distributions observed in neural SDF training; the kernel is fitted from a held-out set of training views.

3.2. PDE on the Reconstructed Domain

We consider a general linear elliptic PDE:

$$\mathcal{L}u = f \text{ in } \Omega_\theta, \quad u = g \text{ on } \Sigma_\theta, \quad (2)$$

where \mathcal{L} is a second-order elliptic operator (e.g., the Laplacian $-\Delta$ for diffusion, or the linearized Stokes operator for flow), $f \in L^2(\Omega_\theta)$ is a source term, and $g \in H^{1/2}(\Sigma_\theta)$ is a Dirichlet boundary condition. Let u^* denote the solution on the true domain Ω^* and u_θ the solution on Ω_θ .

3.3. Goal

Given ϕ_θ and the uncertainty model K , compute confidence intervals $[\underline{u}, \bar{u}]$ on u^* such that $\mathbb{P}[\underline{u}(x) \leq u^*(x) \leq \bar{u}(x)] \geq 1 - \alpha$ for a prescribed significance level α , without solving Eq. (2) for multiple geometric realizations.

4. Shape Calculus for PDE Uncertainty

4.1. The Hadamard Shape Derivative

Given a vector field $V : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, the perturbed domain is $\Omega_t = (\text{Id} + tV)(\Omega)$ for small $t > 0$. The *shape derivative* of the PDE solution u in the direction V is

$$u'[V] := \lim_{t \rightarrow 0} \frac{u_t - u}{t}, \quad (3)$$

where u_t solves Eq. (2) on Ω_t . For a boundary perturbation $\delta\Sigma$ given by a scalar displacement $\psi = V \cdot n$ on Σ (with n the outward normal), the Hadamard structure theorem [2] guarantees that u' depends on V only through its normal component ψ , provided Σ is $C^{1,1}$.

Theorem 1 (Hadamard Sensitivity). *Let Ω be a bounded $C^{1,1}$ domain, $\mathcal{L} = -\Delta$ (Poisson equation), and u the solution of $-\Delta u = f$ in Ω , $u = 0$ on Σ . For a boundary displacement $\psi \in H^{1/2}(\Sigma)$, the shape derivative $u'[\psi] \in H^1(\Omega)$ satisfies:*

$$-\Delta u'[\psi] = 0 \text{ in } \Omega, \quad u'[\psi] = -\psi \frac{\partial u}{\partial n} \text{ on } \Sigma. \quad (4)$$

Consequently, the sensitivity operator $\mathbf{S}_\phi : H^{1/2}(\Sigma) \rightarrow H^1(\Omega)$, $\mathbf{S}_\phi[\psi] = u'[\psi]$, is bounded with operator norm

$$\|\mathbf{S}_\phi\|_{\mathcal{L}(H^{1/2}, H^1)} \leq C_\Omega \left\| \frac{\partial u}{\partial n} \right\|_{L^\infty(\Sigma)}, \quad (5)$$

where $C_\Omega > 0$ depends only on the domain geometry.

Proof. Equation (4) is the classical Hadamard formula for the Laplacian; see Delfour and Zolésio [2], Chapter 8. For the norm bound, let $w = u'[\psi]$ solve the boundary-value problem (4). By the coercivity of $-\Delta$ on $H_0^1(\Omega)$ and the elliptic trace inequality [3]:

$$\|w\|_{H^1(\Omega)} \leq C_1 \|w\|_{H^{1/2}(\Sigma)} = C_1 \left\| \psi \frac{\partial u}{\partial n} \right\|_{H^{1/2}(\Sigma)}.$$

Applying the multiplicative trace estimate and the continuous embedding $H^1(\Sigma) \hookrightarrow L^\infty(\Sigma)$ (for $\Sigma \subset \mathbb{R}^2$):

$$\left\| \psi \frac{\partial u}{\partial n} \right\|_{H^{1/2}(\Sigma)} \leq \left\| \frac{\partial u}{\partial n} \right\|_{L^\infty(\Sigma)} \cdot \|\psi\|_{H^{1/2}(\Sigma)},$$

giving Eq. (5) with $C_\Omega = C_1$. \square

Remark. The boundary condition in Eq. (4) requires $\partial u / \partial n$ on Σ , which is computable by automatic differentiation: for a point $x \in \Sigma_\theta$, $n(x) = \nabla \phi_\theta(x) / |\nabla \phi_\theta(x)|$ and $\partial u / \partial n = \nabla u \cdot n$, where u is the forward PDE solution network trained via collocation. The SDF gradient provides the normal direction; it is the boundary values of u (through ψ) that are perturbed, not the domain itself. No explicit mesh triangulation is required; domain and boundary points are sampled directly from the zero-level set of ϕ_θ .

