NEURAL GEOMETRY FOR PDES: REGULARITY, STABILITY, AND CONVERGENCE GUARANTEES

Samundra Karki
Iowa State University
Ames, Iowa, USA
samundra@iastate.edu

Adarsh Krishnamurthy
Iowa State University
Ames, Iowa, USA
adarshk@iastate.edu

Baskar Ganapathysubramanian*
Iowa State University
Ames, Iowa, USA
baskarg@iastate.edu

ABSTRACT

Implicit Neural Representations (INRs) have emerged as a powerful tool for geometric representation, yet their suitability for physics-based simulation remains underexplored. While metrics like Hausdorff distance quantify surface reconstruction quality, they fail to capture the geometric regularity required for provable numerical performance. This work establishes a theoretical framework connecting INR training errors to Partial Differential Equation (PDE) (specifically, linear elliptic equation) solution accuracy. We define the minimal geometric regularity required for INRs to support well-posed boundary value problems and derive *a priori* error estimates linking the neural network's function approximation error to the finite element discretization error. Our analysis reveals that to match the convergence rate of *linear* finite elements, the INR training loss must scale *quadratically* relative to the mesh size.

1 INTRODUCTION

Accurately representing complex geometry is a foundational requirement in computational science and engineering, underpinning boundary-value problems arising in solid mechanics, fluid flow, transport, and multiphysics coupling. In most classical simulation workflows, geometry is supplied explicitly, for instance, as a CAD boundary representation or a polygonal surface, and subsequently converted into a discretization-ready form, such as a body-fitted mesh. While this pipeline is mature, generating high-quality meshes for complex shapes can be labor-intensive, computationally expensive, and prone to failure or loss of fidelity when geometric complexity increases (Nguyen et al., 2015; Liu et al., 2022; Peskin, 1972; McHenry & Bajcsy, 2008; Chiba et al., 1998).

Implicit neural representations (INRs) have recently emerged as an attractive alternative for representing geometry in a compact, continuous, and differentiable manner (Park et al., 2019; Chibane et al., 2020; Gropp et al., 2020; Sitzmann et al., 2020). An INR encodes a shape implicitly as the zero level set of a neural field $\phi_\theta : \mathbb{R}^n \to \mathbb{R}$, which is commonly trained to approximate a signed distance function (SDF). Compared to explicit surface representations, INRs offer several advantages: (i) they provide a resolution-independent description of geometry, (ii) they can be learned directly from data such as point clouds or images (Wang et al., 2021; Gropp et al., 2020), and (iii) they can be queried at arbitrary spatial locations without requiring explicit surface extraction.

Implicit fields have long been used in numerical analysis pipelines, most notably for geometry processing and mesh generation. For instance, Persson (2005) demonstrated boundary-fitted mesh construction using level-set representations, where the robustness of the procedure depends critically on the well-posedness of geometric quantities derived from the implicit field, such as normals and gradients. More recently, INRs have been integrated directly into embedded and unfitted discretizations.

*Corresponding author

In particular, Karki et al. (2025; 2026) propose direct simulation workflows on octree-based computational grids for prototypical PDEs, including linear elasticity and the incompressible Navier–Stokes equations, where boundary conditions are imposed weakly through the *Shifted Boundary Method (SBM)* (Main & Scovazzi, 2018a;b). A key ingredient in such direct pipelines is ensuring that the learned implicit field remains geometrically well-behaved in a narrow band around the interface. To this end, Karki et al. (2025) employs training strategies that enforce an approximately Eikonal structure ($\|\nabla\phi_\theta\| \approx 1$) in a tubular neighborhood, together with additional normal-based consistency losses, in order to maintain stable boundary projections and well-defined normal directions.

These developments highlight that, beyond surface reconstruction, *simulation on neural geometry fundamentally depends on differential regularity properties of the implicit field*. Nevertheless, a rigorous understanding of this requirement is still missing: *what geometric accuracy and regularity are required for an INR to support stable and convergent numerical simulation?* The majority of INR literature evaluates geometric quality using reconstruction metrics such as Chamfer or Hausdorff distance between an original surface and a reconstructed surface. While such measures quantify surface proximity, they do not, by themselves, control the stability constants of boundary value problems or guarantee the well-posedness of boundary operators used in embedded methods. In particular, a small reconstruction error does not preclude pathological behavior near the interface, such as vanishing gradients, irregular normals, or unbounded curvature, all of which can destabilize boundary condition imposition and degrade the accuracy of variational discretizations.

