Geometry-Aware Learnable Gabor Networks for High-Fidelity Unsigned Distance Field Reconstruction

Xiaohui Jia, Yuan Zhang*, Caiqin Jia, Xiaowen Yang, Min Pang
North University of China, Taiyuan, Shanxi, China
15835290296@163.com, {20040835, 20220076, 20050817, pmin}@nuc.edu.cn

***Abstract*—Unsigned Distance Functions (UDFs) provide flexible representations for non-watertight geometries but struggle with a critical spectral mismatch. Since 3D surfaces exhibit non-stationary frequency distributions, demanding high frequencies for sharp edges but low frequencies for flat regions, traditional global encodings are inherently inefficient. They typically enforce excessive bandwidth globally to resolve local details, inevitably inducing computational redundancy and ripple artifacts in smooth areas. To address this, we propose a geometry-aware Learnable Gabor Network. First, leveraging the time-frequency locality of Gabor wavelets, our network adaptively modulates its spectral bandwidth, concentrating representational capacity on geometric details while preventing noise in flat regions. Second, to mitigate the gradient ambiguity of UDFs, we introduce a Dual-Query Geometric Consistency mechanism. By imposing local differential constraints, we regularize the gradient field direction without relying on ground-truth normal supervision. Experiments demonstrate that our approach improves upon state-of-the-art methods, achieving a superior balance between reconstruction fidelity and efficiency.**

***Index Terms*—3d point cloud, gabor, udf, implicit function, surface reconstruction**

I. Introduction

Neural implicit representations, particularly Signed Distance Functions(SDFs) [1], [2], have excelled in closed manifold reconstruction. However, their watertightness assumption limits applications in open and non-manifold structures. Unsigned Distance Functions(UDFs) [3] naturally handle arbitrary topologies by modeling the zero-level set, emerging as a versatile alternative for general implicit geometric modeling.

Despite relaxing topological constraints, UDFs still suffer from the spectral bias of implicit neural representations. As illustrated in Fig. 1, 3D geometric features exhibit strong spatial-frequency non-stationarity: smooth regions correspond to low-frequency components, whereas corners and thin structures manifest as localized high-frequency signals. Conventional Fourier-based encodings [4], [5], due to their global support, require introducing high-frequency components across the entire domain to approximate such localized details. This results in redundant computation and induces Gibbs-like artifacts in smooth regions.

*Corresponding author: Yuan Zhang (20040835@nuc.edu.cn). Supported by NSFC (62272426), Shanxi NSF (202303021212206), and Shanxi Key R&D (202402020101001, Digital Cultural Heritage Preservation).

Fig. 1. Close-up visualization on the Gargoyle model highlighting the distinct spatial-frequency characteristics: flat regions are dominated by low-frequency components, while sharp edges and fine textures (zoomed patches) require localized high-frequency representation to prevent over-smoothing commonly seen in global Fourier-based UDF methods.

Based on these observations, we propose a Geometry-Aware Learnable Gabor Network to mitigate the scale mismatch in UDF learning via localized time-frequency analysis. We introduce a learnable Gabor encoding that, unlike global isotropic Fourier features, leverages the optimal spatio-frequency locality of Gabor wavelets. By making the basis parameters trainable, the model adaptively concentrates capacity on local high-frequency geometric details while preserving smooth responses in flat regions.

To resolve gradient ambiguity and sign deficiency near the UDF zero level set, we present a Dual-Query Geometric Consistency mechanism. Using local perturbation sampling around each point, it enforces gradient reversal and numerical equivariance in a fully unsupervised manner, effectively preventing degradation in thin structures and complex topologies.

The main contributions are:

- Learnable Gabor-based UDF framework addressing spectral mismatch and Gibbs artifacts;
- Self-supervised Dual-Query Geometric Consistency for stable gradients in non-manifold cases;
- Superior accuracy and efficiency over state-of-the-art on multiple benchmarks.

II. Related Work

A. Topology-Agnostic Implicit Surface Reconstruction

Early neural implicit representations were primarily based on Signed Distance Functions (SDF) or occupancy fields, with representative works such as DeepSDF [6], NeuS [7], and NeuralAngelo [4]. These methods achieved extremely high geometric fidelity in watertight object reconstruction by modeling the signed distance from spatial points to closed manifolds, combined with volume rendering or sphere tracing algorithms. However, the inherent topological closedness assumption of SDF constitutes a fundamental limit in real-world applications.

Raw scan data from the real world, such as LiDAR or Kinect point clouds, often contain open surfaces, non-manifold structures, and internal holes.

To break this limitation, NDF [8] first introduced the Unsigned Distance Function (UDF) as a topology-agnostic general representation. Subsequent research has been dedicated to improving the convergence quality of UDF: CAP [9] proposed a consistency-aware field optimization strategy, while MeshUDF [10] solved the difficulty of extracting zero level set meshes from UDF. Despite relaxing topological constraints, the non-differentiability of gradients near the zero level set of UDF remains an unresolved issue. Existing methods mostly rely on global smoothness constraints to suppress noise, which leads to over-smoothed reconstruction results and loss of key high-frequency details when processing sharp geometric features such as edges of CAD models or thin-walled structures.