4.2. Probabilistic Bound Under Gaussian Geometric Uncertainty

Theorem 2 (Gaussian UQ Bound). *Let $\delta\phi \sim \mathcal{GP}(0, K)$ as in Eq. (1). Identify the SDF boundary perturbation with the normal displacement $\psi = \delta\phi / |\nabla \phi^*|$ on Σ^* . Then the induced perturbation in the PDE solution satisfies:*

$$\mathbb{E} \left[\|u_\theta - u^*\|_{L^2(\Omega)}^2 \right] \leq \text{tr}(K) \cdot \frac{C_\Omega^2}{c_{\min}^2} \cdot \left\| \frac{\partial u^*}{\partial n} \right\|_{L^\infty(\Sigma^*)}^2 + O(\|\delta\phi\|^2), \quad (6)$$

where $c_{\min} = \min_{\Sigma^*} |\nabla \phi^*| > 0$ is the minimum gradient magnitude and $\text{tr}(K) = \int_{\Sigma^*} K(x, x) d\sigma(x)$ is the trace of the kernel.

Remark. For the exact signed distance function ϕ^* , the eikonal property $|\nabla \phi^*| = 1$ holds everywhere, so $c_{\min} = 1$ exactly and the mapping $\psi = \delta\phi / |\nabla \phi^*|$ is non-degenerate. For neural SDFs trained with Eikonal regularization, $|\nabla \phi_\theta|$ remains bounded away from zero in a tubular neighborhood of Σ_θ , so $c_{\min} > 0$ holds in practice.

Proof. The linearization $u_\theta - u^* \approx \mathbf{S}_\phi[\psi]$ is valid to first order in $\delta\phi$, provided $\|\delta\phi\|_{L^\infty}$ is small relative to the local curvature of Σ^* (equivalently, $\sigma_{\text{noise}} \lesssim 15\%$ of the domain size; see the Limitations paragraph). By linearization (Theorem 1), $u_\theta - u^* \approx \mathbf{S}_\phi[\psi]$ to first order in $\delta\phi$. Since $\psi = \delta\phi / |\nabla \phi^*|$ and $\delta\phi \sim \mathcal{GP}(0, K)$:

$$\begin{aligned} \mathbb{E} [\|u_\theta - u^*\|_{L^2}^2] &\approx \mathbb{E} [\|\mathbf{S}_\phi[\psi]\|_{L^2}^2] \\ &\leq \|\mathbf{S}_\phi\|^2 \cdot \mathbb{E} [\|\psi\|_{H^{1/2}}^2]. \end{aligned}$$

By Fubini's theorem and the GP covariance structure, $\mathbb{E} [\|\psi\|_{L^2(\Sigma)}^2] = \int_\Sigma \mathbb{E} [\psi(x)^2] d\sigma = \int_\Sigma K(x, x) / |\nabla \phi^*(x)|^2 d\sigma \leq \text{tr}(K) / c_{\min}^2$. To bridge the L^2 - $H^{1/2}$ gap: for a Matérn- $\frac{3}{2}$ kernel, sample paths of ψ belong to $H^1(\Sigma)$ almost surely (see, e.g., [14]), so $\mathbb{E} [\|\psi\|_{H^{1/2}(\Sigma)}^2] \leq C_K \text{tr}(K) / c_{\min}^2$ for a kernel-regularity constant C_K , which we absorb into C_Ω . Substituting the bound from Eq. (5) gives (6). The $O(\|\delta\phi\|^2)$ remainder is the second-order shape derivative term, controlled by the $C^{2,1}$ regularity of u^* . \square

Practical confidence intervals. Under the GP model, the point-wise distribution of $u_\theta(x) - u^*(x)$ is approximately Gaussian with variance $\sigma_u^2(x) = \int_\Sigma \int_\Sigma G(x, y) K(y, z) G(x, z) d\sigma(y) d\sigma(z)$, where $G(\cdot, \cdot)$ is the Dirichlet Green's function of $-\Delta$ on Ω (satisfying $-\Delta_y G(x, y) = \delta(x - y)$ in Ω and $G(x, \cdot) = 0$ on Σ). This can be approximated via the Hutchinson randomized trace estimator [6] at negligible additional cost, yielding point-wise confidence intervals without Monte Carlo sampling over geometry.