This reveals a gap between *reconstruction fidelity* and *simulation readiness*: there is currently no unified framework that (i) specifies minimal regularity conditions to solve PDE on INR-defined domains, (ii) quantifies how neural approximation error in the geometry propagates into PDE solutions, and (iii) connects training tolerance requirements to discretization accuracy guarantees. The objective of this work is to bridge this gap by developing error analysis that links INR approximation error to PDE solution error for elliptic boundary value problems, thereby clarifying when INRs can serve as reliable geometric proxies for scientific computing. Our analysis focuses on the Poisson-type diffusion model:

$$-\nabla\cdot(\kappa\nabla u) = f \quad \text{in } \Omega, \qquad u = u_g \quad \text{on } \partial\Omega,$$

and its INR-induced counterpart on Ω^p:

$$-\nabla\cdot(\kappa\nabla u^p) = f \quad \text{in } \Omega^p, \qquad u^p = u_g \quad \text{on } \partial\Omega^p,$$

where κ is uniformly elliptic and f is a given forcing term.

Main contributions

- **Minimal geometric regularity requirements.** We identify sufficient geometric conditions on the learned level-set function ψ that ensure Γ^p admits a well-defined tubular neighborhood and a unique normal projection map. Concretely, we require $\psi \in C^{1,1}$ in a neighborhood of the interface, a uniform non-degeneracy bound $|\nabla\psi| \geq c_0 > 0$, and bounded curvature through $\|\nabla^2\psi\|_{op} \leq C_\psi$.
- **Training error ⇒ Hausdorff geometry bound.** Let ϕ denote a reference (ground-truth) level-set function defining the true boundary $\Gamma = \{\phi = 0\}$, that encloses a true domain $\Omega = \{\phi < 0\}$. We quantify INR approximation accuracy by the level-set mismatch in a tubular neighborhood $\mathcal{T}_h := \{x : |\phi(x)| < h_{\text{tube}}\}$:

  $$\varepsilon_\infty := \|\psi - \phi\|_{L^\infty(\mathcal{T}_h)}.$$

  Under gradient non-degeneracy, this implies a Hausdorff perturbation bound of the form

  $$d_H(\Gamma, \Gamma^p) \lesssim \frac{\varepsilon_\infty}{c_0},$$

  making explicit the amplification mechanism caused by vanishing gradients.
- **Main *a priori* total error estimate (INR + FEM).** Let u_h^p denote the finite element solution of polynomial degree k on the perturbed domain Ω^p using mesh size h_{mesh}. Our main result yields the total error decomposition

  $$\|(u - u_h^p)\|_{L^2(D)} \leq C_{\text{geom}}\left(\frac{\varepsilon_\infty}{c_0}\right) + C_{\text{FEM}}\, h_{\text{mesh}}^{k+1}\, |u^p|_{H^{k+1}(\Omega^p)}, \tag{1}$$

where D is a fixed background domain containing $\Omega \cup \Omega^p$ and functions are compared via zero extension. A direct consequence is a practical training guideline: to preserve the $\mathcal{O}(h_{\text{mesh}}^2)$ FEM convergence rate, the INR must satisfy $\varepsilon_\infty = \mathcal{O}(h_{\text{mesh}}^2)$ (for linear basis functions).

Outline The remainder of the paper is organized as follows. In Section 2 we introduce notation, geometric distances, and the functional analytic setting. Section 3 states the minimal geometric regularity assumptions required for well-posed projection and curvature control of INR-defined boundaries. Section 4 derives a Hausdorff-type bound linking uniform INR function error ε_∞ to geometric boundary perturbations. Section 5 combines sharp domain perturbation estimates for elliptic problems with standard finite element approximation theory to prove the total a priori bound Equation (1), yielding explicit requirements on INR training accuracy to avoid degrading simulation convergence.

2 Preliminaries and Notation

We summarize the notation and functional-analytic framework used throughout the paper. This section is primarily added for readers coming from different backgrounds.

Geometric notation Let $\mathbb{R}^n$ denote the n-dimensional Euclidean space. Vectors are denoted by bold symbols $\mathbf{x} \in \mathbb{R}^n$. Let $\Omega \subset \mathbb{R}^n$ be an open, bounded domain with boundary $\Gamma := \partial\Omega$. The neural (INR-induced) approximation of Ω is denoted by Ω^p, with boundary Γ^p.

For $\mathbf{x} \in \mathbb{R}^n$ and $r > 0$, $B_r(\mathbf{x}) := \{\mathbf{y} \in \mathbb{R}^n : |\mathbf{x} - \mathbf{y}| < r\}$ denotes the open ball. The unit sphere is $\mathbb{S}^{n-1} := \{\mathbf{x} \in \mathbb{R}^n : |\mathbf{x}| = 1\}$. To quantify geometric discrepancies between domains, we use the Hausdorff distance

$$d_H(\Omega, \Omega^p) := \max\Big\{ \sup_{\mathbf{x}\in\Omega} \operatorname{dist}(\mathbf{x}, \Omega^p), \sup_{\mathbf{y}\in\Omega^p} \operatorname{dist}(\mathbf{y}, \Omega) \Big\}, \tag{2}$$

where $\operatorname{dist}(\mathbf{x}, S) := \inf_{\mathbf{y}\in S} |\mathbf{x} - \mathbf{y}|$.

