VHEAT-UDF: REGULARIZING NEURAL UNSIGNED DISTANCE FIELDS WITH VECTOR HEAT DIFFUSION

Anonymous authors
Paper under double-blind review

ABSTRACT

We present VHeat-UDF, a method for learning unsigned distance fields regularized by vector heat diffusion. Given potentially inconsistent surface normals, vector heat diffusion propagates geometric orientation into the ambient space, providing effective regularization for UDF learning. To address singular behavior near field singularities, we introduce a rescaled vector Dirichlet energy that remains well-defined in such regions. VHeat-UDF provides accurate distance queries and robust surface reconstruction, even from incomplete or sparse inputs. Experiments comparing our approach with existing methods and scalar heat diffusion-based regularization demonstrate its effectiveness and robustness. Code will be released upon acceptance.

1 INTRODUCTION

Neural distance fields provide continuous, resolution-independent, and differentiable distance queries, making them a popular geometric representation across computer vision, graphics, and robotics for surface reconstruction Ma et al. (2022); Huang et al. (2023; 2022); Chibane et al. (2020), shape modeling Williams et al. (2021); Sellán et al. (2021), neural rendering Liu et al. (2020); Jiang et al. (2020), and collision detection Wang et al. (2024); Qu et al. (2025). Among these distance representations, unsigned distance fields (UDFs) encode the shortest distance to a surface without requiring an inside–outside distinction, making them more flexible than signed distance fields (SDFs), which typically assume watertight surfaces. UDFs are therefore well suited for representing open, self-intersecting, and non-manifold geometries Jones et al. (2006); Zhou et al. (2026).

Learning a neural UDF typically involves two coupled objectives: providing accurate distance queries throughout the spatial domain and preserving a reliable zero-level-set surface. However, achieving both distance accuracy and surface fidelity remains challenging. Existing methods often prioritize field regularization to improve surface reconstruction quality, which can compromise off-surface distance accuracy and limit their applicability to tasks that require reliable spatial distance queries for efficient computation, such as sphere tracing for implicit surface rendering Hart (1996); Vicini et al. (2022); Seyb et al. (2019), physics-based simulation Li et al. (2020); Teschner et al. (2005), and robotic collision evaluation Irshad et al. (2024); Qu et al. (2026).

To train a neural UDF that provides accurate spatial distance queries while maintaining a reliable surface boundary, even under sparse and incomplete observations, we introduce VHeat-UDF, a neural UDF learning method with vector heat diffusion-based regularization. Starting from potentially inconsistent surface normals, vector heat diffusion propagates surface orientation into the ambient space, providing geometric regularization for distance-field learning.

However, directly incorporating vector heat diffusion into UDF training can cause instability near field singularities, where the original vector Dirichlet energy becomes ill-conditioned; see Sec. 3.3 for details. To address this issue, we propose a rescaled vector Dirichlet energy that remains well defined near field singularities and enables stable vector heat diffusion regularization. During training, the diffusion regularization is progressively relaxed, allowing the final field to be optimized subject to the zero-surface boundary and Eikonal constraints and to converge to a high-accuracy UDF. Our contributions are threefold:

- We introduce a vector heat diffusion-based regularization for neural UDF learning that improves both spatial distance accuracy and zero-level-set surface continuity.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

- We provide a theoretical analysis showing that the vector heat formulation is invariant to normal sign inconsistencies, supporting robust learning from inconsistent orientations.
- We propose a rescaled vector Dirichlet energy that remains well defined near field singularities and enables stable vector heat diffusion regularization.

2 Related Work

Neural Distance Field Learning. Existing methods typically learn distance fields from surface observations using boundary constraints together with Eikonal regularization Atzmon & Lipman (2020a;b). Boundary constraints usually apply to observed surface points, while the Eikonal equation only constrains the local gradient magnitude. Directly optimizing the nonlinear Eikonal loss can lead to undesirable local minima, producing high-frequency ghost geometries Ben-Shabat et al. (2022); Wang et al. (2023). To improve distance-field accuracy and surface reconstruction quality, various regularization strategies have been proposed.

For SDF learning, the well-defined gradient across the surface enables normal supervision Gropp et al. (2020); Sitzmann et al. (2020). Beyond first-order constraints, second-order regularizations, including Laplacian-based and Hessian-based formulations, have been introduced to improve geometric accuracy and training stability Ben-Shabat et al. (2022); Yang et al. (2023); Wang et al. (2023); Zhang et al. (2022). Compared with SDF learning, UDF learning is more challenging due to the intrinsic gradient discontinuities across surfaces. Existing methods smooth such discontinuities through hyperbolic scaling Fainstein et al. (2024), incorporate normal-alignment cues Xu et al. (2025), or adopt an alternative squared-distance representation Yang et al. (2025). Some methods also employ hybrid sign and distance learning to handle geometries with complex topologies Wang et al. (2022); Kong et al. (2026), or introduce a specialized network architecture that limits the Lipschitz constant for the Eikonal condition Coiffier & Béthune (2024).

Another line of work supervises distance fields through projection consistency, where query points are expected to project onto the implicit surface using the predicted distance and gradient Ma et al. (2020); Marschner et al. (2023). Subsequent extensions enforce inter-level-set consistency and adapt projection-based supervision to UDF learning Li et al. (2025); Zhou et al. (2024; 2023).

Despite these advances, learning an accurate UDF over the spatial domain while maintaining the surface boundary remains challenging. Most regularization strategies developed for SDF are not directly applicable to UDF, while existing UDF methods, which primarily rely on near-surface constraints or projection-based supervision, often fail to construct reliable spatial distance fields.

Heat Diffusion for Distance Approximation. Heat diffusion establishes a fundamental connection between heat flow and distance through Varadhan's formula Varadhan (1967). Based on this relation, scalar heat methods have been widely used for geodesic distance approximation, either by solving backward-Euler linear systems Crane et al. (2013) or through variational formulations Belyaev & Fayolle (2015; 2020). More recently, scalar heat diffusion has been adopted as a regularizer for neural SDF learning, using Hopf–Cole transformation-based formulations Lipman (2021); Wang et al. (2025) or variational heat energies optimized with neural representations Weidemaier et al. (2026). This scalar heat-based regularization has also recently been extended to neural UDF learning Qu et al. (2026).

Vector heat diffusion extends heat transport from scalar to vector fields Sharp et al. (2019), producing diffused vectors that encode directions associated with shortest paths from the surface. By diffusing boundary orientations using a one-step backward Euler discretization and then approximating a distance field, vector heat diffusion has been applied to signed distance approximation on discrete geometric representations Feng & Crane (2024). A similar strategy has also been used for unsigned distance approximation on discrete spatial grids Kong et al. (2025). However, existing vector heat methods primarily focus on distance approximation over discrete representations rather than neural field learning, and formulating vector heat diffusion as a regularizer for neural UDFs remains nontrivial.

In this work, we investigate vector heat diffusion as a geometric regularizer for neural distance-field learning, derived from its variational formulation. We further compare scalar and vector heat diffusion and show that vector diffusion, by leveraging first-order directional information, provides more robust regularization for UDF learning.

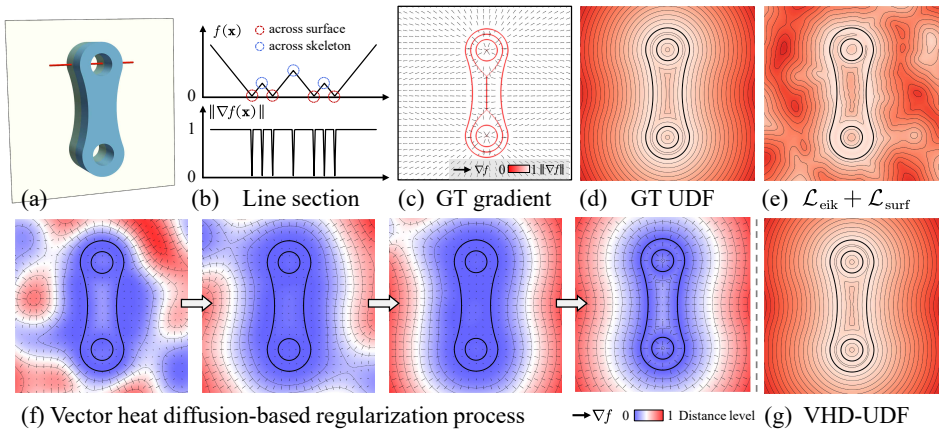

Figure 1: A UDF $f(\mathbf{x})$ has first-order discontinuities across surface boundaries and skeletons, which appear as distance-field singularities with vanishing gradients, i.e., $\|\nabla f(\mathbf{x})\| \approx 0$, illustrated by (a) the model, (b) a line section, and (c,d) slice planes. (e) Direct optimization with only zero-surface and Eikonal constraints cannot reconstruct a reliable distance field and produces spatial ghost geometries. The proposed vector heat diffusion (VHD)-based regularization significantly improves UDF learning: (f) it propagates surface orientations and suppresses ghost geometries during the diffusion process, finally yielding an accurate UDF in (g). See the supplementary video for the VHD-UDF learning process.

3 Method

3.1 UDF Formulation and Challenges

An unsigned distance field (UDF) $f : \Omega \to \mathbb{R}_{\geq 0}$ maps each query point $\mathbf{x} \in \Omega$ to its shortest distance from a surface $\mathcal{S} \subset \Omega$, represented by the zero level set of f. We denote the geometric skeleton by $\mathcal{K}$. The surface may be open, with boundary $\mathcal{C} = \partial\mathcal{S}$. Ideally, f satisfies

$$\|\nabla f(\mathbf{x})\| = 1, \quad \mathbf{x} \in \Omega \setminus (\mathcal{S} \cup \mathcal{K}), \qquad f(\mathbf{x}) = 0, \quad \mathbf{x} \in \mathcal{S}, \qquad \partial_{\mathbf{n}} f^{\pm}(\mathbf{x}) = \pm 1, \quad \mathbf{x} \in \mathcal{S}, \tag{1}$$

where $\mathbf{n}$ is the surface normal and $f^{\pm}$ denote the two one-sided limits across the surface. Neural UDF learning is commonly supervised by a zero-surface loss and an Eikonal loss,

$$\mathcal{L}_{\text{eik}} = \int_{\Omega \setminus \mathcal{S}} |\|\nabla f(\mathbf{x})\| - 1| \, d\mathbf{x}, \qquad \mathcal{L}_{\text{surf}} = \int_{\mathcal{S}} |f(\mathbf{x})|^2 \, d\mathbf{x}. \tag{2}$$

However, these constraints alone are insufficient. Although a UDF satisfies the Eikonal equation almost everywhere, it exhibits first-order discontinuities across surface boundaries and medial axes, where field singularities with vanishing gradients can arise, as shown in Fig. 1(b,c). Moreover, the nonlinear Eikonal equation admits many local minima. As a result, directly optimizing Eq. 2 can introduce spatial ghost geometries, which corrupt distance queries and degrade surface quality, as illustrated in Fig. 1(e). Therefore, additional regularization is needed to construct accurate spatial distance level sets while preserving the zero-level-set surface.

3.2 Vector Heat Diffusion-based UDF Regularization

We propose vector heat diffusion as a geometric regularizer for UDF learning in addition to zero-boundary and Eikonal constraints. Unlike scalar heat diffusion, which propagates zero-order heat values, vector heat diffusion propagates first-order surface orientations and directly regularizes UDF gradients. The vector heat diffusion regularizer is defined below, with a detailed derivation from its variational formulation provided in the Appendix Sec. A.

$$\mathcal{L}_{\text{vd}} = \mathcal{L}_{\text{vd}}^{\text{sm}} + \mathcal{L}_{\text{vd}}^{\text{bn}} = \frac{1}{2} \int_{\Omega \setminus \mathcal{S}} \left\| D\left(a(\mathbf{x}) \widehat{\nabla} f(\mathbf{x}) \right) \right\|^2 d\mathbf{x} + \frac{\lambda}{2} \int_{\mathcal{S}} \left(1 - \left| \widehat{\nabla} f(\mathbf{x}) \cdot \mathbf{n} \right| \right)^2 d\mathbf{x}, \tag{3}$$

where D denotes the covariant derivative of the vector field $\mathbf{u}(\mathbf{x})$, with $(D\mathbf{u})_{ij} = \partial u_i / \partial x_j$, and here $\mathbf{u}(\mathbf{x}) = a(\mathbf{x}) \widehat{\nabla} f(\mathbf{x})$ with $\widehat{\nabla} f = \nabla f / (\|\nabla f\| + \varepsilon)$ for unit-vector normalization. The function

$a(\mathbf{x}) \geq 0$ is the rescaled vector magnitude introduced in Sec. 3.3, and $\lambda > 0$ is the screening parameter. For this vector heat diffusion regularizer, we introduce two propositions: one showing that globally consistent normal orientation is not necessarily required, and the other ensuring stable regularization.

Proposition 1 (Consistency-free normals with minimum-level boundary condition)**.** *Let $f : \Omega \to \mathbb{R}$ be a UDF-like field bounded below by c, with the surface imposed as its minimum-level set:*

$$f(\mathbf{x}) \geq c, \qquad f(\mathbf{p}) = c, \quad \forall \mathbf{p} \in \mathcal{S}. \tag{4}$$

Then f is double-sided around $\mathcal{S}$ and determines only the unoriented normal line $\{\mathbf{n}, -\mathbf{n}\}$ on $\mathcal{S}$. Vector diffusion for distance learning can use surface normals that are not globally consistent.