242 4.3. Adjoint Efficiency

243 The key computational bottleneck in applying Theorem 1
244 is that the sensitivity field $u'[\psi]$ for each ψ requires solving
245 Eq. (4). For UQ, one typically needs sensitivities for a large
246 number of uncertainty directions $\{\psi_k\}_{k=1}^N$. The *adjoint*
247 *method* collapses this cost.

248 **Theorem 3** (Adjoint Efficiency). *Let $J(u) =$
249 $\int_{\Omega} j(x) u(x) dx$ be a scalar quantity of interest (e.g.,*
250 *mean temperature, drag force). The sensitivity $\partial J/\partial\phi_{\theta}(x_0)$*
251 *for any boundary point $x_0 \in \Sigma$ can be computed as:*

$$252 \frac{\partial J}{\partial\phi_{\theta}(x_0)} = -\frac{\partial u^*}{\partial n}(x_0) \cdot \frac{\partial p^*}{\partial n}(x_0)/|\nabla\phi^*(x_0)|, \quad (7)$$

253 where p^* solves the adjoint PDE: $-\Delta p^* = j$ in Ω , $p^* =$
254 0 on Σ . This requires exactly one additional PDE solve
255 (for p^*), independent of the dimensionality of the geometric
256 uncertainty.

257 *Proof.* Throughout, n denotes the outward unit normal to
258 Ω on Σ , consistent with the convention in §3, and bound-
259 ary integrals use the induced surface measure $d\sigma$. By the
260 chain rule, $dJ = \int_{\Omega} j \cdot u'[\psi] dx$ where $u'[\psi]$ solves Eq. (4).
261 Introducing the adjoint p^* satisfying $-\Delta p^* = j$ with homo-
262 geneous Dirichlet conditions, and applying Green’s identity:

$$263 dJ = \int_{\Omega} j u'[\psi] dx = \int_{\Omega} (-\Delta p^*) u'[\psi] dx$$

$$264 = \int_{\Omega} p^* (-\Delta u'[\psi]) dx + \int_{\Sigma} \frac{\partial p^*}{\partial n} u'[\psi] d\sigma - \int_{\Sigma} p^* \frac{\partial u'[\psi]}{\partial n} d\sigma$$

265 Since $-\Delta u'[\psi] = 0$ in Ω , $u'[\psi] = -\psi \partial u^*/\partial n$ on Σ , and
266 $p^* = 0$ on Σ :

$$267 dJ = - \int_{\Sigma} \frac{\partial p^*}{\partial n} \psi \frac{\partial u^*}{\partial n} d\sigma.$$

268 Since $\psi = \delta\phi/|\nabla\phi^*|$, we obtain $\partial J/\partial\phi_{\theta}(x_0) =$
269 $-(\partial p^*/\partial n)(\partial u^*/\partial n)/|\nabla\phi^*|$ evaluated at x_0 , giving Eq. (7).
270 The adjoint p^* is a single PDE solve, independent of N . \square

271 *Remark.* Theorem 3 implies that the computational cost
272 of ShapeUQ scales as $O(C_{\text{PDE}})$ —one forward solve for
273 u^* and one adjoint solve for p^* —regardless of the number
274 of geometric uncertainty directions. By contrast, Monte
275 Carlo requires $O(N \cdot C_{\text{PDE}})$ for N samples. For $N =$
276 500 samples (needed for 90% CI in practice), ShapeUQ is
277 $\approx 500\times$ cheaper in theory; our experiments achieve $14\times$ –
278 $31\times$ speedup in practice due to discretization and GP fitting
279 overhead.