Regularity assumptions We use standard Hölder spaces $C^{k,\alpha}$. A function is in $C^{0,1}$ if it is Lipschitz continuous, and in $C^{1,1}$ if its gradient is Lipschitz continuous. Throughout, we assume the level-set functions ϕ and ψ satisfy the tubular regularity conditions of Assumption 3.1, which imply that the interfaces $\Gamma = \{\phi = 0\}$ and $\Gamma^p = \{\psi = 0\}$ are $C^{1,1}$ hypersurfaces. For the Hessian $\nabla^2\phi$, we denote by $\|\nabla^2\phi\|_{\text{op}}$ its operator (spectral) norm.

Function spaces Let $L^2(\Omega)$ denote the space of square-integrable functions on Ω. We use standard Sobolev spaces $H^k(\Omega)$, with $H_0^1(\Omega)$ denoting functions in $H^1(\Omega)$ that vanish on Γ (in the trace sense), and $H^{-1}(\Omega)$ the dual space of $H_0^1(\Omega)$. In the analysis of elliptic problems on Lipschitz domains, we employ the Besov space $B_{2,\infty}^{3/2}$, characterized (after zero extension to a background domain D) by

$$\|u\|_{B_{2,\infty}^{3/2}(D)} := \|u\|_{H^1(D)} + \sup_{0<|h|\le 1} \frac{\|\nabla u(\cdot + h) - \nabla u(\cdot)\|_{L^2(D)}}{|h|^{1/2}}. \tag{3}$$

This is the sharp regularity scale for elliptic problems posed on Lipschitz domains and plays a central role in the stability estimates in Section 5.

Remark 2.1 (Zero extension)**.** *Functions defined on Ω or Ω^p are extended by zero to a fixed background domain $D \supset \Omega \cup \Omega^p$. This allows solutions on different domains to be compared in a common space $H_0^1(D)$.*

3 Geometric Regularity of Implicit Neural Representations

Let $\psi : \mathbb{R}^n \to \mathbb{R}$ be a scalar field implicitly defining a physical domain Ω and its interface Γ ($\psi = 0$) as shown in Figure 1:

$$\Omega = \{\mathbf{x} \in \mathbb{R}^n : \psi(\mathbf{x}) < 0\}, \qquad \Gamma = \{\mathbf{x} \in \mathbb{R}^n : \psi(\mathbf{x}) = 0\}.$$

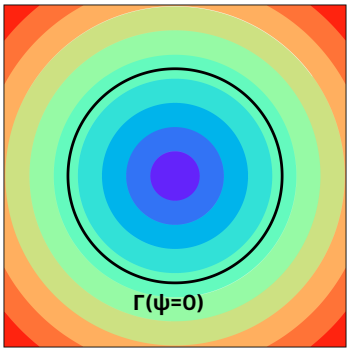

Figure 1: **Level-Set around true geometry** Γ. Implicit field (ψ) around a boundary (Γ) such that $\psi = 0$ at Γ.

For the level-set field ψ to serve as a valid geometric representation for PDE problems with Dirichlet or Neumann boundary conditions imposed on Γ, it must satisfy certain regularity and non-degeneracy properties in a neighborhood of the boundary. We now formally state the assumptions required of ψ.

Assumption 3.1 (Geometric Regularity)**.** *Let $\mathcal{T}_h$ be a tubular neighborhood of width $h > 0$ around Γ (which is zero-level set of a scalar-field ψ). We assume $\psi \in C^{1,1}(\mathcal{T}_h)$, meaning ψ is continuously differentiable with a Lipschitz continuous gradient. Furthermore, there exist constants $c_0 > 0$ and $C_\psi < \infty$ such that for all $\mathbf{x} \in \mathcal{T}_h$:*

$$\|\nabla\psi(\mathbf{x})\| \geq c_0, \qquad \|\nabla^2\psi(\mathbf{x})\|_{\text{op}} \leq C_\psi, \tag{4}$$

where $\|\cdot\|_{\text{op}}$ denotes the spectral norm of the Hessian.

These conditions imply that the unit normal field $\mathbf{n}$ is Lipschitz continuous and the principal curvatures of Γ are bounded by C_ψ/c_0.