The proof is provided in Appendix Sec. B.1, with experimental verification in Appendix Sec. C.3.1. This property allows us to use local surface normals, which can be easily estimated, e.g., by PCA, without enforcing global consistency. The boundary term $\mathcal{L}_{\text{vd}}^{\text{bn}}$ aligns the UDF gradient direction with the unoriented normal line $\{\mathbf{n}, -\mathbf{n}\}$, while the smoothness term $\mathcal{L}_{\text{vd}}^{\text{sm}}$ minimizes the vector Dirichlet energy of the rescaled gradient field. Intuitively, vector heat diffusion propagates surface orientation into the spatial domain. After diffusion, the vector at $\mathbf{x} \in \Omega \backslash \mathcal{S}$ aligns with the orientation transported from its closest surface point, revealing its relation to the distance field. This diffusion process is shown in Fig. 1(f), ultimately producing an accurate UDF in Fig. 1(g).

3.3 Rescaled Vector Dirichlet Energy

Minimizing the Dirichlet energy of the unnormalized gradient field alone favors a trivial solution with $\nabla f(\mathbf{x}) \approx \mathbf{0}$. Alternatively, optimizing a unit-normalized vector field can be unstable, since its vector Dirichlet energy $\mathcal{E} = \frac{1}{2}\int_{\mathcal{Q}} \|D(\widehat{\nabla f}(\mathbf{x}))\|^2 \, d\mathbf{x}$ becomes unbounded near field singularities Knöppel et al. (2013). This is problematic for UDF learning, where surface boundaries and skeletons inherently induce such singular regions.

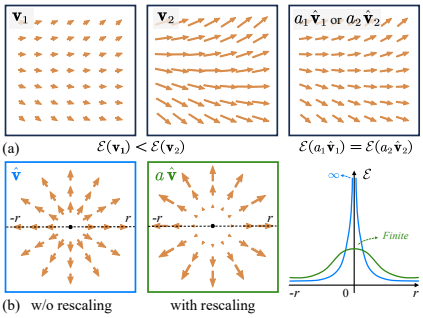

Figure 2: Rescaling removes magnitude bias in vector Dirichlet energy $\mathcal{E}(\mathbf{v}) = \frac{1}{2}\int_{\mathcal{Q}} \|D\mathbf{v}(\mathbf{x})\|^2 \, d\mathbf{x}$ and stabilizes regularization near field singularities.

As illustrated in Fig. 2(a), vector Dirichlet energy can be affected by vector magnitudes: two vector fields with identical smoothness may have different energies if their magnitudes differ. Direct unit normalization removes this magnitude bias but can cause infinite energy near singularities, as shown in Fig. 2(b). To address these issues, we introduce a rescaling strategy that automatically adapts vector magnitudes near field singularities while preserving the average vector scale over the sampled domain. This enables fair smoothness comparison and avoids energy blow-up caused by enforcing unit-length vectors everywhere, thereby stabilizing vector heat diffusion-based regularization.

Proposition 2 (Stable rescaled vector Dirichlet energy)**.** *Let $\mathcal{Q} \subset \Omega$ be a bounded evaluation domain with $0 < |\mathcal{Q}| < \infty$. For $\varepsilon > 0$, define the rescaled magnitude*

$$a(\mathbf{x}) = \frac{|\mathcal{Q}| \, \|\nabla f(\mathbf{x})\|}{\int_{\mathcal{Q}} \|\nabla f(\mathbf{y})\| \, d\mathbf{y} + \varepsilon}. \tag{5}$$

Then $\mathbf{v}(\mathbf{x}) = a(\mathbf{x})\widehat{\nabla f}(\mathbf{x})$ preserves gradient direction and has average magnitude approaching 1 as $\varepsilon \to 0$, while allowing vector scales to automatically shrink near field singularities. The rescaled field and its vector Dirichlet energy are invariant to positive global rescaling of f.

The proof is provided in Appendix Sec. B.2, with experimental verification in Appendix Sec. C.3.1. By using the rescaling strategy, vector diffusion propagates surface orientation while allowing vector magnitudes to vanish near field singularities, as observed in Fig. 1(f) and Fig. 15, enabling stable vector diffusion-based regularization for UDF learning.

3.4 Scalar versus Vector Heat Diffusion-based Regularization

Prior work Wang et al. (2025); Qu et al. (2026) has demonstrated the effectiveness of scalar heat regularization for distance-field learning. In this section, we adapt it to UDF learning, with derivations

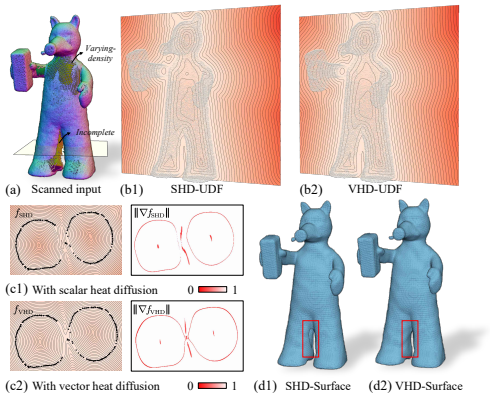

Figure 4: (a) Scanned data with varying density and incomplete input. Both (b1) scalar and (b2) vector heat diffusion effectively regularize UDF learning. VHD is more robust to incomplete input, repairing missing regions, as shown in zoom-in slices (c1,c2) and extracted surfaces (d1,d2).

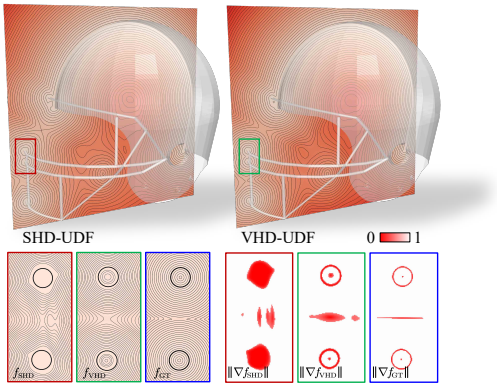

Figure 5: Vector heat diffusion regularization is more robust in small-scale regions. By leveraging first-order geometric information, VHD reconstructs interior distance level sets in narrow regions, whereas SHD treats these regions as thick surfaces.

provided in the Appendix Sec. A, and compare it with its vector counterpart. Theoretically, vector heat regularization leverages first-order information, which is naturally more robust than the zero-order values used in scalar heat methods. Similar advantages of leveraging first-order information have also been demonstrated in prior work Guennebaud & Gross (2007); Chen et al. (2013). Regarding UDF learning, vector heat is more robust with sparse, varying-density, and broken inputs, as well as small geometric features. As shown in Fig. 3, scalar heat may create point-to-point gaps with sparse inputs, whereas the vector formulation propagates distance level sets coherently along surface orientations through smooth vector interpolation. For varying-density and broken geometries, the vector formulation also produces more coherent distance fields.

This advantage is further validated on scanned data from the SRB dataset Williams et al. (2019), as shown in Fig. 4. Both scalar heat diffusion (SHD) and vector heat diffusion (VHD) regularize UDF learning, but VHD better handles varying density and incomplete regions. In particular, VHD learns a more continuous surface level set through smooth vector interpolation, thereby faithfully repairing broken regions, whereas SHD fails to recover missing structures.

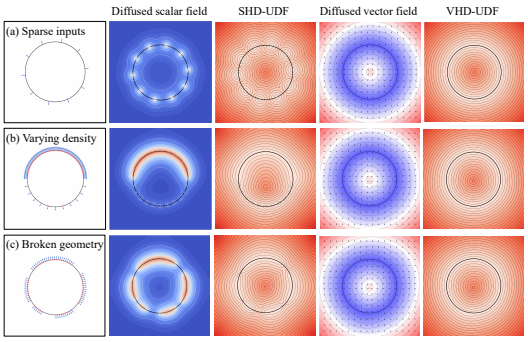

Figure 3: Vector heat diffusion regularization for UDF learning is more robust than scalar heat diffusion for (a) sparse inputs, (b) varying density, and (c) broken geometries.

Vector diffusion also performs better in small-feature regions by leveraging first-order geometric information. Scalar heat diffusion may treat nearby small-scale observations as heat sources, merging them into a thick surface and failing to recover interior distance levels. In contrast, VHD preserves local directional structure and reconstructs distance level sets inside narrow regions, as shown in Fig. 5. Its gradient norm $\|\nabla f_{\mathrm{VHD}}(\mathbf{x})\|$ remains more consistent with the Eikonal property, whereas SHD produces thick volumetric regions with $\|\nabla f_{\mathrm{SHD}}(\mathbf{x})\| \approx 0$.

4 Experiments

We evaluate our method on various models, focusing on spatial distance accuracy and surface reconstruction quality. Additional learning details and results, including extensions to SDF and generalized distance-field learning, are provided in Appendix Sec. C.

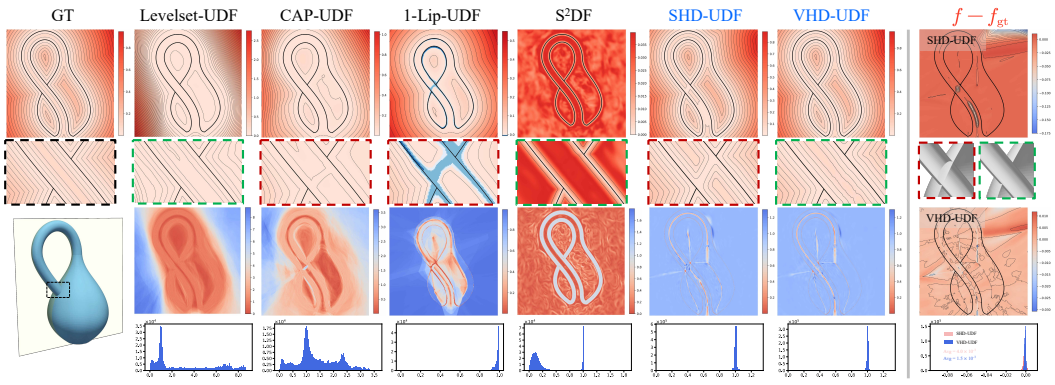

Figure 6: Field-slice comparison on the non-manifold Klein Bottle model. Top: learned UDF f. Second row: zoomed region, where green boxes indicate correct non-manifold structures and red boxes indicate failures. Third row: gradient norm maps $\|\nabla f\|$. Bottom: gradient norm statistics. The right side reports signed distance errors for SHD and VHD variants. Colorbars vary across figures to better show the results with high contrast.

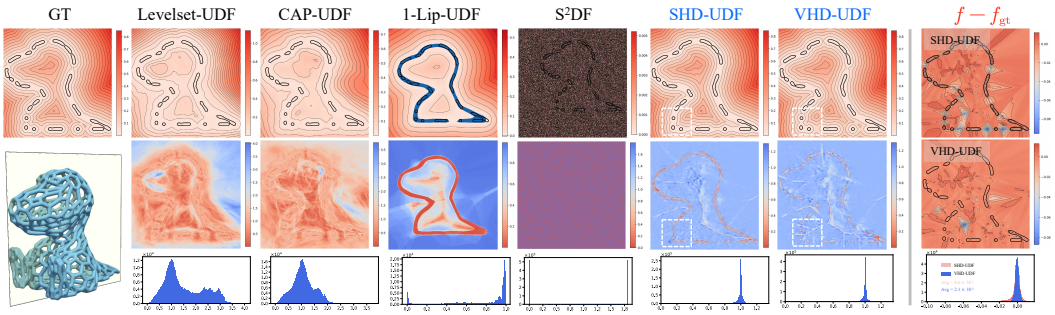

Figure 7: Field-slice comparison on the complex Voronoi-Dinosaur model. SHD-UDF fails in small-feature regions, while VHD-UDF remains accurate and robust, as highlighted by dashed boxes. Colorbars vary across figures to better show the results with high contrast.

Implementation details. We represent the UDF with a five-layer SIREN Sitzmann et al. (2020). Hidden layers use `Sine` activations, and the output layer uses `Softplus`, which is infinitely differentiable and ensures non-negative distance values with approximate zeros on surface. We optimize with Adam Kingma & Ba (2014) using a learning rate of 1×10^{-4}. **Datasets.** We evaluate distance accuracy on models with ground-truth references, including open-surface, non-manifold, and complex geometries from common graphics models, MGN Bhatnagar et al. (2019), and ShapeNet Chang et al. (2015). For surface reconstruction, we use scanned point clouds from SRB Williams et al. (2019), which contain missing geometries, and open-surface models from DeepFashion Zhu et al. (2020). We further validate generality and robustness on CAD models from ABC Koch et al. (2019) and sparse sketch inputs Huang et al. (2019). **Comparisons.** We compare with representative UDF learning methods, including LevelSetUDF Zhou et al. (2023), CAP-UDF Zhou et al. (2024), 1-Lipschitz-UDF Coiffier & Béthune (2024), and S^{2}DF Yang et al. (2025), in terms of distance accuracy and surface reconstruction quality. **Metrics.** All metrics are computed in the same normalized cube $[-1, 1]^3$. For UDF accuracy, we report the signed distance error, mean absolute distance error (ME) to the ground-truth mesh distance, Avg= $\mathrm{Mean}(|f - f_{\mathrm{gt}}|)$, and statistics of the UDF gradient norm to assess the Eikonal property. For surface reconstruction, we report Chamfer distance (CD $\times 10^{-2}$), Hausdorff distance (HD $\times 10^{-2}$), and F-score (%) at a threshold of 0.01, using default units unless explicitly noted. More details are provided in Appendix Sec. C.2.