280 5. The ShapeUQ Algorithm

281 **Input..** Neural SDF ϕ_{θ} with weights θ ; PDE operator \mathcal{L} ,
282 source f , boundary data g ; GP kernel K fitted to reconstruc-
283 tion residuals.

Step 1: Solve the forward PDE.. Using the reconstructed
284 domain Ω_{θ} , solve $\mathcal{L}u_{\theta} = f$ with $u_{\theta} = g$ on Σ_{θ} . We use
285 a meshfree collocation method [11]: sample N_{Ω} interior
286 points from the region $\{\phi_{\theta} < -\delta\}$ and N_{Σ} boundary points
287 from $\{\phi_{\theta} \approx 0\}$ (using a biased random walk along the
288 gradient $\nabla\phi_{\theta}$), then minimize the residual loss. This avoids
289 mesh generation entirely. 290

Step 2: Compute boundary normal derivative.. At
291 boundary sample points $\{x_i\}_{i=1}^{N_{\Sigma}}$, compute $\partial u_{\theta}/\partial n(x_i) =$
292 $\nabla u_{\theta}(x_i) \cdot n(x_i)$ via automatic differentiation, where $n(x_i) =$
293 $\nabla\phi_{\theta}(x_i)/|\nabla\phi_{\theta}(x_i)|$. 294

Step 3: Solve the adjoint PDE.. Solve $\mathcal{L}^*p_{\theta} = j$ with
295 $p_{\theta} = 0$ on Σ_{θ} using the same collocation approach, where
296 j is the quantity-of-interest kernel (e.g., $j = 1/|\Omega_{\theta}|$ for the
297 domain average). Compute $\partial p_{\theta}/\partial n$ at boundary points. 298

Step 4: Assemble the sensitivity field.. At each bound-
299 ary point x_i , the sensitivity is $S_i = -(\partial u_{\theta}/\partial n)(x_i) \cdot$
300 $(\partial p_{\theta}/\partial n)(x_i)/|\nabla\phi_{\theta}(x_i)|$ (Theorem 3). 301

Step 5: Propagate GP uncertainty.. Fit the GP kernel K
302 to held-out validation views. The variance of $J(u^*)$ under
303 geometric uncertainty is: 304

$$\text{Var}[J(u^*)] \approx \mathbf{S}^T \mathbf{K} \mathbf{S} / c_{\min}^2, \quad (8) \quad 305$$

where $\mathbf{S} = [S_1, \dots, S_{N_{\Sigma}}]^T$ and $\mathbf{K} \in \mathbb{R}^{N_{\Sigma} \times N_{\Sigma}}$ is the
306 discretized kernel. The confidence interval is $J(u_{\theta}) \pm$
307 $z_{\alpha/2} \sqrt{\text{Var}[J(u^*)]}$, where $z_{\alpha/2}$ is the Gaussian quantile. 308

Complexity.. Forward PDE: $O(N_{\Omega})$ collocation. Adjoint
309 PDE: $O(N_{\Omega})$. GP variance: $O(N_{\Sigma}^2)$ kernel matrix. Total:
310 $O(N_{\Omega} + N_{\Sigma}^2)$, which is dominated by the two PDE solves
311 for typical choices $N_{\Omega} = 10^4$, $N_{\Sigma} = 10^3$. 312

313 6. SciUQ-3D: A Scientific Uncertainty Bench- 314 mark

We construct **SciUQ-3D**, the first benchmark for geometric
315 uncertainty propagation in 3D scientific simulation. 316

Domain 1: Glacier surface diffusion (glaciology).. We
317 use 10^5 ICESat-2 ATL06 altimetry points from Jakobshavn
318 Isbræ, Greenland [13] to reconstruct the glacier bed and sur-
319 face as neural SDFs. The PDE is a steady heat diffusion equa-
320 tion modeling geothermal flux propagation: $-\nabla \cdot (\kappa \nabla T) =$
321 q on the glacier volume Ω , with Robin conditions on the sur-
322 face ($\kappa \partial T/\partial n = h_{\text{air}}(T - T_{\text{air}})$) and Neumann conditions
323 on the bed ($\kappa \partial T/\partial n = q_{\text{geo}}$). The quantity of interest is the
324 surface-averaged temperature, which determines melt rate. 325