Remark 3.2 (Valid Neural Architectures)**.** *The choice of activation function determines whether an INR satisfies Assumption 3.1:*

1. ***Smooth Activations** (C^∞): Architectures using `tanh`, `sine` (SIREN), `swish`, or `GELU` yield $\psi \in C^\infty$, satisfying the assumption locally provided gradients do not vanish.*

2. ***Higher-Order Rectifiers** ($C^{1,1}$): The Squared-ReLU, $\sigma(x) = (\max\{0, x\})^2$, yields globally $C^{1,1}$ functions with bounded Hessians, making them geometrically safe for PDE solvers.*

3. ***Standard ReLU** ($C^{0,1}$): Standard `ReLU` networks produce piecewise linear boundaries with sharp creases. They fail Assumption 3.1 as curvature is distributional (Dirac delta) and normals are discontinuous, leading to ill-posed boundary projections.*

Remark 3.3 (Spectral Bias and Effective Corner Regularization)**.** *Implicit Neural Representations are known to exhibit a* spectral bias*, whereby low-frequency components of the target function are learned preferentially over high-frequency components. From a geometric standpoint, sharp corners or edges correspond to high-frequency features of the level-set field, characterized by discontinuities in the unit normal and unbounded curvature. As a consequence, INRs trained with finite capacity and finite optimization tolerance typically approximate such features through locally rounded transitions at a small but nonzero length scale.*

This behavior implies that, even when the target geometry contains exact corners, the learned implicit field ψ often exhibits bounded curvature and Lipschitz-continuous normals at the scale resolved by the network, thereby satisfying Assumption 3.1 in an effective sense within a tubular neighborhood of the zero level set. Rather than representing a limitation, this implicit geometric regularization is advantageous for PDE-based simulation, as it restores the regularity required for stable closest-point projections and distance-based formulations.

3.1 WELL-POSEDNESS OF THE BOUNDARY PROJECTION

Fictional-domain methods (e.g., Shifted Boundary Method) require mapping points near the boundary back to Γ (Main & Scovazzi, 2018a;b; Hsu et al., 2016; Jaiswal et al., 2024). We show that Assumption 3.1 guarantees this mapping is unique, even if ψ is not a perfect Signed Distance Function (SDF).

Lemma 3.4 (Local existence and uniqueness of the normal projection)**.** *Let $\psi \in C^{1,1}(\mathcal{T}_h)$ satisfy Assumption 3.1 whose zero-level set represents boundary of domain Ω. Then, there exists a constant $h_{\text{crit}} > 0$, depending only on c_0 and C_ψ, such that for every $h \in (0, h_{\text{crit}})$ the following holds:*

For every $\mathbf{x} \in \mathcal{T}_h$, there exist unique $\mathbf{x}^* \in \Gamma$ and unique $d \in (-h, h)$ such that

$$\mathbf{x} = \mathbf{x}^* + d\,\mathbf{n}(\mathbf{x}^*), \qquad \mathbf{n}(\mathbf{x}^*) = \frac{\nabla\psi(\mathbf{x}^*)}{\|\nabla\psi(\mathbf{x}^*)\|}. \tag{5}$$

Proof. **Step 1: Regularity of Γ and the unit normal.** By Assumption 3.1, $\psi \in C^{1,1}(\mathcal{T}_h)$ and $\|\nabla\psi(x)\| \geq c_0 > 0$ on $\mathcal{T}_h$. Hence the zero level set

$$\Gamma := \{x \in \mathbb{R}^n : \psi(x) = 0\}$$

is a $C^{1,1}$ embedded hypersurface. Define the unit normal field

$$\mathbf{n}(x) := \frac{\nabla\psi(x)}{\|\nabla\psi(x)\|}, \qquad x \in \Gamma.$$

For $x, y \in \Gamma$, using $\left\|\frac{a}{\|a\|} - \frac{b}{\|b\|}\right\| \leq \frac{2}{\min(\|a\|,\|b\|)}\|a - b\|$ and $\|\nabla\psi\| \geq c_0$ on $\Gamma \subset \mathcal{T}_h$, we obtain

$$\|\mathbf{n}(x) - \mathbf{n}(y)\| \leq \frac{2}{c_0}\,\|\nabla\psi(x) - \nabla\psi(y)\| \leq \frac{2C_\psi}{c_0}\,\|x - y\|,$$

since $\nabla\psi$ is Lipschitz on $\mathcal{T}_h$ with Lipschitz constant bounded by C_ψ. Thus $\mathbf{n}$ is Lipschitz on Γ with

$$\text{Lip}(\mathbf{n}) \leq L_n := \frac{2C_\psi}{c_0}.$$

Step 2: Existence of a reach neighborhood and uniqueness of the metric projection. Since Γ is a $C^{1,1}$ hypersurface with Lipschitz unit normal, it has *positive reach*. Consequently, there exists $\rho > 0$, depending only on L_n, such that the nearest-point projection onto Γ is well defined and unique on

$$U_\rho := \{x \in \mathbb{R}^n : \text{dist}(x, \Gamma) < \rho\}.$$

Step 3: Embedding the level-set tube into the reach neighborhood. We show that for h sufficiently small, $\mathcal{T}_h \subset U_\rho$. Fix $x \in \mathcal{T}_h$ and consider the ODE

$$\dot{\gamma}(t) = -\frac{\nabla\psi(\gamma(t))}{\|\nabla\psi(\gamma(t))\|^2}, \qquad \gamma(0) = x. \tag{6}$$

Along this trajectory,

$$\frac{d}{dt}\psi(\gamma(t)) = \nabla\psi(\gamma(t)) \cdot \dot{\gamma}(t) = -1,$$

hence $\psi(\gamma(t)) = \psi(x) - t$. Let $T := |\psi(x)|$. For $t \in [0, T]$ we have $\gamma(T) \in \Gamma$, and

$$\text{dist}(x, \Gamma) \leq \int_0^T \|\dot{\gamma}(t)\|\,dt \leq \frac{|\psi(x)|}{c_0} < \frac{h}{c_0}.$$

Choosing $h_{\text{crit}} := c_0\,\rho$, we have $\mathcal{T}_h \subset U_\rho$ whenever $h < h_{\text{crit}}$.