B. Spectral Encoding and Bias of Neural Fields

Neural Tangent Kernel theory [11] reveals that standard MLPs exhibit strong low-frequency bias, hindering the fitting of complex high-frequency geometric details. Positional Encoding [12] and SIREN [5] mitigate this issue by mapping inputs to high-dimensional Fourier features and employing periodic activations, respectively, greatly improving detail representation. More recent approaches, such as Instant-NGP [13] (multi-resolution hash grids) and 2D Gaussian Splatting [14] (anisotropic kernels), further enable efficient capture of high-frequency content.

Nevertheless, when applied to 3D geometry, existing encoding schemes suffer from a fundamental scale mismatch. Global-support Fourier bases must introduce high frequencies across the entire domain to model spatially localized singularities, resulting in spectral leakage and noticeable Gibbs artifacts. Although Gaussian-based methods exploit spatial locality, their explicit formulation does not easily extend to continuous implicit UDF modeling. Thus, designing an implicit encoding that jointly achieves spatial locality and learnability remains critical to overcoming the current fidelity bottleneck.

C. Geometric Consistency and Self-Supervised Learning

In unsupervised point cloud reconstruction without normal supervision, effective geometric priors are essential. Points2Surf [15] and IGR [16] enforce gradient regularization via Eikonal constraints, while recent UDF methods such as DUDF [3] (hyperbolic scaling) and LevelSetUDF [17] (level-set projection) improve field continuity.

However, these approaches primarily rely on first-order differential properties, which are prone to failure or spurious solutions in distant regions and lack explicit mechanisms to constrain local gradient directions. This often leads to severe degradation in thin-walled structures, including sign cancellation or vanishing zero level sets due to ambiguous directional guidance.

In contrast, our method avoids global manifold assumptions and instead leverages local sampling duality to impose fully self-supervised constraints, achieving robust and precise gradient locking under unsupervised conditions.

III. Method

This section details our proposed unsupervised UDF reconstruction framework, as illustrated in Fig. 2. Addressing the two core challenges of spectral mismatch and gradient ambiguity in implicit geometric modeling, we adopt a divide-and-conquer strategy: on one hand, precise locking of the gradient field direction is achieved under unsupervised conditions via *Dual-Query Geometric Consistency*; on the other hand, adaptive capture of high-frequency details and feature decoupling are realized at the feature level via *Geometry-Aware Learnable Gabor Encoding*.

A. Dual-Query Geometric Consistency

Traditional unsupervised UDF learning relies on nearest-neighbor distance minimization, which lacks directional guidance and often causes zero-level-set disappearance or gradient oscillation in thin-walled structures and complex topologies. To lock the gradient field without requiring normals or watertightness assumptions, we propose a point-centered dual-query sampling strategy.

We emphasize that this strategy does not rely on any global symmetry of the object, but only on local manifold properties. For each observation point $p_i \in P$, we construct a dual query pair by a small perturbation along a random direction:

$$q_i^+ = p_i + \delta_i, \quad q_i^- = p_i - \delta_i. \tag{1}$$

Validity of the dual-query assumption. The DQGC construction assumes that p_i remains the nearest surface projection for both q_i^+ and q_i^-. This assumption may be violated under sparse/noisy sampling, high-curvature regions, or thin-walled structures where a query may snap to the opposite side. To make the assumption hold in practice, we enforce DQGC only within a narrow near-surface band and use an adaptive, truncated perturbation whose magnitude is tied to the local sampling scale.

Let r_i denote the average distance from p_i to its k nearest neighbors in P (a proxy of local sampling scale), and let $u_i \sim \mathcal{U}(\mathbb{S}^2)$ be a random unit direction. We sample $\delta_i = \beta_i r_i u_i$ with $\beta_i \in [\beta_{\min}, \beta_{\max}]$, followed by a hard truncation $\|\delta_i\|_2 \leq \delta_{\max}$.