4.1 UDF Accuracy

We show field slices on representative 3D models: a non-manifold Klein bottle in Fig. 6 and a complex Voronoi dinosaur Mehta et al. (2022) in Fig. 7, with additional results provided in Appendix

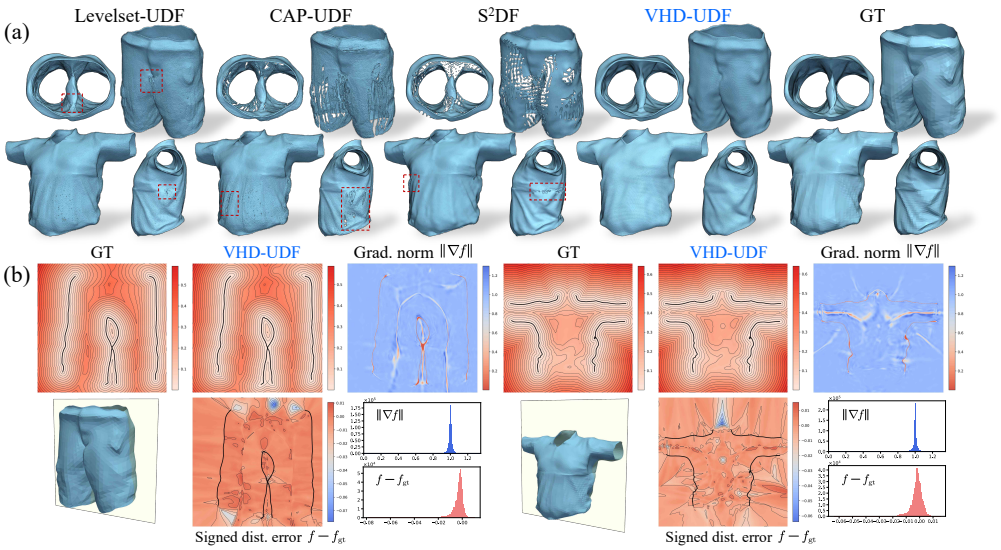

Figure 8: Results on MGN garments with open boundaries and intersections. (a) Surface reconstruction of different methods, with artifacts highlighted by boxes. (b) Field slices from VHD-UDF. Colorbars vary across figures to better show the results with high contrast; additional results are provided in the Appendix.

Sec. C. Since different methods use different normalization strategies and training domains, the models may appear at different scales in the slice visualizations. Colorbars also vary across figures to better display the results with high contrast.

The same-interval distance contours show that previous methods mainly capture a near-surface band but fail to recover accurate ambient-space distances. In contrast, both our scalar heat diffusion (SHD) and vector heat diffusion (VHD) variants produce meaningful distance fields throughout the space. Since previous methods exhibit substantially larger distance errors, we report signed distance error maps, $f - f_{\text{gt}}$, and mean absolute distance error, Avg$= \text{Mean}(|f - f_{\text{gt}}|)$, only for heat variants.

We further evaluate the gradient norm, which should satisfy the Eikonal condition $\|\nabla f(\mathbf{x})\| = 1$ in free space, except at surface boundaries and geometric skeletons where gradients vanish, i.e., $\|\nabla f(\mathbf{x})\| \approx 0$. The histograms of $\|\nabla f(\mathbf{x})\|$ show that only our methods closely match this constraint. In contrast, LevelSetUDF, CAP-UDF, and S^{2}DF deviate substantially from the expected Eikonal behavior except near the surface. 1-Lipschitz-UDF encodes the UDF as an SDF with a small margin, set to 1×10^{-3}, and bounds the gradient norm below 1 through Lipschitz networks. However, it still fails to recover accurate distance level sets and produces signed interior regions larger than the prescribed margin 1×10^{-3}.

For the Klein Bottle in Fig. 6, the zoomed region contains a non-manifold structure. Many methods, including SHD-UDF, fail to recover the correct non-manifold structure, whereas VHD-UDF succeeds. For the high-genus Dinosaur model in Fig. 7, the model contains small geometric features. S^{2}DF fails to reconstruct the near-surface field, while other baselines remain unsatisfactory. Consistent with Fig. 5, SHD-UDF struggles in small-scale regions, whereas VHD-UDF remains robust and accurate, as highlighted by the dashed boxes.

We further evaluate open-boundary garment models from MGN Bhatnagar et al. (2019). We report only VHD-UDF, since previous results already show that SHD-UDF is less robust than VHD-UDF. Fig. 8 shows results, and Table 1 reports reconstruction metrics. Additional field slices of more MGN examples are provided in the Appendix.

Table 1: Metrics on MGN garments using 80K random samples.

Method	Shorts			T-shirts		
	CD↓	HD↓	F-score↑	CD↓	HD↓	F-score↑
LevelSetUDF	2.06	5.30	48.16	1.04	2.72	93.50
CAP-UDF	2.32	7.67	45.19	1.05	**2.68**	92.98
S^{2}DF	2.18	17.42	73.61	0.96	3.83	98.33
VHD-UDF (ours)	**1.25**	**4.98**	**88.86**	**0.91**	3.14	**98.38**

On the Shorts model, CAP-UDF and S^{2}DF fail to construct a continuous surface level set and leave holes. LevelSetUDF produces better surfaces but still exhibits artifacts around intersections. Similar artifacts appear on the T-shirts model. In contrast, VHD-UDF handles both open boundaries and intersections more robustly, which is also reflected in the better overall metrics in Table 1. The field slices show that

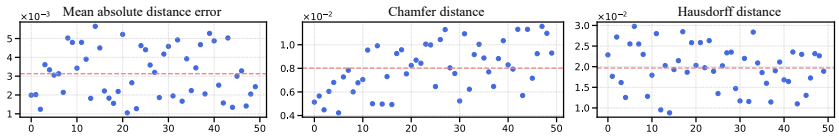

Figure 9: VHD-UDF metrics on 50 sampled ABC models. Dashed lines indicate averages.

Table 2: Metrics on the SRB dataset. Evaluation uses 100K randomly sampled points.

Method	Anchor			Daratech			DC			Gargoyle			Lord Quas		
	CD↓	HD↓	F-score↑	CD↓	HD↓	F-score↑	CD↓	HD↓	F-score↑	CD↓	HD↓	F-score↑	CD↓	HD↓	F-score↑
LevelSetUDF Zhou et al. (2023)	2.11	**21.01**	76.58	1.48	7.86	85.27	1.46	9.48	84.02	1.55	13.29	81.43	1.11	8.82	94.52
CAP-UDF Zhou et al. (2024)	2.22	22.68	76.48	1.50	8.09	85.04	1.45	8.92	84.18	1.55	13.17	81.17	1.08	8.63	94.72
S^2DF Yang et al. (2025)	2.17	22.18	76.08	1.46	8.02	85.97	1.47	9.15	85.39	1.52	13.19	**83.44**	1.07	7.92	95.40
VHD-UDF (ours)	**2.04**	21.99	**78.94**	**1.35**	**6.77**	**87.31**	**1.30**	**5.60**	**88.04**	**1.48**	**11.01**	82.95	**0.99**	**3.71**	**96.41**

gradient norms clearly reveal surfaces and skeletons, while distance errors remain small. Larger errors mainly occur near open boundaries, where exact distance behavior can be ambiguous. These results demonstrate that VHD-UDF effectively learns UDFs for geometries with open boundaries and intersections.

To further evaluate the robustness and generality of VHD-UDF, we test it on 50 sampled models from the ABC Koch et al. (2019) dataset. Fig. 9 reports the metrics, with slice visualizations provided in the Appendix. The mean absolute distance error, computed from 256^3 volumetric UDF queries against ground-truth point-to-mesh distances, averages 3.13×10^{-3} with respect to the $[-1,1]^3$ cube. Surface reconstruction also shows strong performance, with average CD and HD of 0.80×10^{-2} and 1.96×10^{-2}, respectively, using 100K random samples.

4.2 Surface Reconstruction

We evaluate UDF-based surface reconstruction on scanned datasets, including SRB Williams et al. (2019), which contains occlusion-induced missing regions, and DeepFashion Zhu et al. (2020), which further includes open boundaries. These experiments demonstrate that VHD-UDF learns high-fidelity distance fields and continuous surface level sets for repairing missing regions, and can reconstruct surfaces from sparse inputs, including sketch clouds Huang et al. (2019). Fig. 11 shows SRB reconstructions, with quantitative metrics reported in Table 2. In the highlighted regions, VHD-UDF consistently repairs incomplete surfaces through smooth vector interpolation. In contrast, LevelSetUDF and CAP-UDF rely on Chamfer-distance-based losses and cannot recover surface level sets in large missing regions from broken inputs. S^2DF also fails, as it focuses on dense near-surface fitting while suppressing non-near-surface regions. VHD-UDF is the only method that reconstructs faithful surface level sets and thereby repairs these regions, although it may smooth fine details, as discussed in Sec. 5. Fig. 11(b) further shows that VHD-UDF learns reliable distance fields and reconstructs continuous surface level sets in missing regions, as reflected in the gradient-norm maps.

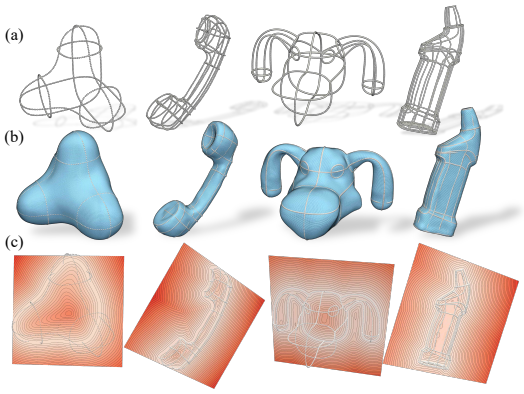

Figure 10: VHD-UDF on sparse sketch data. (a) Sketch point clouds, (b) reconstructed surfaces, and (c) learned UDFs.

Fig. 12 shows results on DeepFashion Zhu et al. (2020), which contains both occlusion-induced missing regions and open boundaries. VHD-UDF repairs missing regions while preserving open boundaries, as also visible in the field slices. In contrast, other methods fail to recover missing geometry and may incorrectly close open regions, as highlighted by the boxes.

As discussed in Sec. 3.4, vector heat diffusion is robust to sparse inputs. We further learn UDFs from sparse curve sketches. We first estimate point-cloud orientations using the method of Lin et al. (2024), then apply VHD-UDF to construct distance fields. Fig. 10 shows the learned UDFs and reconstructed surfaces.

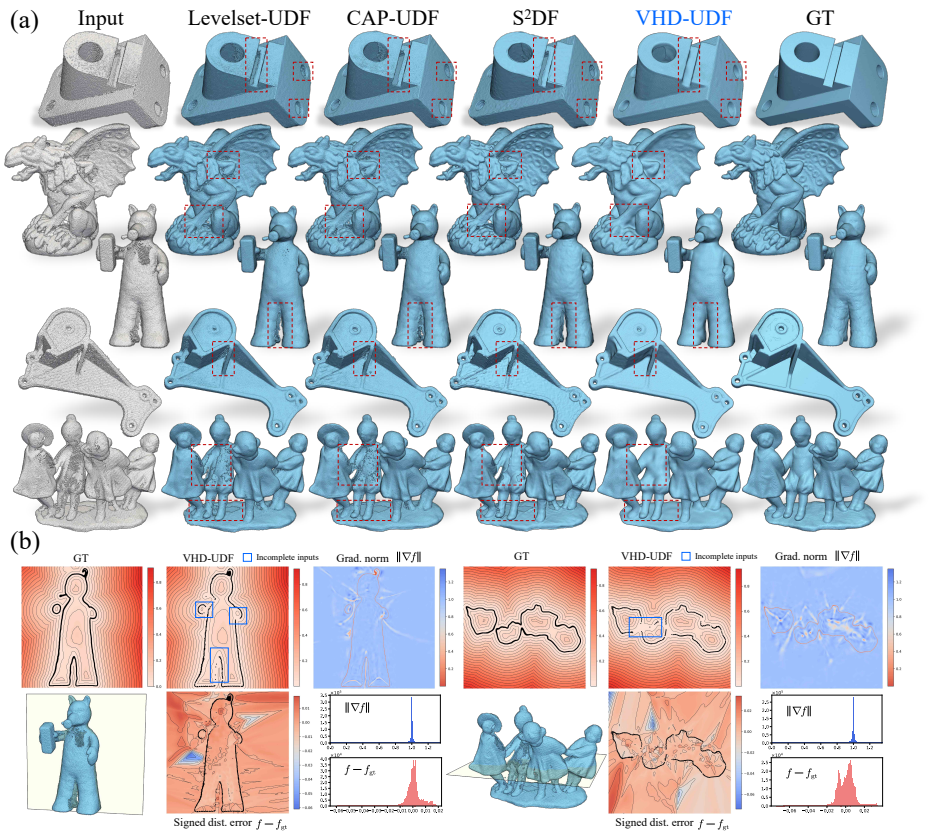

Figure 11: Results on the SRB dataset. (a) Surface reconstruction results. Highlighted boxes indicate regions where VHD-UDF consistently reconstructs faithful surface level sets, thereby repairing geometry missing from the input. (b) Field slices show that VHD-UDF learns accurate UDFs from incomplete inputs. Colorbars vary across figures to better display the results with high contrast; additional field-slice results are provided in the Appendix.

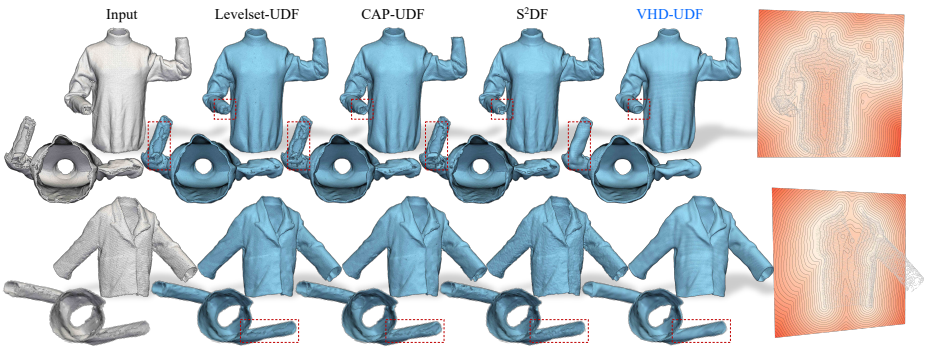

Figure 12: Surface reconstruction on the DeepFashion dataset. Highlighted boxes show regions where VHD-UDF consistently reconstructs faithful surface level sets, thereby repairing missing geometry while maintaining open-boundary regions. VHD-UDF slices are shown on the right.