Step 4: Normal-coordinate representation and uniqueness. For $x \in \mathcal{T}_h$, the unique nearest-point $x^* = \Pi_\Gamma(x) \in \Gamma$ exists, and the first-order optimality condition implies $x - x^*$ is parallel to $\mathbf{n}(x^*)$. Define $d := (x - x^*) \cdot \mathbf{n}(x^*)$, yielding Equation (5). □

This result confirms that INRs need not strictly satisfy the Eikonal equation $\|\nabla\psi\| = 1$. As long as gradients are non-vanishing ($c_0 > 0$) and curvature is controlled, the geometry is valid for simulation.

4 GEOMETRIC PERTURBATION ANALYSIS

In practice, we approximate the true geometry ϕ with a neural network ψ. We now quantify the geometric error introduced by this approximation. Let:

$$\phi(\mathbf{x}) \quad \text{(True Level Set)}, \qquad \psi(\mathbf{x}) \quad \text{(Neural Approximation)}.$$

We assume both fields satisfy Assumption 3.1 in a common tube $\mathcal{T}_h$. Furthermore, we assume the network is trained to a uniform tolerance:

$$\|\psi - \phi\|_{L^\infty(\mathcal{T}_h)} \leq \varepsilon_\infty. \tag{7}$$

Lemma 4.1 (Hausdorff Distance Bound)**.** *Let $\Gamma = \{\phi = 0\}$ represent true zero-level set and $\Gamma^p = \{\psi = 0\}$ represent corresponding zero-level set given by a neural approximation. Under the regularity conditions of Assumption 3.1 and the compatibility Equation* (7)*, the Hausdorff distance is bounded by:*

$$d_H(\Gamma, \Gamma^p) \leq \frac{\varepsilon_\infty}{\min(c_0, \tilde{c}_0)} + \mathcal{O}(\varepsilon_\infty^2), \tag{8}$$

where $c_0, \tilde{c}_0$ are the lower bounds on $\|\nabla\phi\|$ and $\|\nabla\psi\|$ respectively.

Proof. Let $\Gamma = \{\phi = 0\}$ and $\Gamma^p = \{\psi = 0\}$, and assume: (i) $\phi, \psi \in C^{1,1}(\mathcal{T}_h)$ for a common tubular neighborhood $\mathcal{T}_h$; (ii) the nondegeneracy and curvature bounds

$$\|\nabla\phi\| \geq c_0, \quad \|\nabla^2\phi\|_{\text{op}} \leq C_\phi, \qquad \|\nabla\psi\| \geq \tilde{c}_0, \quad \|\nabla^2\psi\|_{\text{op}} \leq \tilde{C}_\psi \quad \text{on } \mathcal{T}_h;$$

(iii) the uniform field error $\|\psi - \phi\|_{L^\infty(\mathcal{T}_h)} \leq \varepsilon_\infty$.

Step 1: From Γ to Γ^p Fix $x \in \Gamma$ so $\phi(x) = 0$. Then $|\psi(x)| \leq \varepsilon_\infty$. Define

$$n_\psi(x) := \frac{\nabla\psi(x)}{\|\nabla\psi(x)\|}, \qquad \sigma := \text{sign}(\psi(x)).$$

Consider the ray $r(s) = x - \sigma s\, n_\psi(x)$ and $g(s) := \sigma\, \psi(r(s))$. Then $g(0) = |\psi(x)| > 0$. By chain rule and Lipschitz continuity of $\nabla\psi$, we obtain

$$g'(s) \leq -\|\nabla\psi(x)\| + \tilde{C}_\psi s.$$

Integrating yields

$$g(s) \leq |\psi(x)| - \|\nabla\psi(x)\|\, s + \tfrac{1}{2}\tilde{C}_\psi s^2.$$

Under suitable smallness conditions on ε_∞, there exists $\bar{s} \in (0, s_{-x}]$ such that $\psi(r(\bar{s})) = 0$, and

$$\text{dist}(x, \Gamma^p) \leq \bar{s} \leq \frac{\varepsilon_\infty}{\tilde{c}_0} + \frac{\tilde{C}_\psi}{2\tilde{c}_0^3}\varepsilon_\infty^2 + \mathcal{O}(\varepsilon_\infty^3).$$