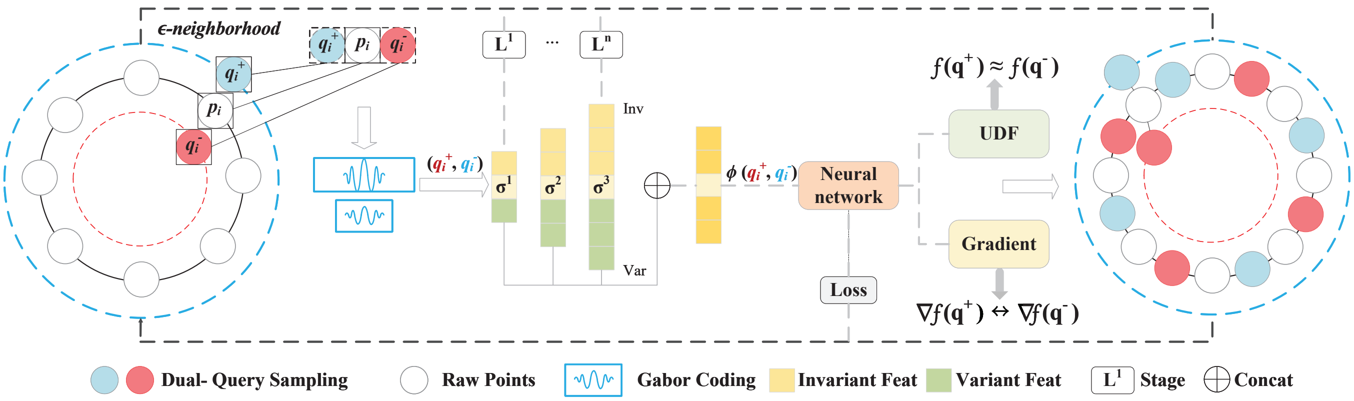

Fig. 2. The pipeline of our proposed Geometry-Aware Gabor Network. We first generate local dual-query pairs via point-centric sampling. These queries are encoded by a multi-scale learnable Gabor layer, where invariant and variant features are decoupled and concatenated. Finally, the network is supervised by Dual-Query Geometric Consistency (DQGC) losses, enforcing value equivalence and gradient contra-directionality.

Based on this construction, the true unsigned distance-to-surface function $f(x)$ satisfies two intrinsic local geometric properties: *(i) numerical equivariance*—q_i^+ and q_i^- are equidistant to p_i, thus $f(q_i^+) = f(q_i^-)$ as $\|\delta_i\|_2 \to 0$; and *(ii) gradient reversal*—the distance gradient points to the fastest distance increase, hence the gradients at q_i^+ and q_i^- should be opposite when they lie on two sides of the same local surface patch. We encourage our network f_θ to satisfy the above local properties through the following self-supervised loss.

Near-field-only directional locking. The contra-directionality prior $\nabla f(q_i^+) \approx -\nabla f(q_i^-)$ is valid only when (q_i^+, q_i^-) are located on two sides of the *same local surface patch*. This typically holds in a narrow near-surface band, but may break near the medial axis where the nearest surface point is not unique and the UDF gradient is discontinuous. Therefore, we enforce DQGC as a *near-field constraint* only.

We define a near-field set Ω_{near} by sampling queries around P using the truncated perturbations above, and apply the DQGC loss only for $(q_i^+, q_i^-) \subset \Omega_{\text{near}}$:

$$L_{\text{dqgc}} = \mathbb{E}_{(q_i^+, q_i^-) \subset \Omega_{\text{near}}} \Bigg[\|f_\theta(q_i^+) - f_\theta(q_i^-)\|_1 + \left\| \frac{\nabla f_\theta(q_i^+)}{\|\nabla f_\theta(q_i^+)\|_2 + \epsilon} + \frac{\nabla f_\theta(q_i^-)}{\|\nabla f_\theta(q_i^-)\|_2 + \epsilon} \right\|_2 \Bigg]. \tag{2}$$

We do *not* impose this constraint on far-field samples to avoid erroneous locking in regions with ambiguous projections.

B. Geometry-Aware Learnable Gabor Encoding

Existing positional encodings such as Fourier Features have global support, forcing the network to increase the frequency bandwidth *globally* to fit spatially localized corners, which results in severe spectral leakage and unnecessary computation. To address this, we design a learnable Gabor encoding that exploits time-frequency locality for efficient allocation of representational capacity.

We define the encoding layer as a filter bank containing K basis functions. For an input coordinate $x \in \mathbb{R}^3$, the response of the k-th Gabor basis is given by the product of a Gaussian envelope and a sinusoidal oscillation:

$$\phi_k(x) = \exp\!\left(-\|M_k(x-\mu_k)\|^2\right) \cdot \cos\!\left(2\pi w_k^\top (x-\mu_k) + \psi_k\right), \tag{3}$$

where μ_k is the spatial center, w_k is the frequency vector, ψ_k is the phase, and M_k controls the spatial bandwidth. Unlike fixed encodings, all parameters in this layer are trainable and optimized jointly with the MLP. In practice, we parameterize M_k as a diagonal matrix for stable anisotropic bandwidth control. By learning M_k, the network can adaptively adjust each basis receptive field—expanding envelopes in smooth regions to capture global shapes, while shrinking envelopes near sharp features to model high-frequency details without spillover noise.

Invariant/variant feature decoupling. To maximize encoding efficiency, we exploit the dual-query symmetry to construct two specialized basis types for feature decoupling: an *invariant* basis G_{inv} for modeling the scalar distance field, and a *variant* basis G_{var} for capturing directional (gradient-related) cues. Since each Gabor basis is centered at μ_k, we define the local coordinate $\tilde{x} = x - \mu_k$ and impose even/odd symmetry in this local frame:

$$\begin{aligned} G_{inv}(x) &\Leftrightarrow \phi_k(\tilde{x}) = \phi_k(-\tilde{x}), \\ G_{var}(x) &\Leftrightarrow \phi_k(\tilde{x}) = -\phi_k(-\tilde{x}). \end{aligned} \tag{4}$$

This explicit decoupling prevents the MLP from having to implicitly separate distance and direction information in deeper layers.

