DUDF: Differentiable Unsigned Distance Fields with Hyperbolic Scaling

Miguel Fainstein[1] miguelon.f98@gmail.com
Viviana Siless[1] viviana.siless@utdt.edu
Emmanuel Iarussi[1,2] emmanuel.iarussi@utdt.edu
[1]Universidad Torcuato Di Tella, [2]CONICET

Abstract

In recent years, there has been a growing interest in training Neural Networks to approximate Unsigned Distance Fields (UDFs) for representing open surfaces in the context of 3D reconstruction. However, UDFs are non-differentiable at the zero level set which leads to significant errors in distances and gradients, generally resulting in fragmented and discontinuous surfaces. In this paper, we propose to learn a hyperbolic scaling of the unsigned distance field, which defines a new Eikonal problem with distinct boundary conditions. This allows our formulation to integrate seamlessly with state-of-the-art continuously differentiable implicit neural representation networks, largely applied in the literature to represent signed distance fields. Our approach not only addresses the challenge of open surface representation but also demonstrates significant improvement in reconstruction quality and training performance. Moreover, the unlocked field's differentiability allows the accurate computation of essential topological properties such as normal directions and curvatures, pervasive in downstream tasks such as rendering. Through extensive experiments, we validate our approach across various data sets and against competitive baselines. The results demonstrate enhanced accuracy and up to an order of magnitude increase in speed compared to previous methods.

1. Introduction

Surface representation is a fundamental aspect in the field of 3D geometry processing, with explicit methods such as meshes, point clouds, and voxelized representations being traditional choices. Implicit surface representations, on the other hand, have been an integral part of the graphics pipeline for many decades. They encapsulate surfaces as the zero-level set of a function, providing a compact and continuous geometry representation. The novelty in recent years has emerged from parameterizing these implicit functions with Neural Networks (NNs), combining their learning capabilities with the advantages of implicit representations.

Signed Distance Functions (SDFs) have traditionally been the chosen formulation for implicit surface representation [4, 7] due to their well-defined gradients and the ease they offer for computing constructive solid geometry operations and mesh reconstruction. However, they are inherently limited to closed surfaces, which poses a significant challenge for representing open surfaces with implicit methods. This limitation arises from the inside/outside sign flip on the zero level set, which is impossible to define for surfaces that do not enclose a volume.

The advent of Unsigned Distance Functions (UDFs) extended representation capabilities to open surfaces. However, this advancement introduced a challenge: the non-differentiable nature of the distance functions at the zero level set. Such non-differentiability leads to inaccuracies in distances and gradients learned by NNs, particularly near the surface where precision is paramount. This paper tackles the challenge of representing open surfaces with neural networks, focusing on a formulation that ensures the surface continuity and smoothness.

2. Proposed Approach

2.1. Mathematical background

Methods addressing closed surfaces approach the learning of SDFs as finding the solution to a system of Partial Differential Equations (PDEs) governed by the homogeneous *Eikonal* equation with *Dirichlet* and *Neumann* boundary conditions. Although SDFs are not differentiable at every point and represent only a weak solution to the *Eikonal* equation, recent research has demonstrated considerable success in solving this problem with continuously differentiable implicit neural representation networks [3, 5, 8, 10]. These architectures use periodic activations, facilitating smoother optimization processes and improved control over the solution's gradient field. This is achievable because locations where the signed distance lacks differentiability are distant from the isosurface, hence the approximation tends to be good in a close neighborhood of the zero level set. In particular, these networks enable accurate computation of critical topological properties, includ-

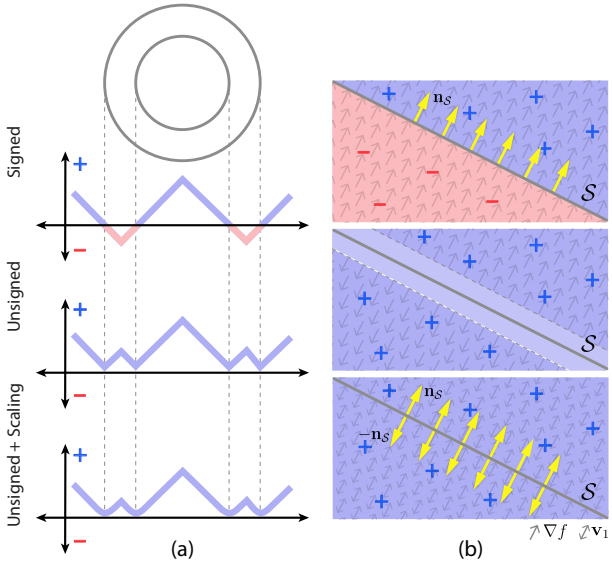

Figure 1. In (a), view of different distance fields for a 2D slice of a torus. Note the effect of hyperbolic scaling near the isosurface (bottom). In (b), sign and gradient for each distance field f. In signed distance, the gradient ∇f at the isosurface is equal to the surface's normal field $\mathbf{n}_\mathcal{S}$ (top). In unsigned distance, the gradient is undefined at the isosurface due to the change in orientation (middle). Our maximum curvature field vectors $\mathbf{v}_1$ align with the surface's unoriented normals (bottom).

ing mean and Gaussian curvature [8]. This contrasts with networks featuring piecewise linear activations, such as *ReLUs*, which have null second-order derivatives [10].

In the context of representing UDFs, these functions cannot be a solution to the *Eikonal* equation both in the interior and exterior regions of the surface, without losing differentiability at the zero level set. This presents challenges for continuously differentiable NNs in achieving satisfactory outcomes, being that the approximation errors happen at the isosurface where accuracy is paramount.