Step 2: From Γ^p to Γ By symmetry, we obtain

$$\sup_{y \in \Gamma^p} \text{dist}(y, \Gamma) \leq \frac{\varepsilon_\infty}{c_0} + \frac{C_\phi}{2c_0^3}\varepsilon_\infty^2 + \mathcal{O}(\varepsilon_\infty^3).$$

Step 3: Hausdorff distance Taking the maximum of the two directed distances yields the result. □

Remark 4.2 (Impact of Vanishing Gradients)**.** *The bound is scaled by $1/c_0$. This highlights a critical training requirement: merely minimizing the residual $|\psi - \phi|$ is insufficient. The training loss must also enforce the Eikonal constraint or gradient regularization to prevent $c_0 \to 0$. If the INR gradients vanish at the interface, the geometric error d_H becomes unbounded, destroying simulation accuracy.*

5 STABILITY AND DISCRETIZATION ON PERTURBED DOMAINS

In this section we quantify how geometric approximation errors in an INR-defined domain propagate to the solution of elliptic boundary value problems. Our presentation follows the quantitative domain perturbation framework of Savaré and Schimperna (Savaré & Schimperna, 2002), which yields sharp stability estimates under minimal geometric regularity (Lipschitz boundaries).

5.1 PROBLEM FORMULATION

Let $D \subset \mathbb{R}^n$ be a bounded background domain containing both the true domain Ω and the perturbed domain Ω^p. We consider the Dirichlet problems

$$\text{(True)} \quad \begin{cases} -\nabla \cdot (\kappa \nabla u) = f & \text{in } \Omega, \\ u = 0 & \text{on } \Gamma := \partial\Omega, \end{cases} \qquad \text{(Perturbed)} \quad \begin{cases} -\nabla \cdot (\kappa \nabla u^p) = f & \text{in } \Omega^p, \\ u^p = 0 & \text{on } \Gamma^p := \partial\Omega^p, \end{cases} \tag{9}$$

where $\kappa \in L^\infty(D)$ satisfies uniform ellipticity

$$0 < \kappa_{\min} \le \kappa(x) \le \kappa_{\max} < \infty,$$

and $f \in L^2(D)$. We extend solutions by zero outside their domains so that $u, u^p \in H_0^1(D)$ for comparison on D.

5.2 GEOMETRIC REGULARITY: UNIFORM CONE CONDITION

The stability theory of Savaré & Schimperna (2002) requires a minimal boundary regularity assumption formulated via a uniform cone condition.

Assumption 5.1 (Uniform (ρ, θ)-cone condition)**.** *The domain $\Omega \subset \mathbb{R}^n$ satisfies a uniform (ρ, θ)-cone condition if there exist constants $\rho > 0$ and $\theta \in (0, \pi/2)$ such that for every $x_0 \in \partial\Omega$ there exists a unit vector $n(x_0) \in \mathbb{S}^{n-1}$ such that, for all*

$$h \in C_{\rho,\theta}(n(x_0)) := \big\{ h \in \mathbb{R}^n : 0 < |h| < \rho,\ h \cdot n(x_0) > |h| \cos\theta \big\},$$

the inclusions

$$(B_{3\rho}(x_0) \cap \Omega) - h \subset \Omega, \qquad (B_{3\rho}(x_0) \setminus \Omega) + h \subset \mathbb{R}^n \setminus \Omega \tag{10}$$

hold.

Remark 5.2 (Relation to INR regularity)**.** *Our level-set regularity assumption (Assumption 3.1) implies that $\partial\Omega$ is a $C^{1,1}$ hypersurface, hence Lipschitz, and therefore satisfies Assumption 5.1.*

5.3 MAIN STABILITY ESTIMATE

Let $d_\star := d_H(\Gamma, \Gamma^p)$ denote the Hausdorff distance between boundaries.

Theorem 5.3 (Stability under domain perturbation (Savaré & Schimperna, 2002))**.** *Assume that $\Omega \subset \mathbb{R}^n$ satisfies Assumption 5.1. Suppose that the perturbation is sufficiently small:*

$$d_\star \le \tfrac{1}{2}\, \rho \sin\theta. \tag{11}$$

Let $u \in H_0^1(\Omega)$ and $u^p \in H_0^1(\Omega^p)$ be the weak solutions of Equation (9)*. Then, after extension by zero to D, the following estimates hold:*

$$\|\nabla u - \nabla u^p\|_{L^2(D)} \le C_1 \, \|f\|_{L^2(D)}^{1/2} \|f\|_{H^{-1}(D)}^{1/2} \left(\frac{d_\star}{\rho \sin\theta} \right)^{1/2}, \tag{12}$$

$$\|u - u^p\|_{L^2(D)} \le C_2 \, \|f\|_{L^2(D)}^{1/2} \|f\|_{H^{-1}(D)}^{1/2} \left(\frac{d_\star}{\rho \sin\theta} \right), \tag{13}$$

where C_1 and C_2 depend only on n, $\kappa_{\max}/\kappa_{\min}$, and $\operatorname{diam}(D)$.