C. Multi-Scale Hierarchy and Adaptive Sampling

Although the Gabor encoding mitigates spectral leakage via time-frequency locality, its compact support may introduce gradient truncation: if a query lies far from all Gaussian envelopes, the encoding response vanishes and gradients become weak.

Intrinsic trade-off between Gabor and Fourier. Fourier bases have global support, which ensures gradients propagate across the entire domain but often produces ripple (Gibbs-like) artifacts when fitting spatially localized high-frequency details. In contrast, Gabor bases achieve spatial localization through Gaussian envelopes, enabling precise fitting of sharp corners without spillover noise. However, their localized envelopes may decay to near-zero responses for queries outside the effective support, leading to vanishing gradients.

Multi-scale Gabor encoding. To balance local precision and global coverage, we construct a multi-scale basis pyramid by partitioning the K learnable Gabor filters into three groups: fine, middle, and global levels. Fine-level bases use narrow Gaussian envelopes to capture highly localized sharp details. Global-level bases use broad envelopes with low effective bandwidth to provide spatially widespread responses, ensuring non-vanishing effective gradients throughout the domain. Middle-level bases bridge the spectrum and model transition regions, preventing either over-smoothing (too global) or ripple artifacts (too local).

Distance-aware hybrid sampling. To explicitly activate different scales during training, we employ a distance-aware hybrid sampling strategy that separates query points into near-field and far-field sets. We define the near-field set as

$$\Omega_{\text{near}} = \{p_i \pm \delta_i \mid p_i \in P,\ \|\delta_i\|_2 \le \delta_{\max}\}, \tag{5}$$

which concentrates samples in a narrow band around the observed surface to drive fine- and mid-level bases for detail refinement. In addition, we sample a far-field set uniformly in the normalized space to activate global-level bases and promote field continuity:

$$\Omega_{\text{far}} \sim \mathcal{U}([-1,1]^3). \tag{6}$$

We form the overall sampling domain as $\Omega = \Omega_{\text{near}} \cup \Omega_{\text{far}}$. In practice, DQGC (2) and the near-field correction step (7) are applied only on Ω_{near}, while the Eikonal regularization (9) is evaluated over Ω to stabilize the learned field in both near- and far-field regions.

D. Loss Function

Global Surface Alignment Loss. To align the implicit zero-level set $\mathcal{Z} = \{x \mid f_\theta(x) = 0\}$ with the input point cloud P, we use a Chamfer-style alignment loss based on a one-step *approximate near-field correction*. Unlike SDFs, UDFs are sign-ambiguous and may not form a valid distance metric during early optimization; therefore, we do not treat this update as an exact projection and apply it *only to near-field queries* $q \in \Omega_{\text{near}}$. Specifically, we perform a one-step correction:

$$q^{corr} = q - f_\theta(q) \cdot \frac{\nabla f_\theta(q)}{\|\nabla f_\theta(q)\|_2 + \epsilon}, \quad q \in \Omega_{\text{near}}. \tag{7}$$

We then minimize the bidirectional spatial discrepancy between the corrected query set Q^{corr} and the input point set P:

$$L_{cd} = \frac{1}{N} \sum_{q \in Q^{corr}} \min_{p \in P} \|q - p\|_2 + \frac{1}{M} \sum_{p \in P} \min_{q \in Q^{corr}} \|p - q\|_2, \tag{8}$$

where N and M denote the number of points in Q^{corr} and P, respectively.

Eikonal Gradient Regularization. To ensure the learned scalar field f_θ represents a physically valid distance-to-surface function rather than an arbitrary scalar distribution, we impose the Eikonal constraint. A fundamental property of Euclidean distance fields is that the gradient magnitude must be unity almost everywhere (i.e., $\|\nabla f(x)\|_2 = 1$), representing the unit speed of distance propagation. We enforce this property over the sampling space Ω via a quadratic penalty:

$$L_{reg} = \mathbb{E}_{q \in \Omega} \left[(\|\nabla f_\theta(q)\|_2 - 1)^2 \right]. \tag{9}$$

Overall Objective. Our final training objective is

$$L = \lambda_{cd} L_{cd} + \lambda_{reg} L_{reg} + \lambda_{dqgc} L_{\text{dqgc}}. \tag{10}$$

Here ϵ is a small constant for numerical stability.

IV. Experiments

A. Experimental Setup

Datasets. We evaluate on five diverse benchmarks: Thingi10k [18] for sharp watertight meshes, PCPNet [19] for open surfaces, 3DScene [20] for complex non-manifold scenes, and SRB [21] for real-world scans. Additionally, FAMOUS [15] is used for ablation studies.