5 Conclusion and Discussion

In this paper, we present neural UDF learning with vector heat diffusion-based regularization. Experiments on non-manifold, open-surface, sparse, and complex geometries show that VHD-UDF enables robust and accurate spatial distance-field learning and surface reconstruction. *Limitation.* A main limitation is that VHD-UDF may smooth surfaces, as it incorporates accurate spatial distance levels over the full domain rather than dense near-surface regularization. Possible solutions include hybrid low-/high-frequency representations for preserving surface detail Wang et al. (2021).

AI USE STATEMENT

In this work, we used generative AI tools primarily to assist with section formatting, language refinement, and the checking of mathematical proofs for two propositions. We did not use generative AI tools for method implementation or algorithm design. The authors take full responsibility for the final content of this work, including all text, claims, and artifacts produced with the assistance of generative AI.

REFERENCES

Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2565–2574, 2020a.

Matan Atzmon and Yaron Lipman. Sald: Sign agnostic learning with derivatives. *arXiv preprint arXiv:2006.05400*, 2020b.

Alexander Belyaev and Pierre-Alain Fayolle. An admm-based scheme for distance function approximation. *Numerical Algorithms*, 84(3):983–996, 2020.

Alexander G Belyaev and Pierre-Alain Fayolle. On variational and pde-based distance function approximations. In *Computer Graphics Forum*, volume 34, pp. 104–118. Wiley Online Library, 2015.

Yizhak Ben-Shabat, Chamin Hewa Koneputugodage, and Stephen Gould. Digs: Divergence guided shape implicit neural representation for unoriented point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19323–19332, 2022.

Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5420–5430, 2019.

Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.

Jiazhou Chen, Gaël Guennebaud, Pascal Barla, and Xavier Granier. Non-oriented mls gradient fields. In *Computer graphics forum*, volume 32, pp. 98–109. Wiley Online Library, 2013.

Julian Chibane, Gerard Pons-Moll, et al. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems*, 33:21638–21652, 2020.

Guillaume Coiffier and Louis Béthune. 1-lipschitz neural distance fields. In *Computer Graphics Forum*, volume 43, pp. e15128. Wiley Online Library, 2024.

Julian D Cole. On a quasi-linear parabolic equation occurring in aerodynamics. *Quarterly of applied mathematics*, 9(3):225–236, 1951.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (ToG)*, 32(5):1–11, 2013.

Miguel Fainstein, Viviana Siless, and Emmanuel Iarussi. Dudf: Differentiable unsigned distance fields with hyperbolic scaling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4484–4493, 2024.

Nicole Feng and Keenan Crane. A heat method for generalized signed distance. *ACM Transactions on Graphics*, 43(4), 2024.

Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.

Gaël Guennebaud and Markus Gross. Algebraic point set surfaces. In *ACM siggraph 2007 papers*, pp. 23–es. 2007.

John C Hart. Sphere tracing: A geometric method for the antialiased ray tracing of implicit surfaces. *The Visual Computer*, 12(10):527–545, 1996.

Eberhard Hopf. The partial differential equation. 1950.

Jiahui Huang, Hao-Xiang Chen, and Shi-Min Hu. A neural galerkin solver for accurate surface reconstruction. *ACM Transactions on Graphics (TOG)*, 41(6):1–16, 2022.

Jiahui Huang, Zan Gojcic, Matan Atzmon, Or Litany, Sanja Fidler, and Francis Williams. Neural kernel surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4369–4379, 2023.

Zhiyang Huang, Nathan Carr, and Tao Ju. Variational implicit point set surfaces. *ACM Transactions on Graphics (TOG)*, 38(4):1–13, 2019.

Muhammad Zubair Irshad, Mauro Comi, Yen-Chen Lin, Nick Heppert, Abhinav Valada, Rares Ambrus, Zsolt Kira, and Jonathan Tremblay. Neural fields in robotics: A survey. *arXiv preprint arXiv:2410.20220*, 2024.

Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1248–1258. IEEE, 2020.

Mark W Jones, J Andreas Baerentzen, and Milos Sramek. 3d distance fields: A survey of techniques and applications. *IEEE Transactions on visualization and Computer Graphics*, 12(4):581–599, 2006.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. Globally optimal direction fields. *ACM Transactions on Graphics (ToG)*, 32(4):1–10, 2013.

Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9601–9611, 2019.

Jiayi Kong, Chen Zong, Junkai Deng, Xuhui Chen, Fei Hou, Shiqing Xin, Junhui Hou, Chen Qian, and Ying He. Voronoi-assisted diffusion for computing unsigned distance fields from unoriented points. *arXiv preprint arXiv:2510.12524*, 2025.

Jiayi Kong, Xuhui Chen, Chen Zong, Fei Hou, Junhui Hou, Wenping Wang, and Ying He. Metric–phase fields: Decoupling distance and sign for thin-structure reconstruction from unoriented point clouds. *arXiv preprint arXiv:2605.25503*, 2026.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M Kaufman. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM transactions on graphics*, 2020.

Shengtao Li, Yudong Liu, Ge Gao, Ming Gu, and Yu-Shen Liu. I-filtering: Implicit filtering for learning neural distance functions from 3d point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Siyou Lin, Zuoqiang Shi, and Yebin Liu. Fast and globally consistent normal orientation based on the winding number normal consistency. *ACM Transactions on Graphics (TOG)*, 43(6):1–19, 2024.

Yaron Lipman. Phase transitions, distance functions, and implicit neural representations. *arXiv preprint arXiv:2106.07689*, 2021.

Shaohui Liu, Yinda Zhang, Songyou Peng, Boxin Shi, Marc Pollefeys, and Zhaopeng Cui. Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2016–2025. IEEE, 2020.

Marios Loizou, Melinos Averkiou, and Evangelos Kalogerakis. Learning part boundaries from 3d point clouds. In *Computer Graphics Forum*, volume 39, pp. 183–195. Wiley Online Library, 2020.

Baorui Ma, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Neural-pull: Learning signed distance functions from point clouds by learning to pull space onto surfaces. *arXiv preprint arXiv:2011.13495*, 2020.

Baorui Ma, Yu-Shen Liu, Matthias Zwicker, and Zhizhong Han. Surface reconstruction from point clouds by learning predictive context priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6326–6337, 2022.

Zoë Marschner, Silvia Sellán, Hsueh-Ti Derek Liu, and Alec Jacobson. Constructive solid geometry on neural signed distance fields. In *SIGGRAPH Asia 2023 conference papers*, pp. 1–12, 2023.

Ishit Mehta, Manmohan Chandraker, and Ravi Ramamoorthi. A level set theory for neural implicit evolution under explicit flows. In *European conference on computer vision*, pp. 711–729. Springer, 2022.

Jiasheng Qu, Zhuo Huang, Dezhao Guo, Hailin Sun, Aoran Lyu, Chengkai Dai, Yeung Yam, and Guoxin Fang. Inf-3dp: implicit neural fields for collision-free multi-axis 3d printing. *ACM Transactions on Graphics (TOG)*, 44(6):1–18, 2025.

Jiasheng Qu, Zhikai Shen, Chenyu Xu, Hailin Sun, Chengkai Dai, Yuhu Guo, Junpeng Wang, Yeung Yam, and Guoxin Fang. Wildfab: Multi-axis 3d printing from models in the wild. *ACM Transactions on Graphics (TOG)*, pp. 1–19, 2026. To appear in SIGGRAPH Asia 2026.

Silvia Sellán, Noam Aigerman, and Alec Jacobson. Swept volumes via spacetime numerical continuation. *ACM Transactions on Graphics (TOG)*, 40(4):1–11, 2021.

Dario Seyb, Alec Jacobson, Derek Nowrouzezahrai, and Wojciech Jarosz. Non-linear sphere tracing for rendering deformed signed distance fields. *ACM Transactions on Graphics*, 38(6), 2019.

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. The vector heat method. *ACM Transactions on Graphics (TOG)*, 38(3):1–19, 2019.

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.

Matthias Teschner, Stefan Kimmerle, Bruno Heidelberger, Gabriel Zachmann, Laks Raghupathi, Arnulph Fuhrmann, M-P Cani, François Faure, Nadia Magnenat-Thalmann, Wolfgang Strasser, et al. Collision detection for deformable objects. In *Computer graphics forum*, volume 24, pp. 61–81. Wiley Online Library, 2005.

Sathamangalam R Srinivasa Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics*, 20(2): 431–455, 1967.

Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Differentiable signed distance function rendering. *ACM Transactions on Graphics (ToG)*, 41(4):1–18, 2022.

Jingping Wang, Tingrui Zhang, Qixuan Zhang, Chuxiao Zeng, Jingyi Yu, Chao Xu, Lan Xu, and Fei Gao. Implicit swept volume sdf: Enabling continuous collision-free trajectory generation for arbitrary shapes. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024.

Li Wang, Weikai Chen, Xiaoxu Meng, Bo Yang, Jintao Li, Lin Gao, et al. Hsdf: Hybrid sign and distance field for modeling surfaces with arbitrary topologies. *Advances in Neural Information Processing Systems*, 35:32172–32185, 2022.

Yifan Wang, Lukas Rahmann, and Olga Sorkine-hornung. Geometry-consistent neural shape representation with implicit displacement fields. In *International Conference on Learning Representations*, 2021.

Zimo Wang, Cheng Wang, Taiki Yoshino, Sirui Tao, Ziyang Fu, and Tzu-Mao Li. Hotspot: Signed distance function optimization with an asymptotically sufficient condition. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 1276–1286, 2025.

Zixiong Wang, Yunxiao Zhang, Rui Xu, Fan Zhang, Peng-Shuai Wang, Shuangmin Chen, Shiqing Xin, Wenping Wang, and Changhe Tu. Neural-singular-hessian: Implicit neural representation of unoriented point clouds by enforcing singular hessian. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023.

Samuel Weidemaier, Florine Hartwig, Josua Sassen, Sergio Conti, Mirela Ben-Chen, and Martin Rumpf. Sdfs from unoriented point clouds using neural variational heat distances. In *Computer Graphics Forum*, pp. e70296. Wiley Online Library, 2026.

Francis Williams, Teseo Schneider, Claudio Silva, Denis Zorin, Joan Bruna, and Daniele Panozzo. Deep geometric prior for surface reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10130–10139, 2019.

Francis Williams, Matthew Trager, Joan Bruna, and Denis Zorin. Neural splines: Fitting 3d surfaces with infinitely-wide neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9949–9958, 2021.

Cheng Xu, Fei Hou, Wencheng Wang, Hong Qin, Zhebin Zhang, and Ying He. Details enhancement in unsigned distance field learning for high-fidelity 3d surface reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 8806–8814, 2025.

Chuanxiang Yang, Yuanfeng Zhou, Guangshun Wei, Long Ma, Junhui Hou, Yuan Liu, and Wenping Wang. Monge-ampere regularization for learning arbitrary shapes from point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Huizong Yang, Yuxin Sun, Ganesh Sundaramoorthi, and Anthony Yezzi. Steik: Stabilizing the optimization of neural signed distance functions and finer shape representation. *Advances in neural information processing systems*, 36:13993–14004, 2023.

Jingyang Zhang, Yao Yao, Shiwei Li, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Critical regularizations for neural surface reconstruction in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6270–6279, 2022.

Junsheng Zhou, Baorui Ma, Shujuan Li, Yu-Shen Liu, and Zhizhong Han. Learning a more continuous zero level set in unsigned distance fields through level set projection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 3181–3192, 2023.

Junsheng Zhou, Baorui Ma, Shujuan Li, Yu-Shen Liu, Yi Fang, and Zhizhong Han. Cap-udf: Learning unsigned distance functions progressively from raw point clouds with consistency-aware field optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):7475–7492, 2024.

Junsheng Zhou, Weiqi Zhang, Baorui Ma, Kanle Shi, Yu-Shen Liu, and Zhizhong Han. Udfstudio: A unified framework of datasets, benchmarks and generative models for unsigned distance functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.

Heming Zhu, Yu Cao, Hang Jin, Weikai Chen, Dong Du, Zhangye Wang, Shuguang Cui, and Xiaoguang Han. Deep fashion3d: A dataset and benchmark for 3d garment reconstruction from single images. In *European Conference on Computer Vision*, pp. 512–530. Springer, 2020.

A Heat Diffusion-based Regularization from Variational Formulations

We derive both scalar and vector heat diffusion-based regularizations for UDF learning from variational formulations of the heat PDEs.