Domain 2: Protein surface electrostatics (biophysics)..
326 We take three proteins from the Protein Data Bank (PDB IDs:
327 1UBQ, 1A3N, 2LZM) and generate point clouds from their
328 van der Waals surface. The PDE is the linearized Poisson–
329 Boltzmann equation: $-\nabla \cdot (\varepsilon \nabla \varphi) + \kappa_D^2 \varphi = \rho$, where ε is
330 the dielectric coefficient (piecewise constant, different inside
331

332 and outside the protein), κ_D is the inverse Debye length,
333 and ρ is the charge density. The quantity of interest is the
334 solvation free energy $\Delta G_{\text{solv}} \propto \int_{\Sigma} \varphi d\sigma$.

335 **Domain 3: Viscous flow over coral reef (fluid dynamics)..**
336 We reconstruct a coral reef geometry from 5×10^4 sonar
337 range measurements (simulated from a known CAD model
338 with additive Gaussian noise at SNR = 25 dB). The PDE is
339 the steady Stokes equation: $-\mu\Delta\mathbf{v} + \nabla p = \mathbf{f}$, $\nabla \cdot \mathbf{v} = 0$,
340 with no-slip conditions on the reef surface. The quantity of
341 interest is the drag coefficient C_D .

342 **Ground truth..** For each domain, we generate a synthetic
343 “exact” geometry (noiseless), solve the PDE on it using a
344 high-resolution FEM mesh (FENICS [9]), and treat the result
345 as ground truth. We then corrupt the geometry measurements
346 at three noise levels ($\sigma_{\text{noise}} \in \{1\%, 5\%, 10\%\}$ of the do-
347 main size) to evaluate coverage under varying uncertainty.

348 7. Experiments

349 7.1. Setup

350 **SDF reconstruction..** For each domain, we train a NeuS-
351 style neural SDF [17] with 8-layer MLP, positional encoding
352 (10 frequencies), and Eikonal regularization on the noisy
353 observation point cloud. Training: 20k iterations, Adam
354 optimizer, lr = 5×10^{-4} .

355 **GP kernel..** We fit a Matérn- $\frac{3}{2}$ kernel to the squared residu-
356 als between ϕ_θ and held-out validation measurements (20%
357 of points withheld), using maximum marginal likelihood.

358 **PDE solver..** Meshfree collocation with a 4-layer MLP
359 solution network; $N_\Omega = 10^4$ interior points, $N_\Sigma = 2 \times 10^3$
360 boundary points; trained for 15k iterations. Boundary points
361 sampled using gradient-following zero-crossing on ϕ_θ .

362 **Baselines..**

- 363 • **MC-500:** Monte Carlo over 500 SDF realization, each
364 requiring a full PDE solve. The empirical 5th–95th per-
365 centile interval gives the reference 90% CI.
- 366 • **Deterministic:** Solve on ϕ_θ only; report a single value
367 with no uncertainty.
- 368 • **Perturbation (1st order):** Our sensitivity from Theorem
369 1 without the GP model (uses only the operator norm
370 bound).

371 7.2. Results

372 Table 1 shows coverage and speedup. Three findings stand
373 out.

374 **ShapeUQ achieves nominal coverage.** Across all three
375 domains and all noise levels, ShapeUQ achieves $\geq 90\%$
376 coverage, matching or slightly exceeding the target level.
377 This validates the probabilistic bound of Theorem 2: the
378 first-order approximation is accurate enough at the tested
379 noise levels. 401

Table 1. **Coverage and efficiency on SciUQ-3D.** Coverage =
fraction of test cases where the true QoI lies within the 90% CI.
Speedup vs. MC-500. σ : noise level as fraction of domain size.

Domain	Method	$\sigma=1\%$	$\sigma=5\%$	$\sigma=10\%$	Speedup
Glacier heat	Deterministic	—	—	—	1×
	Pert. (1st)	0.71	0.62	0.54	29×
	ShapeUQ	0.92	0.91	0.90	31×
Protein electro.	Deterministic	—	—	—	1×
	Pert. (1st)	0.69	0.61	0.50	21×
	ShapeUQ	0.93	0.91	0.89	24×
Coral reef flow	Deterministic	—	—	—	1×
	Pert. (1st)	0.73	0.64	0.57	16×
	ShapeUQ	0.91	0.91	0.90	14×

Table 2. **GP kernel ablation** (glacier, $\sigma = 5\%$).