Baseline Methods. We compare our method with representative approaches for implicit surface reconstruction. For UDF-based topology-agnostic reconstruction, we include NDF [8], CAP [9], and LevelSetUDF [17]. We also compare against additional implicit-field baselines that improve scalability or detail via alternative optimization/representation strategies, including GridPull (GP) [22], SAP [23], and Multi-Pull [24].

Evaluation Metrics. We employ Chamfer Distance (L2CD) to measure the mean squared error between reconstructed and ground truth points. F-Score (FS) evaluates structural completeness with a threshold τ; a lower threshold indicates a stricter standard. Normal Consistency (NC) calculates the cosine similarity of surface normals to assess geometric fidelity.

Implementation Details. Our network consists of a Gabor encoding layer followed by an 8-layer MLP (256 hidden units). The Gabor layer has $K = 64$ learnable basis functions. We use the Adam optimizer with an initial learning rate of 1×10^{-4}.

B. Comparison with State-of-the-Art

Performance on Watertight & Open Surface. Table I and Fig. 3 demonstrate our superior performance on challenging geometries. Quantitatively, our method consistently achieves the lowest L2CD and highest F-Score compared to state-of-the-art baselines. Our method reconstructs smoother surfaces with fewer artifacts.

Performance on Real-World Scans. Table II shows quantitative results on SRB, as illustrated in Fig. 4. Our method maintains low L2CD error while achieving a high NC of 0.961, proving the superior time-frequency locality of learnable Gabor encoding in handling noise.

Performance on Complex Scenes. Table III shows 3DScene results, as visualized in Fig. 5. Our method achieves an F-Score of 92.3%, significantly outperforming LevelSetUDF. This shows the effectiveness of the multi-scale Gabor design.

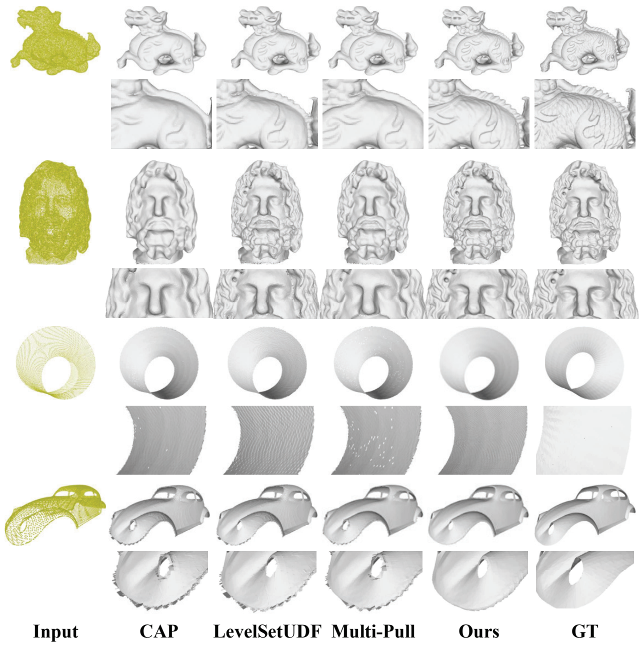

Fig. 3. Qualitative comparison on complex closed surfaces (rows 1–2) and challenging open or non-manifold geometries (rows 3–4).

TABLE I
WATERTIGHT & OPEN SURFACE RECONSTRUCTION COMPARISON ON THINGI10K AND PCPNET DATASETS

Method	L2CD$\times 10^3$	FS$^{0.05}$
CAP	0.0167	0.966
LevelSetUDF	0.0139	0.932
Multi-Pull	0.0122	0.957
Ours	**0.0073**	**0.972**

C. Ablation Study

Effectiveness of Learnable Gabor Encoding. As visualized in Fig. 6, we analyzed the internal representation by plotting the activation magnitude on a 2D slice. Unlike the global high-frequency response of Fourier bases, Gabor features exhibit a spatially sparse pattern, selectively activating only near geometric surfaces. From left to right: Fourier encoding, mid-scale Gabor encoding, and multi-scale Gabor encoding. Fourier features exhibit globally distributed responses, while Gabor features yield more localized activations near geometric boundaries. The multi-scale Gabor encoding produces the sparsest and sharpest boundary pattern.

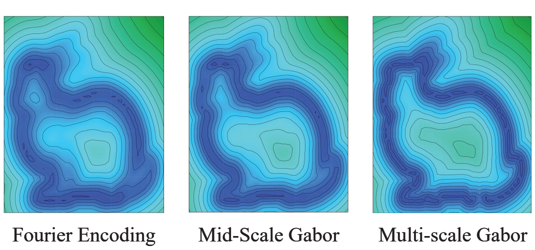

Fig. 6. Feature activation on a 2D slice.

Impact of Dual-Query Geometric Consistency. Removing the DQGC constraint caused a noticeable performance drop. As detailed in Table IV, the F-Score decreases by approximately 4% (comparing Var. B and Var. A). Qualitatively, as shown in Fig. 7, this leads to broken topologies in thin structures, confirming its necessity for locking unsupervised gradient directions.