5.4 CONNECTION TO INR GEOMETRY

Theorem 5.3 bridges INR approximation quality and PDE solution accuracy. Combining it with the geometric perturbation estimate from Section 4 yields an explicit INR-to-PDE propagation bound.

Corollary 5.4 (INR-to-PDE error propagation in L^2)**.** *Under Assumptions 3.1 and 5.1, let $\Omega = \{\phi < 0\}$ and $\Omega^p = \{\psi < 0\}$. Assume the INR field ψ satisfies, on a tubular neighborhood $\mathcal{T}_h$,*

$$\|\phi - \psi\|_{L^\infty(\mathcal{T}_h)} \le \varepsilon_\infty.$$

Then the corresponding solutions u on Ω and u^p on Ω^p satisfy

$$\|u - u^p\|_{L^2(D)} \le \tilde{C}_{L^2} \, \|f\|_{L^2(D)}^{1/2} \, \|f\|_{H^{-1}(D)}^{1/2} \left(\frac{\varepsilon_\infty}{c_0 \, \rho \sin\theta} \right), \tag{14}$$

where $c_0 := \inf_{x \in \Gamma} \|\nabla \phi(x)\|$.

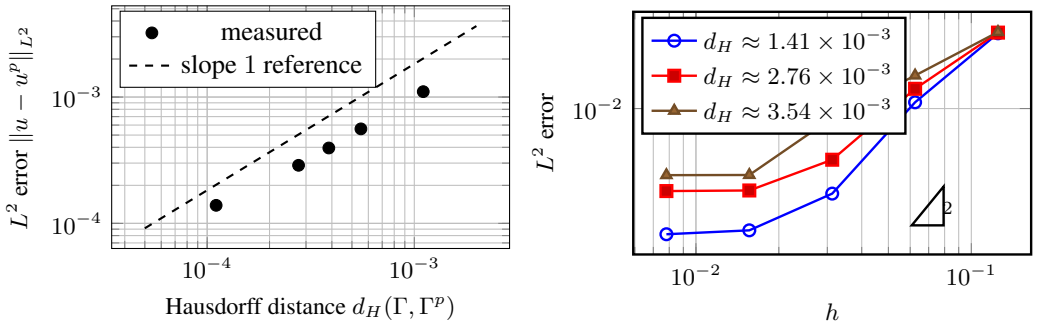

(a) L^2 error vs. Hausdorff distance.

(b) Mesh refinement under fixed d_H.

Figure 2: **Solution error under geometric perturbations.** Left: error vs. Hausdorff distance. Right: mesh refinement convergence.

5.5 Total error: geometry + discretization

In practice we solve the perturbed-domain problem numerically on Ω^p using a finite element discretization with mesh size h_{mesh}. Let u_h^p denote the discrete solution. The total error measured on D admits the decomposition

$$\|u - u_h^p\|_{L^2(D)} \leq \|u - u^p\|_{L^2(D)} + \|u^p - u_h^p\|_{L^2(\Omega^p)}.$$

Theorem 5.5 (A priori total L^2 error estimate)**.** *Under the assumptions of Corollary 5.4, assume conforming finite elements of polynomial degree k are used on Ω^p and that $u^p \in H^{k+1}(\Omega^p)$. Then there exist constants $C_{\text{geom}}, C_{\text{FEM}} > 0$, independent of h_{mesh} and ε_∞, such that*

$$\|u - u_h^p\|_{L^2(D)} \leq C_{\text{geom}} \, \|f\|_{L^2(D)}^{1/2} \, \|f\|_{H^{-1}(D)}^{1/2} \left(\frac{\varepsilon_\infty}{c_0}\right) + C_{\text{FEM}} \, h_{\text{mesh}}^{k+1} \, |u^p|_{H^{k+1}(\Omega^p)}. \tag{15}$$

Remark 5.6 (Balancing geometry and discretization errors)**.** *Assume the forcing norms $\|f\|_{L^2(D)}$, $\|f\|_{H^{-1}(D)}$ and the regularity seminorm $|u^p|_{H^{k+1}(\Omega^p)}$ are $\mathcal{O}(1)$, and that the constants C_{geom} and C_{FEM} are comparable. Balancing the two terms in Equation* (15) *yields the guideline*

$$\varepsilon_\infty \sim h_{\text{mesh}}^{k+1}.$$

In particular, for $k = 1$ one expects $\varepsilon_\infty \sim h_{\text{mesh}}^2$, while for $k = 2$ one expects $\varepsilon_\infty \sim h_{\text{mesh}}^3$.