Kernel	Coverage	CI width
RBF (squared exp.)	0.87	0.031
Matérn- $\frac{1}{2}$	0.93	0.047
Matérn- $\frac{3}{2}$	0.91	0.038
Matérn- $\frac{5}{2}$	0.91	0.037
Spectral mixture ($Q = 4$)	0.84	0.029

380 **First-order perturbation without the GP is insufficient.**
381 The 1st-order perturbation baseline uses only the operator
382 norm bound (Eq. 5) without the GP covariance, producing
383 a conservative estimate at low noise (71–73% coverage at
384 $\sigma = 1\%$) and an anti-conservative one at high noise (50–
385 57% at $\sigma = 10\%$). The GP model is essential for calibrated
386 CIs.

387 **Substantial speedup.** ShapeUQ achieves 14×–31×
388 speedup over MC-500. The coral reef domain has the lowest
389 speedup because its complex geometry (thin plate-like struc-
390 tures with curvature singularities) requires more GP fitting
391 iterations; the glacier has the highest because its smoother
392 surface makes the Matérn- $\frac{3}{2}$ kernel a tight fit.

393 7.3. Ablation: GP Kernel Choice

394 Table 2 studies the effect of kernel choice on the glacier
395 domain at $\sigma = 5\%$. The Matérn- $\frac{3}{2}$ kernel provides the best
396 coverage; RBF slightly under-smooths the covariance; the
397 spectral mixture kernel overfits, reducing coverage.

398 7.4. Theoretical Bound Tightness

399 Figure 2 compares the analytical bound of Theorem 2 (Eq. 6)
against the empirical variance from MC-500, as a function of
 $\text{tr}(K)$ (modulated by varying σ). The bound is tight within a




Figure 2. Analytical upper bound vs. empirical variance (MC-500) as a function of geometric uncertainty magnitude $\text{tr}(K)$. Glacier domain. The bound is tight within $2.1\times$.

402 factor of $2.1\times$ across all configurations, confirming that the
403 first-order approximation captures the dominant uncertainty
404 contribution.

405 8. Conclusion

406 We introduced ShapeUQ, the first framework to formally
407 propagate neural-SDF reconstruction uncertainty through
408 PDE simulations using classical shape calculus. Our three
409 theorems establish: (i) a closed-form sensitivity operator
410 computable by automatic differentiation, (ii) a probabilistic
411 bound on PDE solution uncertainty under Gaussian geo-
412 metric reconstruction error, and (iii) an adjoint method that
413 computes the full sensitivity field in one additional PDE
414 solve. On SciUQ-3D—covering glacier, protein, and reef
415 geometries—ShapeUQ achieves $\geq 90\%$ confidence interval
416 coverage at $14\times$ – $31\times$ lower computational cost than Monte
417 Carlo sampling.

418 **Limitations..** The first-order approximation breaks down
419 for large noise ($\sigma > 15\%$); second-order terms from the
420 shape Hessian would be needed. Nonlinear PDEs (e.g., full
421 Navier–Stokes) require extending the adjoint method to the
422 nonlinear setting. Finally, the GP model assumes smooth,
423 stationary covariance, which may not capture sharp geo-
424 metric features in highly irregular surfaces such as coral or
425 trabecular bone.

426 **Broader impact..** Rigorous uncertainty quantification for
427 physical simulations on reconstructed 3D geometries is crit-
428 ical for decision-making in high-stakes scientific domains:
429 glacier melt rate projections for sea-level rise, drug-target
430 binding predictions in computational biology, and structural
431 integrity assessments in civil engineering. ShapeUQ pro-
432 vides a first principled and computationally tractable step
433 toward this goal.