Necessity of Multi-Scale Hierarchy. Fig. 7 illustrates that using only single-scale basis functions leads to either outlier artifacts (at small scales) or detail loss (at large scales), whereas our multi-scale strategy effectively balances reconstruction fidelity.

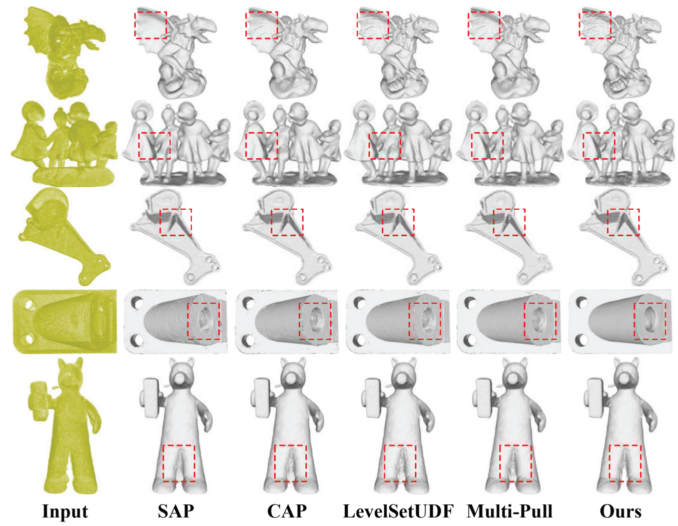

Fig. 4. Qualitative comparison on the SRB dataset of real-world scans.

D. Complexity & Robustness Analysis

Addressing concerns regarding potential computational overhead and parameter sensitivity introduced by the learnable Gabor layer, this section provides a detailed data analysis.

Spectral Efficiency and Model Capacity Analysis. To validate the superiority of learnable Gabor encoding over traditional Fourier encoding in feature representation, we designed a controlled variable experiment while keeping the MLP backbone identical. We compared parameter counts, inference latency, and reconstruction accuracy under different basis function counts (K) (Table V). The results demonstrate that Gabor encoding exhibits high spectral efficiency—achieving lower reconstruction error (L2-CD) than Fourier encoding even when the latter uses $4\times$ more basis functions.

Inference Efficiency. Our method introduces a learnable encoding module and additional geometric constraints. Consequently, the optimization time under per-shape overfitting increases compared to lightweight baselines, especially when gradient-based supervision is enabled. However, this overhead is strictly limited to the training phase. As shown in Table VI, the forward-only latency for evaluating a batch of query points ($B = 5000$) remains highly competitive (approx. 21.5ms), being comparable to CAP and faster than Multi-Pull. Unlike Multi-Pull, which requires iterative frequency refinement, our

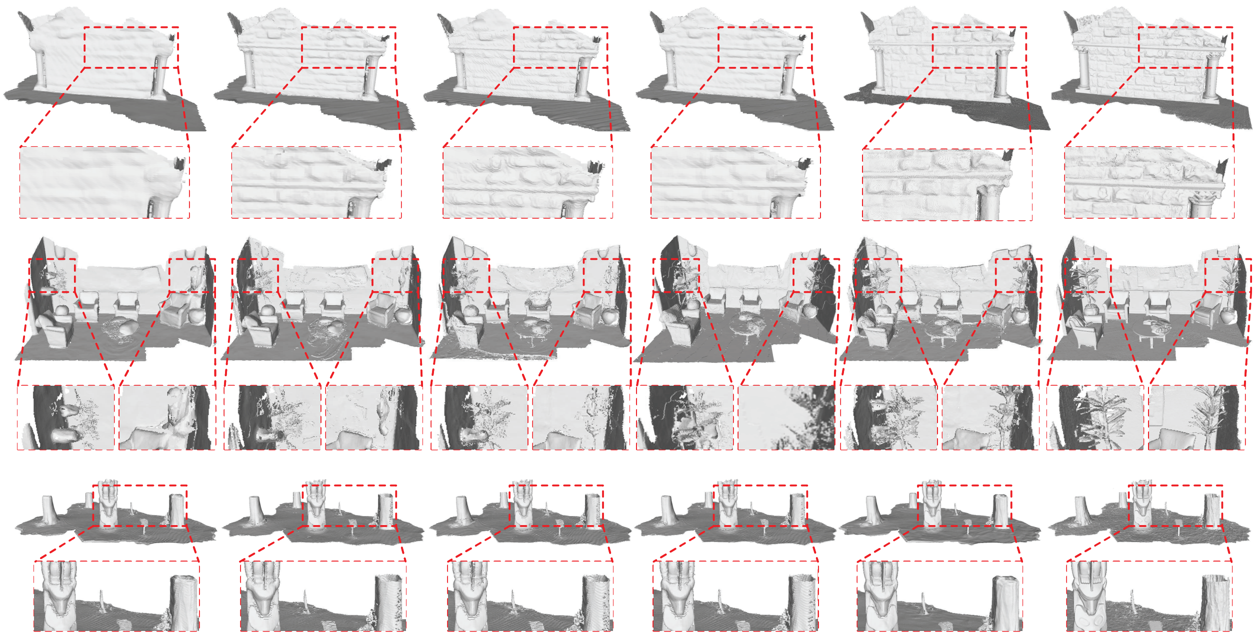

Fig. 5. Qualitative comparison on the 3DScene dataset of complex scenes.