A.1 Scalar Heat Diffusion-based Regularization

Consider the scalar heat equation

$$\begin{aligned} \partial_t u_t &= \Delta u_t, \quad \mathbf{x} \in \Omega \setminus \mathcal{S},\ t > 0, \\ \lim_{t \to 0} u_t &= \mathbb{1}_{\mathcal{S}}, \end{aligned} \tag{6}$$

where u_t is the heat field, Δ is the Euclidean Laplacian, and $\mathbb{1}_{\mathcal{S}}$ is the indicator of $\mathcal{S}$. A backward Euler step leads to a screened Poisson problem Crane et al. (2013), whose solution equivalently minimizes the variational screened Dirichlet energy

$$\mathcal{E}_{\text{scalar}}(u) = \frac{1}{2} \int_{\Omega} \|\nabla u\|^2 + \lambda (u - u_0)^2 \, d\mathbf{x}, \tag{7}$$

where u_0 encodes the initial or boundary heat source and $\lambda > 0$ is the screening parameter. The first term is the scalar Dirichlet energy, and the second term constrains u to be close to u_0. Using the Hopf–Cole transformation $u(\mathbf{x}) = \exp(-\sqrt{\lambda} f(\mathbf{x}))$ Hopf (1950); Cole (1951), the energy can be written in terms of the distance field $f(\mathbf{x})$ as

$$\mathcal{E}_{\text{scalar}}(f) = \frac{1}{2} \int_{\Omega \setminus \mathcal{S}} \left\| \nabla e^{-\sqrt{\lambda} f(\mathbf{x})} \right\|^2 d\mathbf{x} + \frac{\lambda}{2} \int_{\mathcal{S}} \left(e^{-\sqrt{\lambda} f(\mathbf{x})} - 1 \right)^2 d\mathbf{x}. \tag{8}$$

In our formulation of scalar heat diffusion-based regularization, this energy is used as the scalar heat regularizer,

$$\mathcal{L}_{\text{sd}} = \mathcal{L}_{\text{sd}}^{\text{sm}} + \mathcal{L}_{\text{sd}}^{\text{bn}} = \frac{1}{2} \int_{\Omega \setminus \mathcal{S}} \left\| \nabla e^{-\sqrt{\lambda} f(\mathbf{x})} \right\|^2 d\mathbf{x} + \frac{\lambda}{2} \int_{\mathcal{S}} \left(e^{-\sqrt{\lambda} f(\mathbf{x})} - 1 \right)^2 d\mathbf{x}, \tag{9}$$

together with other constraints, which are used in vector heat diffusion-based UDF learning.

A.2 Vector Heat Diffusion-based Regularization

Vector heat diffusion propagates vector data instead of scalar heat values. It is given by the vector heat PDE Sharp et al. (2019),

$$\begin{aligned} \partial_t \mathbf{u}_t &= \Delta^D \mathbf{u}_t, \quad \mathbf{x} \in \Omega \setminus \mathcal{S},\ t > 0, \\ \lim_{t \to 0} \mathbf{u}_t &= \mathbf{u}_0, \end{aligned} \tag{10}$$

where D is the covariant derivative and Δ^D is the connection Laplacian Sharp et al. (2019). In flat Euclidean space, the covariant derivative D reduces to the ordinary derivative, $(D\mathbf{u})_{ij} = \partial u_i / \partial x_j$, and the connection Laplacian reduces to the component-wise Laplacian, $\Delta^D \mathbf{u} = \left(\Delta u^1, \Delta u^2, \Delta u^3\right)^\top$. As $t \to 0$, the diffused vector at $\mathbf{x} \in \Omega \setminus \mathcal{S}$ aligns with the vector transported from its closest point on $\mathcal{S}$, thereby encoding distance-field information Feng & Crane (2024).

Similarly, a variational form of the one-step Euler method for vector heat diffusion is

$$\mathcal{E}_{\text{vector}}(\mathbf{u}) = \frac{1}{2} \int_{\Omega} \|D\mathbf{u}\|^2 + \lambda \|\mathbf{u} - \mathbf{u}_0\|^2 \, d\mathbf{x}. \tag{11}$$

Here, $\|D\mathbf{u}\|^2$ is the vector Dirichlet energy for the vector field $\mathbf{u}(\mathbf{x})$, and the screened term enforces consistency with the prescribed vector source $\mathbf{u}_0$. We instantiate the vector field using the UDF gradient, rescaled as $\mathbf{u}(\mathbf{x}) = a(\mathbf{x}) \widehat{\nabla f}(\mathbf{x})$, where the role of the magnitude $a(\mathbf{x})$ is explained in Appendix Sec. B.2. This gives the distance-field-based energy

$$\mathcal{E}_{\text{vector}}(f) = \frac{1}{2} \int_{\Omega \setminus \mathcal{S}} \left\| D\left(a(\mathbf{x}) \widehat{\nabla f}(\mathbf{x})\right) \right\|^2 d\mathbf{x} + \frac{\lambda}{2} \int_{\mathcal{S}} \left(1 - \left| \widehat{\nabla f}(\mathbf{x}) \cdot \mathbf{n} \right| \right)^2 d\mathbf{x}. \tag{12}$$

This yields the vector heat diffusion regularizer in Eq. 3, where the first term is the rescaled vector smoothness energy and the second term aligns with consistency-free surface orientation. This regularizer is optimized together with other constraints for UDF learning.

B Proofs of Propositions

B.1 Consistency-Free Normals with Minimum-Level Boundary Condition

We prove Proposition 1, with experimental verification in Sec. C.3.1. A UDF-like field is double-sided and orientation-free: it measures the same distance on both sides of the surface. Therefore, vector heat diffusion does not require globally consistent normals when the surface is imposed as the minimum-level set.

Let f be bounded below by c, with $f(\mathbf{x}) \geq c$ and $f(\mathbf{p}) = c$ for all $\mathbf{p} \in \mathcal{S}$. The standard UDF corresponds to $c = 0$,

$$d(\mathbf{x}, \mathcal{S}) = \inf_{\mathbf{y} \in \mathcal{S}} \|\mathbf{x} - \mathbf{y}\|, \qquad d|_{\mathcal{S}} = 0. \tag{13}$$

At a smooth point $\mathbf{p} \in \mathcal{S}$, choose any unit normal $\mathbf{n}$. For sufficiently small $r > 0$, define $\mathbf{x}^+ = \mathbf{p} + r\mathbf{n}$ and $\mathbf{x}^- = \mathbf{p} - r\mathbf{n}$. A UDF-like field satisfies $f(\mathbf{x}^+) = f(\mathbf{x}^-) = c + r$. Thus, flipping $\mathbf{n}$ only swaps the two sides. The field determines the unoriented normal line $\{\mathbf{n}, -\mathbf{n}\}$, rather than a globally oriented normal field.

This is sufficient for Eq. 3, whose boundary term aligns ∇f direction with the normal line rather than with an oriented normal, i.e., $1 - |\widehat{\nabla f} \cdot \mathbf{n}| = 1 - |\widehat{\nabla f} \cdot (-\mathbf{n})|$. This expression is invariant under $\mathbf{n} \mapsto -\mathbf{n}$, so arbitrary pointwise normal flips do not change the boundary alignment condition.

The minimum-level condition is essential to the double-sided distance-field structure. If f is differentiable and attains its lower bound on $\mathcal{S}$, then

$$\nabla f(\mathbf{p}) = \mathbf{0}, \qquad \mathbf{p} \in \mathcal{S}. \tag{14}$$

In the non-smooth UDF limit, the corresponding local behavior is $f(\mathbf{p} + t\mathbf{n}) = c + |t|$. In either case, the field grows away from its minimum on both sides of the surface.

Suppose that f is differentiable but the surface is not placed at its minimum level. A double-sided distance-like field would need to grow away from $\mathcal{S}$ in both normal directions. However, differentiability gives

$$\partial_{-\mathbf{n}} f(\mathbf{p}) = -\partial_{\mathbf{n}} f(\mathbf{p}). \tag{15}$$

Hence, if the field increases to first order in one normal direction, it decreases in the opposite direction, and the two directional derivatives cannot both be positive. A double-sided distance-like field must therefore place $\mathcal{S}$ at its minimum level, either with $\nabla f(\mathbf{p}) = \mathbf{0}$ in a smooth relaxation or with the limiting UDF form $f(\mathbf{p} + t\mathbf{n}) = c + |t|$.

Thus, the minimum-level condition provides the double-sided UDF structure, while the absolute normal-alignment term makes the diffusion boundary constraint invariant under normal flips. In contrast, a signed normal constraint, $1 - \widehat{\nabla f} \cdot \mathbf{n}$, distinguishes $\mathbf{n}$ from $-\mathbf{n}$ and therefore requires consistent normal orientations.

B.2 Rescaling Strategy for Vector Dirichlet Energy

We prove Proposition 2, with experimental verification in Sec. C.3.1. Write $g(\mathbf{x}) = \|\nabla f(\mathbf{x})\|$. Wherever $g > 0$, the unit direction $\widehat{\nabla f} = \nabla f / g$ satisfies

$$D\widehat{\nabla f} = \frac{D^2 f}{g} - \frac{\nabla f \otimes \nabla g}{g^2}, \tag{16}$$

using the convention $(D\mathbf{v})_{ij} = \partial v_i / \partial x_j$, with $(\mathbf{p} \otimes \mathbf{q})_{ij} = p_i q_j$. The local inverse-gradient factors can cause singular behavior near zeros of g, which correspond to field singularities. Adding a fixed positive constant to the pointwise denominator regularizes these factors but can still produce very large derivatives. We instead control the average vector magnitude and allow the magnitude to vanish at singularities. Let

$$I_f = \int_{\mathcal{Q}} g(\mathbf{x})\, d\mathbf{x}, \qquad C_{\mathcal{Q},f,\varepsilon} = \frac{|\mathcal{Q}|}{I_f + \varepsilon}. \tag{17}$$

For $\varepsilon > 0$, $C_{\mathcal{Q},f,\varepsilon}$ is finite and positive, including when $I_f = 0$. The proposed rescaling magnitude is

$$a(\mathbf{x}) = C_{\mathcal{Q},f,\varepsilon} g(\mathbf{x}) \geq 0. \tag{18}$$

For $g(\mathbf{x}) > 0$, the magnitude–direction product simplifies to

$$\mathbf{v}(\mathbf{x}) = a(\mathbf{x})\widehat{\nabla f}(\mathbf{x}) = C_{\mathcal{Q},f,\varepsilon} g(\mathbf{x}) \frac{\nabla f(\mathbf{x})}{g(\mathbf{x})} = C_{\mathcal{Q},f,\varepsilon} \nabla f(\mathbf{x}). \tag{19}$$

At $g(\mathbf{x}) = 0$, both sides are zero by the convention in Sec. 3.2. Thus Eq. 19 holds throughout $\mathcal{Q}$, and

$$\|\mathbf{v}(\mathbf{x})\| = a(\mathbf{x}), \qquad \frac{1}{|\mathcal{Q}|}\int_{\mathcal{Q}} \|\mathbf{v}(\mathbf{x})\|\, d\mathbf{x} = \frac{I_f}{I_f + \varepsilon} \le 1. \tag{20}$$

For $I_f > 0$, this average approaches 1 as $\varepsilon \to 0$. Since $C_{\mathcal{Q},f,\varepsilon} > 0$, the field preserves gradient directions and relative magnitudes. For a fixed smooth field, as $\mathbf{x}$ approaches a critical point, $g(\mathbf{x}) \to 0$ implies $a(\mathbf{x}) = \|\mathbf{v}(\mathbf{x})\| \to 0$, presenting that the vector magnitude shrinks automatically near field singularities.

The global factor $C_{\mathcal{Q},f,\varepsilon}$ depends on f and $\mathcal{Q}$, but not on the spatial evaluation coordinate. Wherever $g > 0$, and under the convention $(D\mathbf{v})_{ij} = \partial v_i / \partial x_j$, the componentwise product rule gives $\partial_j\big(a(\widehat{\nabla f})_i\big) = (\widehat{\nabla f})_i\, \partial_j a + a\, \partial_j (\widehat{\nabla f})_i$. Using this rule and substituting the identities above, along with $a = C_{\mathcal{Q},f,\varepsilon} g$, $\widehat{\nabla f} = \nabla f / g$, and $\nabla a = C_{\mathcal{Q},f,\varepsilon} \nabla g$, yields

$$\begin{aligned} D\mathbf{v} &= \widehat{\nabla f} \otimes \nabla a + a D\widehat{\nabla f} \\ &= C_{\mathcal{Q},f,\varepsilon}\left[\frac{\nabla f \otimes \nabla g}{g} + g\left(\frac{D^2 f}{g} - \frac{\nabla f \otimes \nabla g}{g^2}\right)\right] \\ &= C_{\mathcal{Q},f,\varepsilon} D^2 f. \end{aligned} \tag{21}$$

Since $D\mathbf{v} = C_{\mathcal{Q},f,\varepsilon} D^2 f$, the derivative contains no division by the local gradient magnitude $g = \|\nabla f\|$. Thus, vanishing gradients do not introduce numerical singularities into the rescaled Dirichlet energy, thereby avoiding energy blowup.

Finally, we show that the rescaling strategy makes the energy invariant to positive global rescaling of f, thereby eliminating magnitude bias in the measurement of vector-field smoothness. Consider the ideal normalization $\varepsilon = 0$ with $I_f > 0$. If $\widetilde{f} = kf + b$ for constants $k > 0$ and b, then $I_{\widetilde{f}} = kI_f$, $C_{\mathcal{Q},\widetilde{f},0} = C_{\mathcal{Q},f,0}/k$, and

$$\mathbf{v}_{\widetilde{f}} = C_{\mathcal{Q},\widetilde{f},0} \nabla \widetilde{f} = C_{\mathcal{Q},f,0} \nabla f = \mathbf{v}_f. \tag{22}$$

Hence, the rescaled field and its Dirichlet energy are invariant to positive global scaling, allowing the smoothness of the vector field to be measured without magnitude bias, as shown in Fig. 2(a).