6 Results

This section provides numerical evidence supporting our final theorem. We connect geometric mismatch to downstream PDE accuracy by solving a Poisson problem on perturbed geometries with an unfitted octree discretization and shifted boundary method (SBM) boundary treatment, demonstrating how solution error scales with Hausdorff distance and how mesh refinement saturates when geometry error dominates.

We empirically examine how geometric boundary perturbations influence the numerical solution of an elliptic boundary value problem. The reference domain Ω is the unit disk. We consider the Poisson problem with manufactured solution $u(x, y) = \sin(\pi x)$.

Figure 2(left) shows the measured L^2 error versus Hausdorff distance. The data exhibits approximately linear scaling, consistent with theoretical predictions. Figure 2(right) shows convergence with mesh refinement for fixed perturbation levels. The curves plateau when the mesh becomes sufficiently fine, indicating geometry-limited error floors.

7 Conclusions

We have developed a unified theoretical framework that links the approximation quality of Implicit Neural Representations (INRs) to the stability and accuracy of elliptic PDE solutions posed on INR-defined domains. Our main contributions are:

1. **Minimal geometric regularity for well-posed projection.** We identified sufficient regularity conditions ($\psi \in C^{1,1}$, gradient non-degeneracy $\|\nabla\psi\| \geq c_0 > 0$, bounded curvature) that guarantee existence and uniqueness of normal projection.
2. **Training error ⇒ Hausdorff boundary perturbation.** We derived bounds $d_H(\Gamma, \Gamma^p) \lesssim \varepsilon_\infty / c_0$, making explicit the role of gradient magnitude.
3. **Geometry-to-solution stability and training guidelines.** We obtained explicit a priori bounds showing how geometric perturbations propagate into PDE solution error, yielding practical training requirements.

Overall, our results provide a rigorous framework for using INRs as a reliable geometric primitive in high-fidelity scientific computing, clarifying which regularity and training requirements are necessary for stable, convergent PDE simulation.

REFERENCES

N. Chiba, I. Nishigaki, Y. Yamashita, C. Takizawa, and K. Fujishiro. A flexible automatic hexahedral mesh generation by boundary-fit method. *Computer Methods in Applied Mechanics and Engineering*, 161(1):145–154, 1998. ISSN 0045-7825.

Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6970–6981, virtual, 2020. IEEE.

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*, ICML'20, pp. 1–12, virtual, 2020. JMLR.org.

Ming-Chen Hsu, Chenglong Wang, Fei Xu, Austin J Herrema, and Adarsh Krishnamurthy. Direct immersogeometric fluid flow analysis using B-rep CAD models. *Computer Aided Geometric Design*, 43:143–158, 2016.

Monu Jaiswal, Ashton M. Corpuz, and Ming-Chen Hsu. Mesh-driven resampling and regularization for robust point cloud-based flow analysis directly on scanned objects. *Computer Methods in Applied Mechanics and Engineering*, 432:117426, 2024. ISSN 0045-7825. doi: https://doi.org/10.1016/j.cma.2024.117426. URL https://www.sciencedirect.com/science/article/pii/S0045782524006819.

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, 2026. ISSN 0010-4485. doi: https://doi.org/10.1016/j.cad.2025.103978. URL https://www.sciencedirect.com/science/article/pii/S0010448525001393.

Wing Kam Liu, Shaofan Li, and Harold S Park. Eighty years of the finite element method: birth, evolution, and future. *Archives of Computational Methods in Engineering*, 29(6):4431–4453, 2022.

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, 2018a.

Alex Main and Guglielmo Scovazzi. The shifted boundary method for embedded domain computations. Part II: Linear advection–diffusion and incompressible Navier–Stokes equations. *Journal of Computational Physics*, 372:996–1026, 2018b.

Kenton McHenry and Peter Bajcsy. An overview of 3D data content, file formats and viewers. Technical Report ISDA08-002, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, 1205 W Clark St, Urbana, IL 61801, October 2008. Technical Report, Image Spatial Data Analysis Group.

Vinh Phu Nguyen, Cosmin Anitescu, Stéphane PA Bordas, and Timon Rabczuk. Isogeometric analysis: An overview and computer implementation aspects. *Mathematics and Computers in Simulation*, 117:89–116, 2015.

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. 1–12, Long Beach, June 2019. IEEE.

Per-Olof Persson. *Mesh generation for implicit geometries*. PhD thesis, Massachusetts Institute of Technology, 2005.

Charles S Peskin. Flow patterns around heart valves: A numerical method. *Journal of Computational Physics*, 10(2):252–271, 1972.

Giuseppe Savaré and Giulio Schimperna. Domain perturbations and estimates for the solutions of second order elliptic equations. *Journal de mathématiques pures et appliquées*, 81(11):1071–1112, 2002.

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 7462–7473, virtual, 2020. Curran Associates, Inc. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/53c04118df112c13a8c34b38343b9c10-Paper.pdf.

Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, pp. 27171–27183, 2021.