TABLE II
REAL SCAN RECONSTRUCTION COMPARISON: CHAMFER-L2 (L2CD $\times 10^2$), F-SCORE (THRESHOLD=0.01) AND NORMAL CONSISTENCY (NC).

Method	Anchor			Daratech			Dc			Gargoyle			Lord quas		
	L2CD	FS	NC	L2CD	FS	NC	L2CD	FS	NC	L2CD	FS	NC	L2CD	FS	NC
SAP	0.0875	0.761	0.714	0.1075	0.701	0.642	0.1275	0.694	0.753	0.1132	0.811	0.717	0.1875	0.851	0.742
CAP	0.0433	0.981	0.832	0.0814	0.936	0.713	0.0901	0.913	0.811	0.0555	0.922	0.893	0.0730	0.924	0.812
LevelSetUDF	0.0462	0.985	0.832	0.0801	0.935	0.726	0.0921	0.927	0.816	0.0412	0.936	0.883	0.0760	0.922	0.821
Multi-Pull	0.0422	0.982	0.854	0.0623	0.946	0.753	0.0883	0.962	0.861	0.0328	0.961	0.899	0.0620	0.938	0.816
Ours	**0.0361**	**0.989**	**0.961**	**0.0621**	**0.958**	**0.755**	**0.0865**	**0.963**	**0.866**	**0.0311**	**0.966**	**0.911**	**0.0600**	**0.971**	**0.838**

TABLE III
SCENE RECONSTRUCTION COMPARISON (L2CD $\times 10^3$)

Method	L2CD	FS
PCPNet	0.0765	0.867
GP	0.0852	0.751
CAP	0.0501	0.913
LevelSetUDF	0.0597	0.912
Ours	**0.0421**	**0.923**

TABLE IV
ABLATION STUDY OF DIFFERENT COMPONENTS ON THINGI10K DATASET.

Model	Gabor	DQGC	Multi-Scale	L2-CD ↓	FS ↑
Base	(Standard Fourier)			0.0654	0.869
Var. A	✓			0.0522	0.901
Var. B	✓	✓		0.0415	0.942
Ours	✓	✓	✓	0.0312	0.966

pipeline maintains a concise single-pass inference path, ensuring real-time responsiveness during surface extraction.

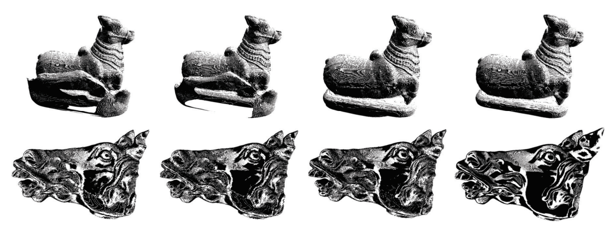

Fig. 7. Qualitative Ablation Study on FAMOUS dataset. Without DQGC, thin structures (the boundary between black and white) suffer from topological breaks. Single-scale models result in either high-frequency artifacts or detail loss. Our full model robustly reconstructs thin-walled regions with smooth surfaces.

TABLE V
SPECTRAL EFFICIENCY ANALYSIS. COMPARISON OF BASIS COUNTS (K) VS. PERFORMANCE ON FAMOUS DATASET.

Method	Bases (K)	Params (M)	Infer. Time (ms/batch)	L2CD ($\times 10^{-4}$) ↓	FS ($\tau = 0.01$) ↑
Baseline	64	0.41	18.2	0.0661	0.887
	128	0.43	18.8	0.0612	0.871
	256	0.46	19.9	0.0677	0.912
Ours	64	0.43	18.9	0.0412	0.936

TABLE VI

COMPUTATIONAL COMPLEXITY UNDER PER-SHAPE OVERFITTING. OPTIMIZATION TIME IS MEASURED FOR 80K ITERATIONS ON A SINGLE GPU. FORWARD LATENCY IS MEASURED WITH BATCH SIZE B=5000.

Method	Params (M)	Opt. Time (min)	Forward (ms)
CAP	0.41	145.3	18.9
LevelSetUDF	0.41	167.1	18.9
Multi-Pull	0.66	261.5	23.1
Ours	0.43	218.9	21.5

Sensitivity to Initialization. Gabor filters are often considered sensitive to initialization parameters. To verify robustness, we tested three initialization strategies: Uniform, Gaussian, and Xavier. Experimental results (Table VII) indicate that the fluctuation of the final F-Score is minimal ($< 0.5\%$). Regardless of initialization, the learnable spatial center μ_k and bandwidth $\mathbf{M}_k$ rapidly migrate to optimal geometric features driven by the gradient flow.