Connection to Globally Optimal Direction Fields Knöppel et al. (2013). Our strategy is closely related to the relaxation principle used in globally optimal direction fields. In our notation, instead of optimizing a pointwise unit-norm vector field $\hat{\mathbf{v}}$ directly, one introduces a magnitude a and considers

$$\min_{\|\hat{\mathbf{v}}\|=1} \left(\min_{a \ge 0,\ \|a\|=1} \int_{\mathcal{Q}} \|D(a\hat{\mathbf{v}})\|^2\, d\mathbf{x} \right). \tag{23}$$

Then the problem is relaxed to

$$\min_{\|\mathbf{w}\|=1} \int_{\mathcal{Q}} \|D\mathbf{w}\|^2\, d\mathbf{x}, \tag{24}$$

with $\mathbf{w} = a\hat{\mathbf{v}}$, which leads to the eigenvalue problem $\Delta^D \mathbf{w} = \lambda \mathbf{w}$. The essential idea is to optimize a smooth vector-valued representative $\mathbf{w}$ with a global normalization, i.e. $\|\mathbf{w}\| = 1$, rather than enforcing pointwise unit norm during optimization. The global constraint prevents collapse to the zero field, while the amplitude $\|\mathbf{w}\|$ is free to shrink near singularities.

Our method follows the same spirit. We do not optimize $\hat{\mathbf{v}}(\mathbf{x}) = \nabla f / \|\nabla f\|$ directly. Instead, we optimize the rescaled representative

$$\mathbf{v} = a\widehat{\nabla f}, \tag{25}$$

where a provides a globally controlled but spatially adaptive magnitude, while maintaining the averaged vector magnitudes.

C More Experimental Details and Results

We provide additional implementation details and the computation of the metrics, followed by additional qualitative and quantitative comparative results. We also show extensions beyond UDF to SDF learning and generalized distance-field learning.

C.1 Open Boundary Constraint

Since the input geometry may contain open boundaries, we first collect boundaries using a point-cloud boundary detection algorithm Loizou et al. (2020), and then simply label open-boundary points based on the local distribution of neighboring points. To preserve these open boundaries, we constrain the distance-field gradient near open-boundary regions using a relaxed Neumann condition, as illustrated in Fig. 13.

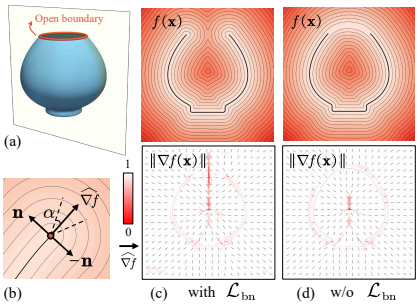

Figure 13: Effect of $\mathcal{L}_{\text{bn}}$ on open boundaries.

$$\mathcal{L}_{\text{bn}} = \int_{\mathcal{C}} \text{ReLU}\left(\left|\widehat{\nabla f}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})\right| - \sin\alpha\right) d\mathbf{x}, \quad (26)$$

where $\mathcal{C}$ denotes the open boundary, $\pm\mathbf{n}(\mathbf{x})$ is the boundary normal line, and $\alpha \in [0°, 180°]$ is the relaxation angle. We set $\alpha = 45°$ in all experiments. With this constraint, the learned surface preserves open boundaries and avoids artificial caps, as shown in Fig. 13(c). Without it, open regions may develop undesired cap-like extensions, as in Fig. 4(c1), or become incorrectly closed, as shown in Fig. 13(d).

C.2 Implementation Details and Metrics Computation

Hardware Information and Training Details. We implement our method in PyTorch. All experiments are conducted on a desktop PC equipped with an Intel Core i9-14900K CPU, an NVIDIA RTX 4090 GPU with 24 GB memory, and 64 GB RAM, running Ubuntu 22.04.4 LTS. We separate each diffusion regularizer into smoothness and boundary terms for flexible weighting during training. During training, diffusion regularization is used only in the early stage, together with other constraints. We use surface samples and off-surface samples at a 1:10 ratio. By default, 10% of the training schedule is used for the scalar or vector heat-diffusion stage, while the remaining epochs are used for distance-field optimization without diffusion regularization. The diffusion boundary weights, $w_{\text{sd}}^{\text{bn}}$ and $w_{\text{vd}}^{\text{bn}}$, are set large to approximate hard boundary constraints, while the smoothness weights, $w_{\text{sd}}^{\text{sm}}$ and $w_{\text{vd}}^{\text{sm}}$, are linearly decayed to zero. During the remaining 90% of training, we optimize only the Eikonal, zero-surface, and optional open-boundary losses to converge to an accurate distance field. For very sparse inputs, such as sketch data, we increase the diffusion stage to 20% of the total schedule, leaving 80% of the schedule without regularization.

Metrics Computation. Both the extracted surface point set $\mathcal{P}$ and ground truth point set $\mathcal{G}$ are recovered into the same $[-1, 1]^3$ cube, with the largest dimension spanning $[-1, 1]$ and the aspect ratio preserved. For evaluation, both point sets are randomly sampled to the same cardinality.

Let $d_{\mathcal{P}\to\mathcal{G}}(\mathbf{p}) = \min_{\mathbf{g}\in\mathcal{G}} \|\mathbf{p} - \mathbf{g}\|_2$ and $d_{\mathcal{G}\to\mathcal{P}}(\mathbf{g}) = \min_{\mathbf{p}\in\mathcal{P}} \|\mathbf{g} - \mathbf{p}\|_2$. We compute Chamfer and Hausdorff distances as

$$\text{CD} = \frac{1}{|\mathcal{P}|}\sum_{\mathbf{p}\in\mathcal{P}} d_{\mathcal{P}\to\mathcal{G}}(\mathbf{p}) + \frac{1}{|\mathcal{G}|}\sum_{\mathbf{g}\in\mathcal{G}} d_{\mathcal{G}\to\mathcal{P}}(\mathbf{g}).$$

$$\text{HD} = \max\left(\max_{\mathbf{p}\in\mathcal{P}} d_{\mathcal{P}\to\mathcal{G}}(\mathbf{p}), \max_{\mathbf{g}\in\mathcal{G}} d_{\mathcal{G}\to\mathcal{P}}(\mathbf{g})\right).$$

We additionally report the F-score with threshold 0.01, corresponding to 0.5% of the normalized cube side length.

Surface Point Localization from Learned UDF. Since these metrics are computed between point sets, we avoid biases from UDF meshing algorithms, which may introduce double layers, broken

Algorithm 1 Surface point localization

Input: Neural UDF f_θ, grid resolution R, M samples along a ray segment, near distance threshold d, convergence threshold τ, maximum number of iterations T, minimum ray length $\ell_{\min}$
Output: Extracted surface points $\mathcal{P}$
$\mathcal{S}(\mathbf{x}, \mathbf{v}, \ell, M)$: M samples on the line segment starting at $\mathbf{x}$, along direction $\mathbf{v}$, with length ℓ

initialize candidates $\mathcal{X}^0$ from an R^3 grid
$\mathcal{X}^0 \leftarrow \{\mathbf{x} \in \mathcal{X}^0 \mid f_\theta(\mathbf{x}) \le d\}$
$\mathcal{A} \leftarrow \mathcal{X}^0$ ▷ active candidates
for $t = 0, \dots, T-1$ **do**
  parallel for all $\mathbf{x}_k^t \in \mathcal{A}$:
    $\mathbf{v}_k^t \leftarrow -\nabla f_\theta(\mathbf{x}_k^t)/(\|\nabla f_\theta(\mathbf{x}_k^t)\| + \epsilon)$
    $\ell_k^t \leftarrow \max\left(2 f_\theta(\mathbf{x}_k^t), \ell_{\min}\right)$
    $\mathbf{x}_k^{t+1} \leftarrow \arg\min_{\mathbf{y} \in \mathcal{S}(\mathbf{x}_k^t, \mathbf{v}_k^t, \ell_k^t, M)} f_\theta(\mathbf{y})$
  $\mathcal{A} \leftarrow \{\mathbf{x}_k^{t+1} \mid \mathbf{x}_k^t \in \mathcal{A},\ f_\theta(\mathbf{x}_k^{t+1}) \ge \tau\}$
  break if $\mathcal{A} = \emptyset$
end for
$\mathcal{P} \leftarrow \{\mathbf{x}_k^{t+1} \mid \mathbf{x}_k^0 \in \mathcal{X}^0\}$
return $\mathcal{P}$

geometry, or difficulties with non-manifold structures. Instead of sampling from extracted meshes, we obtain surface points by back-tracing along the negative UDF gradient direction to the minimum UDF level and use these points for evaluation. The procedure is detailed in Algorithm 1. For a fair comparison, we report the best available results from either the original papers or our reproduced evaluations.

Specifically, let f_θ denote the learned UDF. We first initialize candidate points from an R^3 grid and discard candidates whose UDF value is larger than the near-surface threshold $d = 0.02$. The remaining candidates are refined in parallel. For each active candidate $\mathbf{x}_k^t$, we construct a local line segment along the inverse UDF-gradient direction, $\mathbf{v}_k^t = -\nabla f_\theta(\mathbf{x}_k^t)/(\|\nabla f_\theta(\mathbf{x}_k^t)\| + \epsilon)$. The segment length is dynamically updated as $\ell_k^t = \max(2 f_\theta(\mathbf{x}_k^t), \ell_{\min})$. We then sample M points on this segment and update the candidate to the sample with the minimum UDF value. Candidates whose UDF values fall below the convergence threshold τ are marked inactive, while the remaining candidates continue to be updated until convergence or the maximum number of iterations T. In our evaluation, we use $R = 256$, $T = 4$, $M = 100$, and $\tau = 0.001$ to obtain enough samples for metric evaluation.

C.3 Additional Experimental Results

In this section, we provide additional experiments to verify the two propositions. We also extend our proposed vector heat diffusion-based regularization from UDF learning to SDF learning and generalized distance-field learning. Finally, we present more results on UDF accuracy and surface reconstruction.

C.3.1 Experimental Verification of Propositions

Randomly flipped surface orientations. Proposition 1 theoretically shows that the double-sided UDF formulation does not require globally consistent surface normals. We further verify this property experimentally by randomly flipping the input surface normals while keeping all other settings unchanged. As shown in Fig. 14, the learned UDF and the extracted surface remain almost identical to those obtained with the globally consistent normals. CD is evaluated using 100K random samples. This confirms that vector heat diffusion-based UDF learning is robust to inconsistent surface orientations.

Rescaling strategy for vector Dirichlet energy. We experimentally verify Proposition 2 by comparing vector diffusion with and without the proposed rescaling strategy. As shown in Fig. 15, without rescaling, directly applying vector unit normalization causes the vector Dirichlet energy $\mathcal{L}_{\mathrm{vd}}^{\mathrm{sm}}$ to blow up and the Eikonal loss $\mathcal{L}_{\mathrm{eik}}$ to fail to converge. This leads to a noisy UDF and reconstructed surfaces with visible creases. In contrast, the proposed rescaling automatically reduces vector mag-

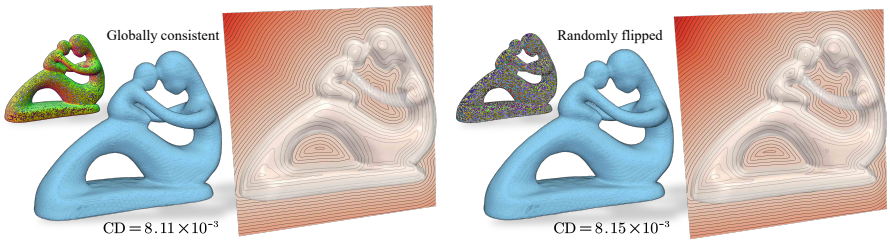

Figure 14: Random normal flips have little effect on UDF learning.

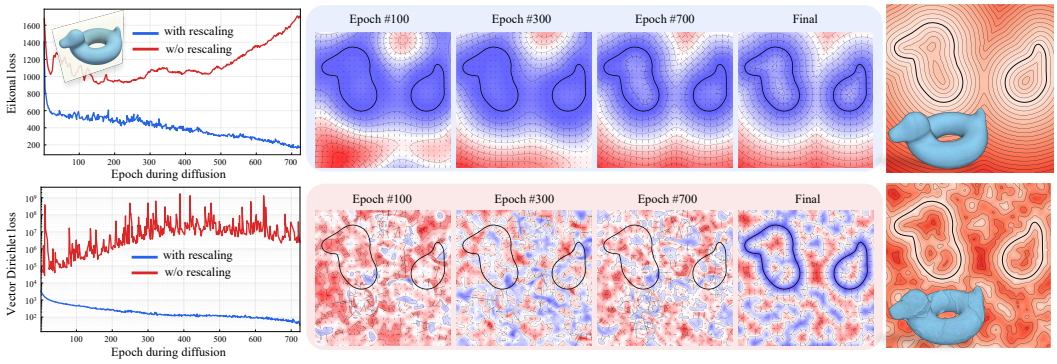

Figure 15: Effect of the rescaling strategy on vector diffusion regularization. Without the proposed rescaling strategy, directly applying vector unit normalization causes the vector Dirichlet energy to blow up and the Eikonal loss to fail to converge, producing a noisy UDF and creased surface reconstruction. With rescaling, vector magnitudes automatically shrink near singularity regions, stabilizing training and yielding an accurate UDF and high-quality surface reconstruction. See the supplementary video for the UDF learning process with and without the proposed rescaling strategy.

nitudes near singularity regions, such as surface boundaries and geometric skeletons. This stabilizes vector diffusion regularization and yields an accurate distance field. The resulting CD is 8.25×10^{-3} with rescaling and 2.31×10^{-2} without rescaling, evaluated using 100K random samples.

C.3.2 Extensions to SDF and Generalized Distance Fields

Although designed for UDF learning, our vector heat diffusion regularization also extends to other distance-field settings. We demonstrate two examples: neural SDF learning and generalized distance-field learning. Since these settings are not double-sided as in Proposition 1, they require globally consistent normals.