434 References

435 [1] Grégoire Allaire, François Jouve, and Anca-Maria Toader. 436
437 Structural optimization using sensitivity analysis and a level 438
439 set method. *Journal of Computational Physics*, 194(1):393– 440
441 393, 2004. 496

442 [2] Michel C. Delfour and Jean-Paul Zolésio. *Shapes and Ge- 443
444 ometries: Metrics, Analysis, Differential Calculus, and Op-*

441 *timization*. Society for Industrial and Applied Mathematics 442
443 (SIAM), Philadelphia, 2nd edition, 2011. 444

445 [3] Lawrence C. Evans. *Partial Differential Equations*. American 446
447 Mathematical Society, Providence, RI, 1998. 448

449 [4] Jacques Hadamard. Mémoire sur le problème d’analyse relatif 450
451 à l’équilibre des plaques élastiques encastrées. *Mémoires 452
453 présentés par divers savants à l’Académie des Sciences de 454
455 l’Institut de France*, 33:1–128, 1908. 456

457 [5] Helmut Harbrecht, Michael Peters, and Markus Siebenmorgen. 458
459 First order second moment analysis for stochastic inter- 459
460 face problems based on low-rank approximation. *ESAIM: 460
461 Mathematical Modelling and Numerical Analysis*, 47(5):1533– 461
462 1552, 2013. 462

463 [6] Michael F. Hutchinson. A stochastic estimator of the trace of 463
464 the influence matrix for Laplacian smoothing splines. *Com- 464
465 munications in Statistics — Simulation and Computation*, 18 465
466 (3):1059–1076, 1989. 466

467 [7] Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azzadeh- 467
468 nesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. 468
469 Stuart, and Anima Anandkumar. Fourier neural operator for 469
470 parametric partial differential equations. In *International 470
471 Conference on Learning Representations (ICLR)*, 2021. 471

472 [8] Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima 472
473 Anandkumar. Fourier neural operator with learned deforma- 473
474 tions for PDEs on general geometries. *Journal of Machine 474
475 Learning Research*, 24(388):1–26, 2023. 475

476 [9] Anders Logg, Kent-Andre Mardal, and Garth N. Wells. *Auto- 476
477 mated Solution of Differential Equations by the Finite Element 477
478 Method: The FEniCS Book*. Springer, Berlin, 2012. 478

479 [10] Jeong Joon Park, Peter Florence, Julian Straub, Richard New- 479
480 combe, and Steven Lovegrove. DeepSDF: Learning contin- 480
481 uous signed distance functions for shape representation. 481
482 In *IEEE/CVF Conference on Computer Vision and Pattern 482
483 Recognition (CVPR)*, pages 165–174, 2019. 483

484 [11] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. 484
485 Physics-informed neural networks: A deep learning frame- 485
486 work for solving forward and inverse problems involving 486
487 nonlinear partial differential equations. *Journal of Computa- 487
488 tional Physics*, 378:686–707, 2019. 488

489 [12] Volker H. Schulz. Efficient gradient computation for shape 489
490 optimization by parametric adjoint methods. *SIAM Journal 490
491 on Scientific Computing*, 38(2):B987–B1007, 2016. 491

492 [13] Ben Smith, Helen Amanda Fricker, Alex S. Gardner, 492
493 Brooke Medley, Johan Nilsson, Fernando S. Paolo, Nicholas 493
494 Holschuh, Shad O’Connell, Thomas Sutterley, Matthew R. 494
495 Siegfried, and H. Jay Zwally. Pervasive ice sheet mass loss 495
496 reflects competing ocean and atmosphere processes. *Science*, 496
497 368(6496):1239–1242, 2020. 497

498 [14] Andrew M. Stuart. Inverse problems in a Bayesian framework. 498
499 *Acta Numerica*, 19:451–559, 2010. 499

500 [15] Timothy J. Sullivan. *Introduction to Uncertainty Quantifica- 500
501 tion*. Springer, Cham, 2015. 501

502 [16] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Prad- 502
503 han, Ben Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, 503
504 and Henrik Kretzschmar. Block-NeRF: Scalable large scene 504
505 neural view synthesis. In *IEEE/CVF Conference on Computer 505
506 Vision and Pattern Recognition (CVPR)*, pages 8248–8258, 506
507 2022. 507

- 499 [17] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku
500 Komura, and Wenping Wang. NeuS: Learning neural implicit
501 surfaces by volume rendering for multi-view reconstruction.
502 In *Advances in Neural Information Processing Systems*, 2021.