TABLE VII

ROBUSTNESS ANALYSIS OF GABOR INITIALIZATION STRATEGIES.

Init. Strategy	L2-CD (↓)	FS (↑)	NC (↑)
Uniform $\mathcal{U}(-1, 1)$	0.0412	0.936	0.961
Gaussian $\mathcal{N}(0, 1)$	0.0415	0.934	0.959
Xavier Normal	0.0411	0.937	0.962
Max Fluctuation	*¡ 1%*	*¡ 0.5%*	*¡ 0.5%*

V. CONCLUSION

In this paper, we present a Geometry-Aware Learnable Gabor Network to enhance unsupervised implicit field estimation. By utilizing the unique modulation properties of Gabor wavelets, our method dynamically adapts spectral bandwidth, effectively capturing high-frequency geometric details while mitigating spectral leakage in smooth regions. Combined with Dual-Query Geometric Consistency for gradient stabilization, our approach achieves robust reconstruction of complex non-manifold topologies. Extensive experiments validate the effectiveness of our modulation strategy in significantly improving reconstruction fidelity and efficiency compared to state-of-the-art methods. Future work will explore extending this adaptive spectral analysis to large-scale scene reconstruction and dynamic shape modeling.

REFERENCES

[1] Y. Lu, L. Wan, N. Ding, et al., "Unsigned orthogonal distance fields: An accurate neural implicit representation for diverse 3D shapes," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024, pp. 20551–20560.

[2] M. Berger, A. Li, J. A. Levine, et al., "Neural fields for implicit shape modeling: A survey," *Comput. Graph. Forum*, vol. 42, no. 4, pp. 1–25, 2023.

[3] M. Fainstein, V. Russo, et al., "DUDF: Differentiable unsigned distance fields with hyperbolic scaling," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024.

[4] Z. Li, T. Müller, A. Evans, R. H. Taylor, M. Unberath, M.-Y. Liu, and C.-H. Lin, "Neuralangelo: High-fidelity neural surface reconstruction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023.

[5] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 7462–7473.

[6] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 165–174.

[7] Y. Wang, P. Long, T. Wang, et al., "NeuS: Learning compact neural signed distance functions for high-quality 3D reconstruction," *ACM Trans. Graph.*, vol. 41, no. 4, pp. 1–15, 2022.

[8] J. Chibane, T. Alldieck, and G. Pons-Moll, "Neural unsigned distance fields for implicit function learning," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 21638–21652.

[9] J. Zhou et al., "CAP-UDF: Learning unsigned distance functions progressively from raw point clouds with consistency-aware field optimization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 12, pp. 7475–7492, 2024.

[10] B. Guillard, F. Stella, and P. Fua, "MeshUDF: Fast and differentiable meshing of unsigned distance field networks," in *Eur. Conf. Comput. Vis. (ECCV)*, 2022, pp. 576–592.

[11] A. Jacot, F. Gabriel, and C. Hongler, "Neural tangent kernel: Convergence and generalization in neural networks," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 31, 2018, pp. 8571–8580.

[12] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," in *Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 405–421.

[13] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Trans. Graph.*, vol. 41, no. 4, pp. 102:1–102:15, 2022.

[14] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 1–14, 2023.

[15] P. Erler, P. Guerrero, S. Ohrhallinger, N. J. Mitra, and M. Wimmer, "Points2Surf: Learning implicit surfaces from point clouds," in *Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 108–124.

[16] A. Gropp, L. Yariv, N. Haim, et al., "Implicit geometric regularization for learning shapes," in *Int. Conf. Mach. Learn. (ICML)*, 2020, pp. 3789–3799.

[17] J. Zhou, B. Ma, S. Li, et al., "Learning a more continuous zero level set in unsigned distance fields through level set projection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2023, pp. 3181–3192.

[18] Q. Zhou and A. Jacobson, "Thingi10K: A dataset of 10,000 3D-printing models," arXiv preprint arXiv:1605.04797, 2016.

[19] P. Guerrero, Y. Kleiman, M. Ovsjanikov, and N. J. Mitra, "PCPNet: Learning local shape properties from raw point clouds," *Comput. Graph. Forum* (Proc. Eurographics), vol. 37, no. 2, pp. 75–85, 2018.

[20] Q.-Y. Zhou and V. Koltun, "Dense scene reconstruction with points of interest," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 1–8, 2013.

[21] F. Williams, T. Schneider, C. Silva, D. Zorin, J. Bruna, and D. Panozzo, "Deep geometric prior for surface reconstruction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 10130–10139.

[22] C. Chen, Y.-S. Liu, and Z. Han, "GridPull: Towards scalability in learning implicit representations from 3D point clouds," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2023.

[23] S. Peng, C. Jiang, Y. Liao, et al., "Shape as points: A differentiable Poisson solver," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 34, 2021, pp. 13032–13044.

[24] T. Noda, C. Chen, W. Zhang, X. Liu, Y.-S. Liu, and Z. Han, "MultiPull: Detailing signed distance functions by pulling multi-level queries at multi-step," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2024.