SDF learning with vector heat diffusion. For SDF learning, assuming that the input surface is watertight and globally oriented, we remove the non-negative `Softplus` output activation and use the same loss terms and weights as in UDF learning. This yields a neural SDF learned with vector heat diffusion regularization, as shown in Fig. 16.

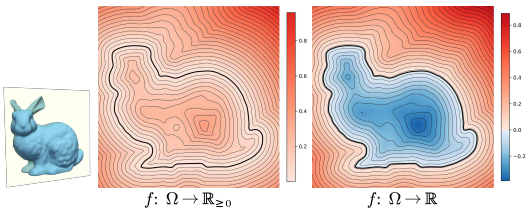

Figure 16: Extension to neural SDF learning.

Generalized distance-field learning without zero-level-surface boundary. The same framework can also learn generalized distance fields without enforcing a zero-level-surface boundary. We remove the zero-level-set constraint $\mathcal{L}_{\text{surf}}$ while retaining the vector diffusion boundary term $\mathcal{L}_{\text{vd}}^{\text{bn}}$ after vector diffusion regularization ends. As a result, different surface layers can have different level sets rather than being constrained to the same zero level set.

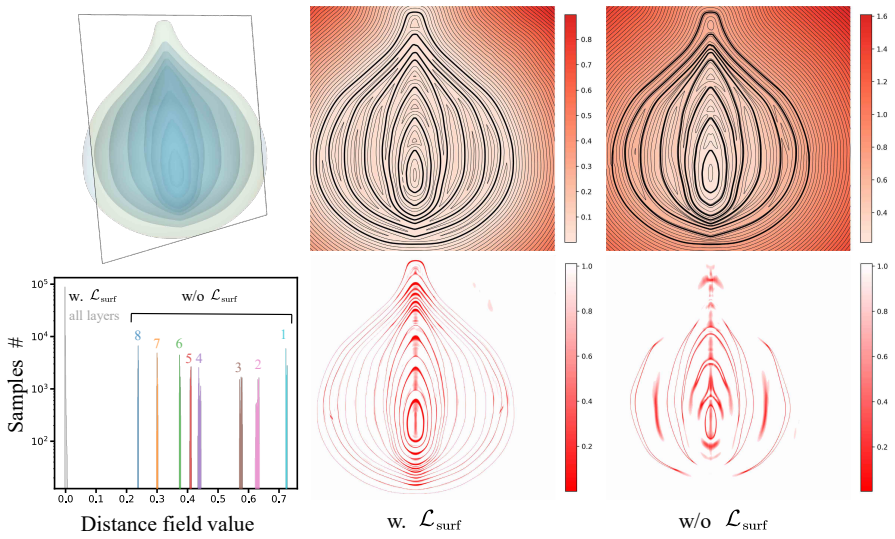

Figure 17: Generalized distance-field learning on nested surfaces. For an onion model with eight nested layers, standard UDF learning constrains all layers to the zero level set, leading to a single cluster of surface values and more spatial singularities in the gradient-norm field. By removing the zero-level-set constraint $\mathcal{L}_{\text{surf}}$ while retaining the vector diffusion boundary term $\mathcal{L}_{\text{vd}}^{\text{bn}}$, the generalized distance field separates the layers into eight distinct level-set clusters, producing a smoother field that better reflects the nested-layer structure.

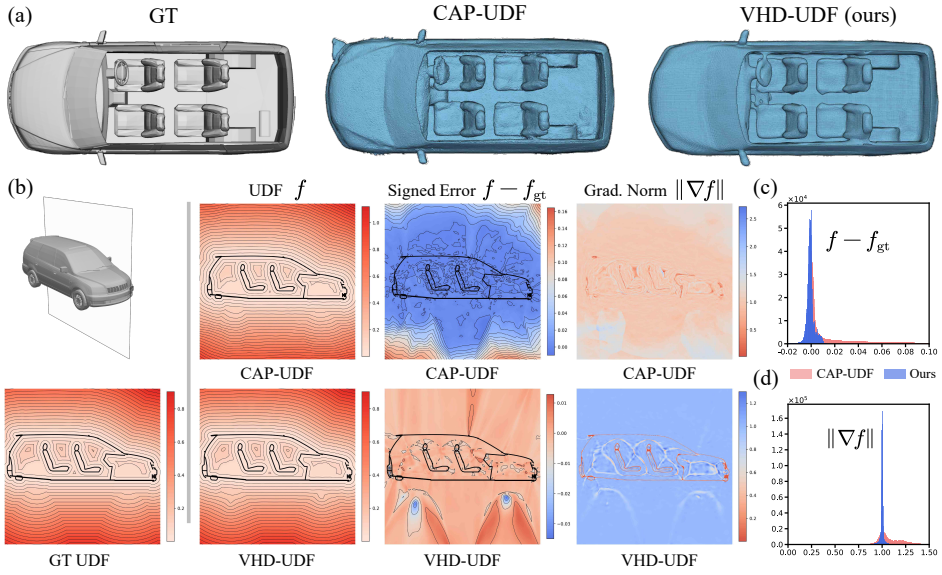

Figure 18: Results on a ShapeNet car with complex interior structures. (a) Surface reconstruction and (b) field slices: the ground truth is shown on the left, CAP-UDF in the upper middle row, and our VHD-UDF in the lower row. Histograms show (c) signed distance errors and (d) UDF gradient norms. Colorbars vary across figures to better show the results with high contrast.

Fig. 17 shows an onion model with eight nested layers. Standard UDF learning constrains all layers to the zero level set, producing a single cluster of surface values. For nested geometry, this constraint can introduce additional spatial singularities, as observed in the gradient-norm slices. In contrast, removing $\mathcal{L}_{\text{surf}}$ but preserving surface orientations through the boundary condition $\mathcal{L}_{\text{vd}}^{\text{bn}}$ separates the eight layers into distinct level-set clusters. The resulting field has fewer spatial singularities and better reflects the nested-layer structure.

C.3.3 Additional Experimental Results on UDF Accuracy and Surface Reconstruction

Models with complex interiors. We evaluate a ShapeNet car with complex interior structures and compare our VHD-UDF with CAP-UDF Zhou et al. (2024), the strongest baseline reported on this model, as shown in Fig. 18. Table 3 reports both surface reconstruction and distance-field metrics. VHD-UDF produces fewer artifacts and more accurate distance level sets. Its gradient norms better satisfy the Eikonal condition and clearly reveal surface boundaries and skeletons, while CAP-UDF shows larger deviations. Although the improvements in CD, HD, and F-score are modest, the mean distance error is substantially reduced from 1.87×10^{-2} to 0.24×10^{-2}, as also reflected in the signed error maps and histograms in Fig. 18.

Additional watertight models. Fig. 19 shows additional UDF accuracy results on watertight models with small geometric features, including Dragon and Happy Buddha. VHD-UDF remains more accurate than SHD-UDF, while other baselines fail to construct reliable ambient-space distance level sets. The gradient norm maps further show that our heat diffusion-based methods better satisfy the Eikonal condition and clearly reveal surface boundaries and geometric skeletons, where vanishing gradients appear.

Table 3: Metrics on the car model using 100K random samples. ME ($\times 10^{-2}$) denotes mean absolute distance error.

Method	CD↓	HD↓	F-score↑	ME↓
CAP-UDF	1.15	5.70	92.63	1.87
VHD-UDF	**1.13**	**4.70**	**93.14**	**0.24**

Additional open-surface models. Fig. 20 shows additional results on MGN models with open boundaries and intersections. Baseline methods fail to recover accurate ambient-space distance levels and deviate from the Eikonal condition, as shown by the gradient norm maps. In contrast, VHD-UDF produces accurate distance fields and handles open boundaries and intersections robustly. The main distance errors occur near open-boundary regions.

Gallery of ABC models. To further evaluate robustness and generality, we test VHD-UDF on additional ABC models. Fig. 21 provides field-slice visualizations, and Fig. 9 reports the metrics.

Field slices on SRB models. Field slices on SRB models are shown in Fig. 22. VHD-UDF reconstructs coherent surfaces from scanned data, learns faithful surface level sets that repair large missing regions, and recovers accurate spatial distance level sets. The main errors occur near broken regions.

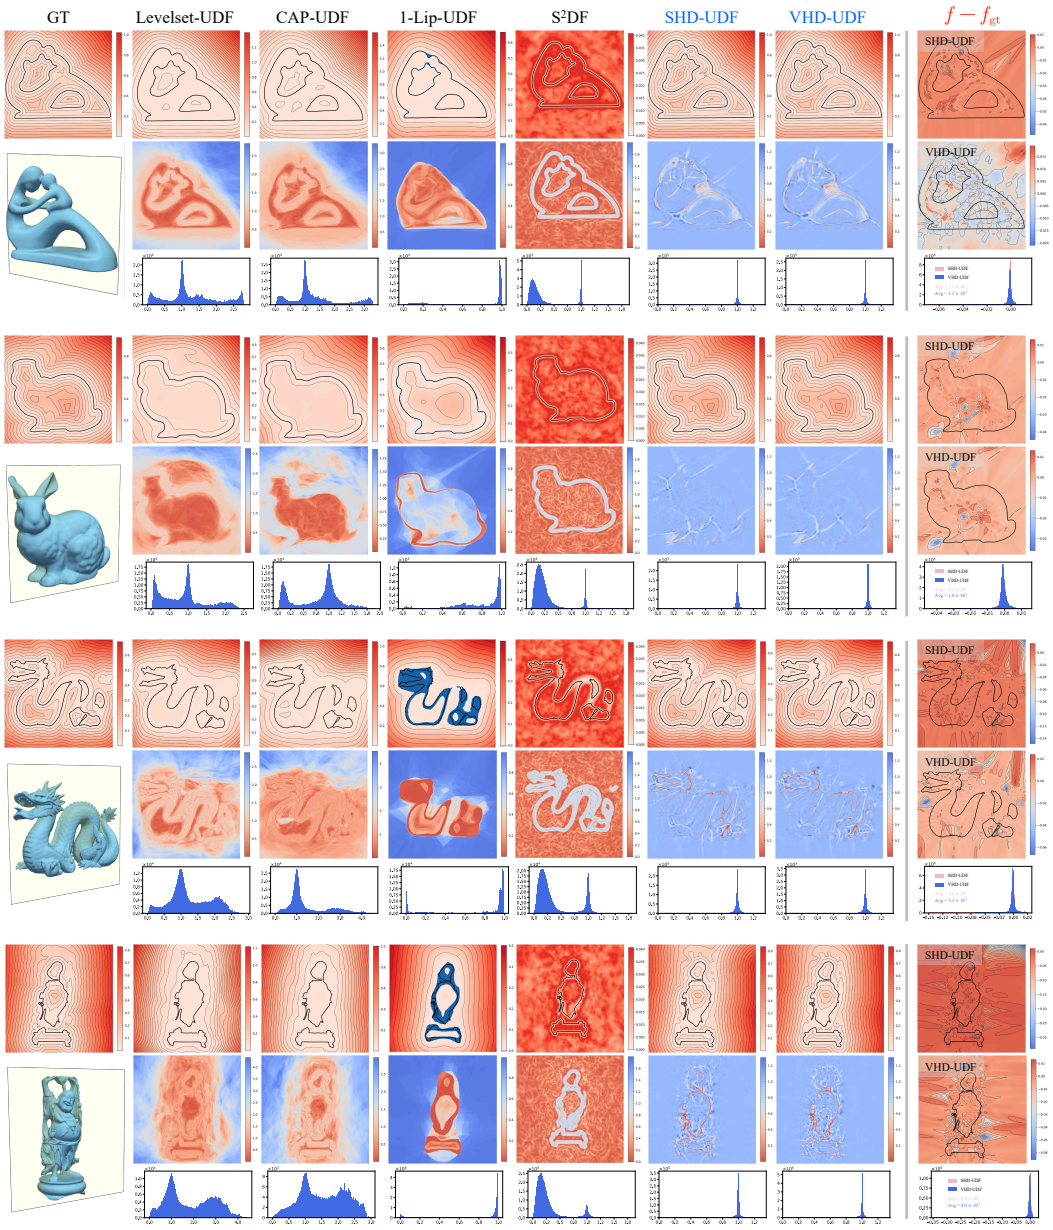

Figure 19: Field slices of watertight models. Top: learned UDF f. Second row: gradient norm maps $\|\nabla f\|$. Bottom: gradient norm statistics. The right side reports signed distance errors for our SHD and VHD variants. Colorbars vary across figures to better show the results with high contrast.

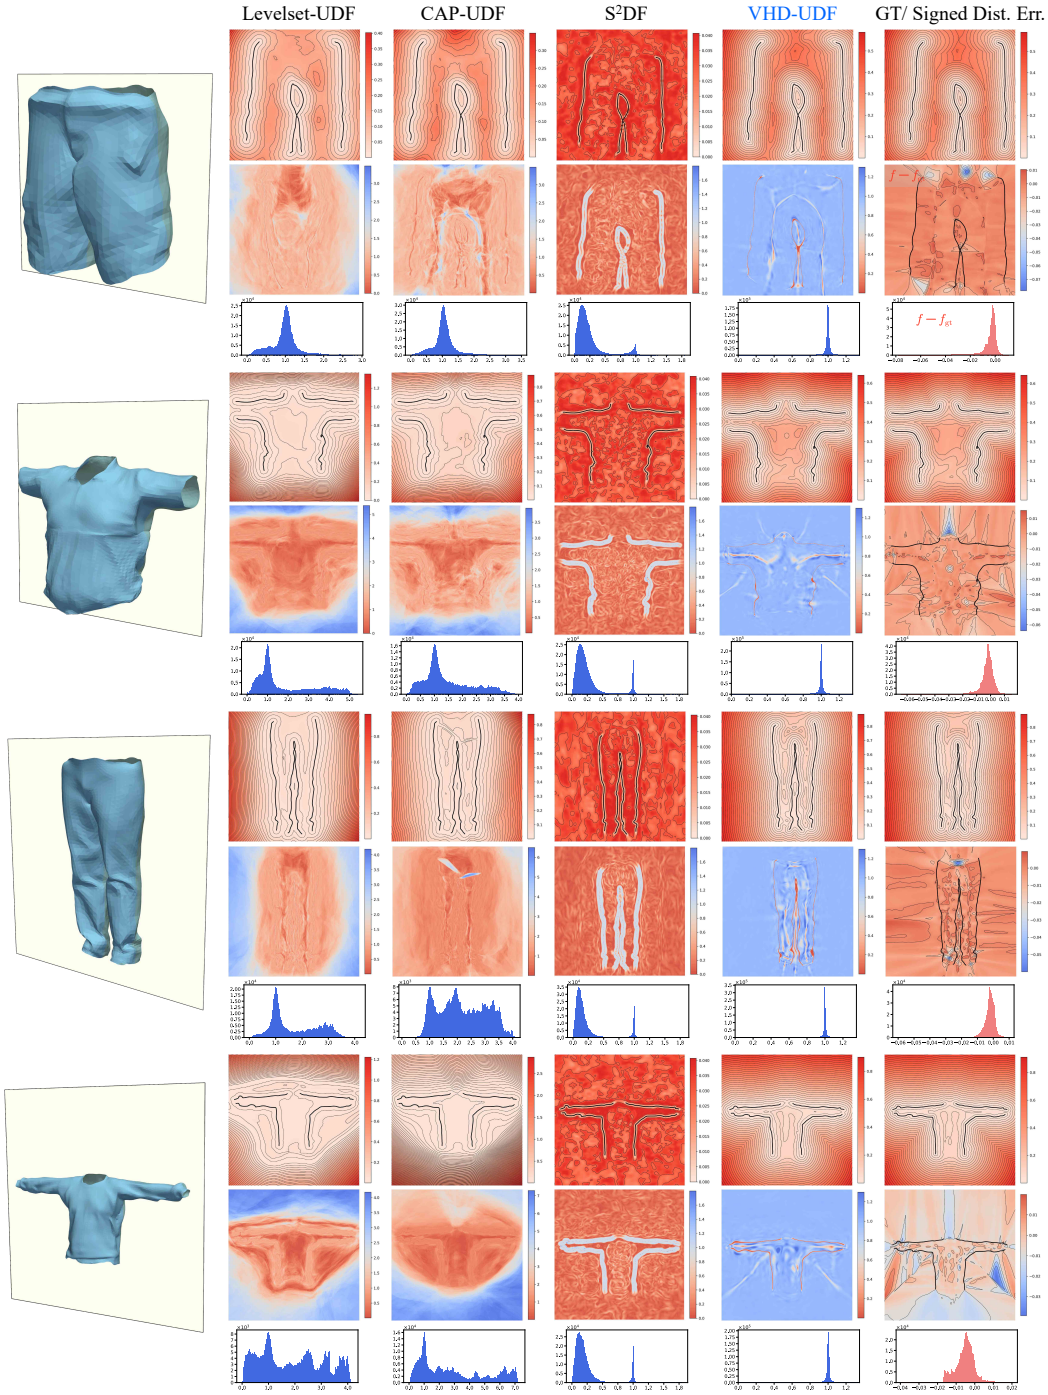

Figure 20: Field slices of MGN models. Top: learned UDF f. Second row: gradient norm maps $\|\nabla f\|$. Bottom: gradient norm statistics. The right side shows the ground-truth UDF and signed distance error map of VHD-UDF. Colorbars vary across figures to better show the results with high contrast.

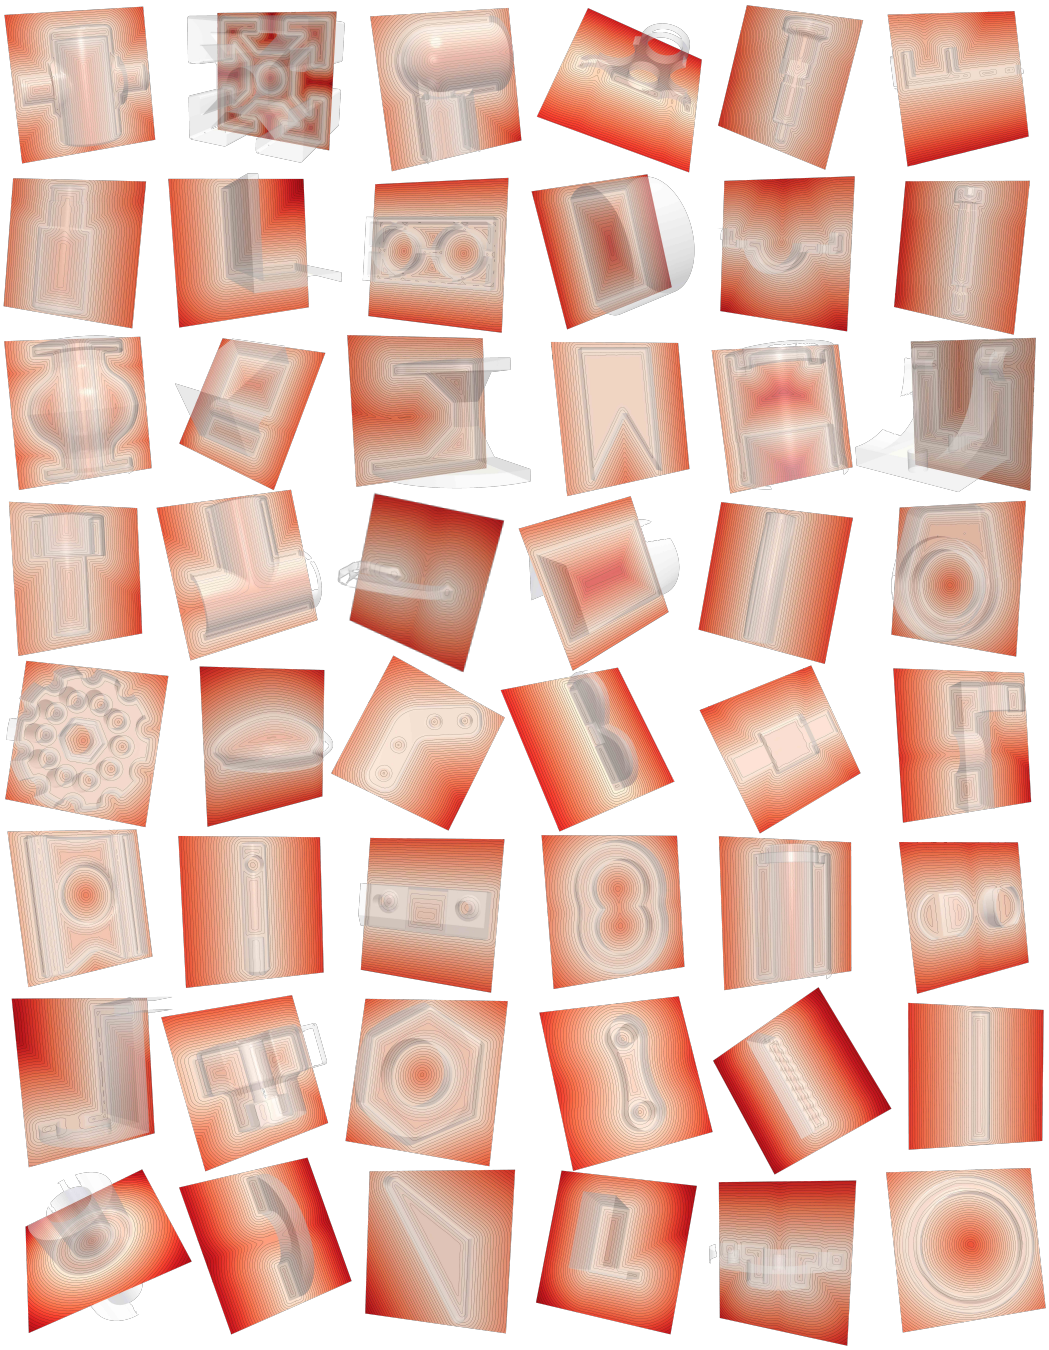

Figure 21: Gallery of VHD-UDF field slices on ABC models.

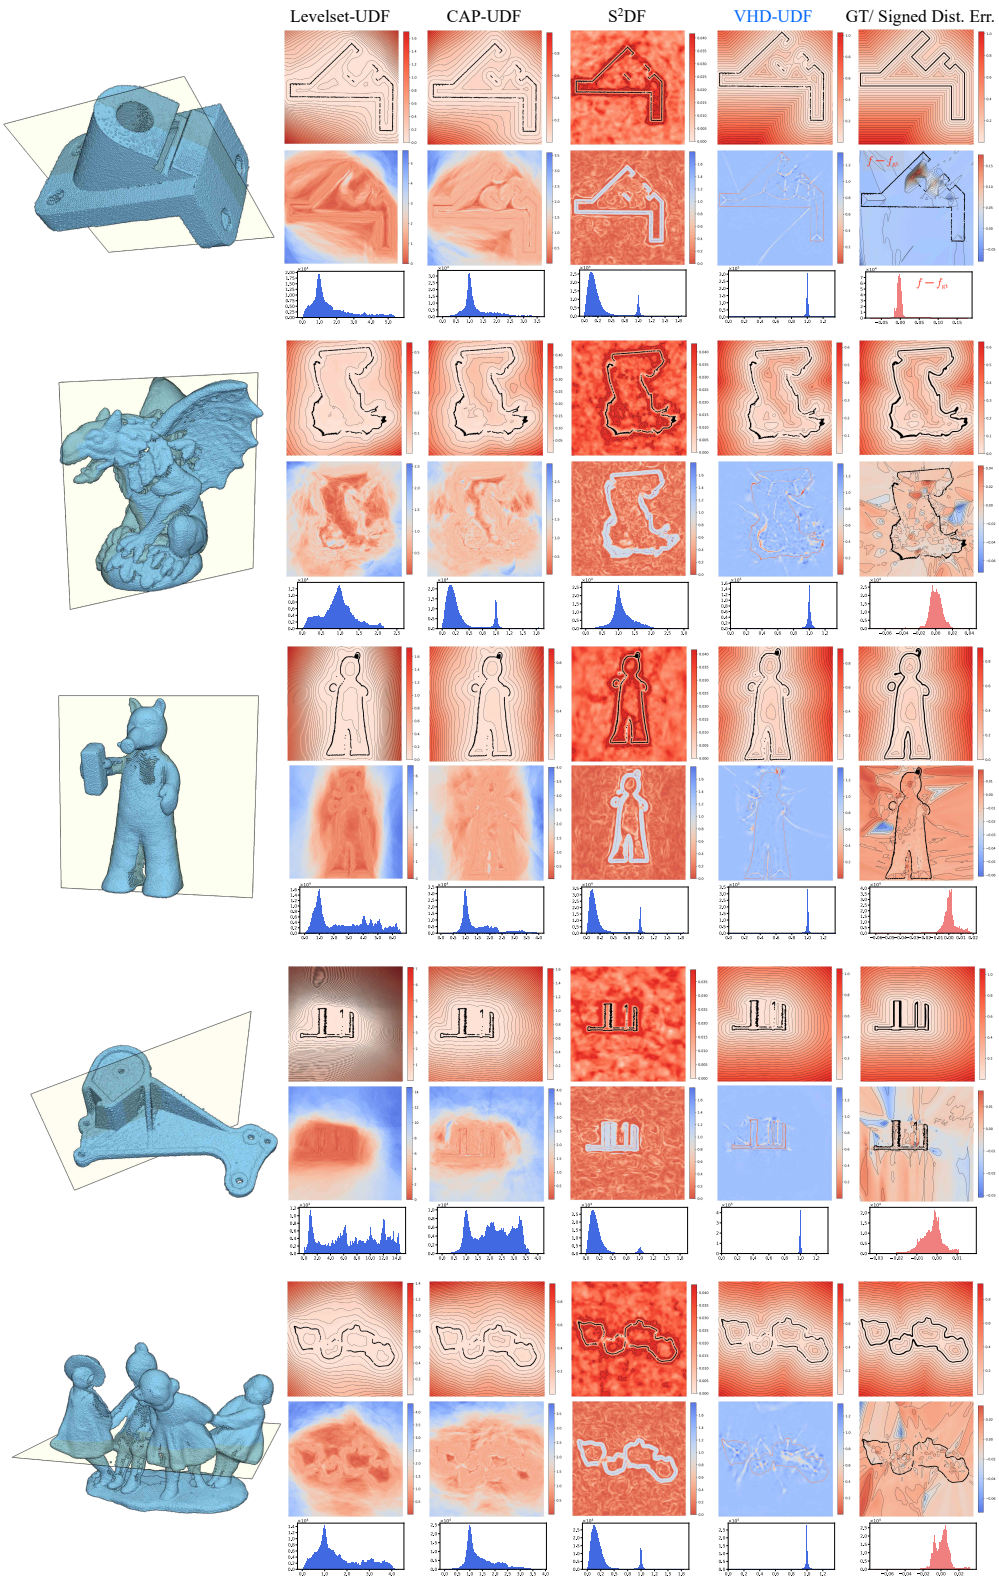

Figure 22: Field slices of SRB models. Top: learned UDF f. Second row: gradient norm maps $\|\nabla f\|$. Bottom: gradient norm statistics. The right side shows the ground-truth UDF and signed distance error map of VHD-UDF. Colorbars vary across figures to better show the results with high contrast.