2.2. Problem statement

Our key insight is to overcome these challenges by redefining the unsigned distance field, particularly through the application of a hyperbolic scaling. In light of this understanding, we propose to learn the parameters θ of a neural network f_θ with periodic activation functions [10] to approximate the function:

$$t_\mathcal{S}(\mathbf{x}) = d_\mathcal{S}(\mathbf{x}) \tanh(\alpha d_\mathcal{S}(\mathbf{x})), \tag{1}$$

where $d_\mathcal{S}$ is the unsigned euclidean distance to surface $\mathcal{S}$; and α is a constant value. The function $t_\mathcal{S}$ is a differentiable approximation of $d_\mathcal{S}$, whose zero level set is the surface $\mathcal{S}$. Fig. 1 (a) illustrates hyperbolic scaling's quadratic smoothing near the isosurface and the linear behavior in the distance.

The scaling enables the application of continuously differentiable implicit neural representation networks to solve an *Eikonal* equation, while retaining the UDF's open surface representation capabilities. For this task, we aim to address a heterogeneous *Eikonal* equation, for which we know $t_\mathcal{S}$ (Eq. 1) is a weak solution:

$$\begin{cases} \|\nabla f\| = \phi \\ f_{|\mathcal{S}} = 0 \\ \nabla f_{|\mathcal{S}} = \mathbf{0} \\ \mathbf{v_1}(f_{|\mathcal{S}}) = \pm \mathbf{n}_\mathcal{S} \end{cases} \tag{2}$$

where ϕ is defined as the L2 norm of the gradient of Eq. 1, formally:

$$\phi(\mathbf{x}) = \tanh(\alpha d_\mathcal{S}(\mathbf{x})) + \alpha d_\mathcal{S}(\mathbf{x})(1 - \tanh^2(\alpha d_\mathcal{S}(\mathbf{x}))). \tag{3}$$

And function $\mathbf{v_1}(f_{|\mathcal{S}})$ represents the principal curvature direction of f at surface $\mathcal{S}$. This coincides with the unitary eigenvector associated with the highest eigenvalue of hessian matrix $\mathbf{H}_{f_\theta}(\mathbf{s})$. This formula allows to compute surface normals during inference and reconstruction, which is a key difference with previous methods that relied on expensive and imprecise numerical approximations.

2.3. Implicit function learning

Following the aforementioned definitions, our neural networks are trained to minimize the following loss function:

$$\mathcal{L} = \lambda_e \mathcal{L}_{Eikonal} + \lambda_d \mathcal{L}_{Dirichlet} + \lambda_n \mathcal{L}_{Neumann} + \lambda_g \mathcal{L}_{MCurv}, \tag{4}$$

with λ_i constant weights controlling the relevance of each term, defined as:

$$\mathcal{L}_{Eikonal} = \int_\mathcal{C} |\|\nabla f_\theta(\mathbf{x})\| - \phi(\mathbf{x})| \quad d\mathbf{x}$$

$$\mathcal{L}_{Dirichlet} = \int_\mathcal{S} |f_\theta(\mathbf{x})| \quad d\mathbf{x}$$

$$\mathcal{L}_{Neumann} = \int_\mathcal{S} \|\nabla f_\theta(\mathbf{x})\| \quad d\mathbf{x}$$

$$\mathcal{L}_{MCurv} = \int_\mathcal{S} 1 - |\mathbf{v_1}(\mathbf{x}) \cdot \mathbf{n}_\mathcal{S}(\mathbf{x})| \quad d\mathbf{x}$$

Similarly to previous works [3, 8], we extend the *Dirichlet* loss function to points far from the surface $\mathcal{S}$. This is achieved by computing an approximation of the function $t_\mathcal{S}$ (Eq. 1) based on nearest neighbors.

3. Results and evaluation

3.1. Experimental setup

We conducted a series of experiments to assess the performance of our method. To target a broad amount of surfaces, we experimented on three well-known data sets: ShapeNet cars [2], Multi-Garment [1], and DeepFashion

Method	DeepFashion [13]				Multi-Garment [1]				ShapeNet cars [2]			
	time(s) ↓	L1CD ↓	L2CD ↓	NC ↓	time(s) ↓	L1CD ↓	L2CD ↓	NC ↓	time(s) ↓	L1CD ↓	L2CD ↓	NC ↓
SIREN [10]	376	27.3	1.980	0.107	374	40.5	8.810	0.094	751	16.9	0.170	**0.240**
CAP-UDF [12]	1390	18.1	1.110	0.080	1440	18.5	1.190	0.083	1040	**10.9**	0.071	0.288
Ours (MC1 [12])	**326**	**9.01**	**0.025**	0.024	**318**	**8.70**	**0.024**	0.026	**317**	12.3	**0.057**	0.387
Ours (MC2 [6])		9.14	0.027	**0.020**		8.82	0.026	**0.021**		13.7	0.081	0.304

Table 1. Training time, L1 and L2 mean Chamfer distances ($\times 10^3$), and Normal Consistency (NC) for the evaluated open surface data sets. SIREN was trained as described in the original paper, while CAP-UDF and our method were trained on unsigned distances.

[13]. We undertake mesh reconstruction using two gradient-based Marching Cubes algorithms, referred here as MC1 [12] and MC2 [6].

3.2. Surface reconstruction

3.2.1 Closed shapes

We benchmarked DUDF against state-of-the-art methods for representing closed surfaces, where signed distances are well-defined. Given that our primary focus is not on closed surface representation, we limited this comparison to the ShapeNet car dataset [2], modifying the meshes for closure and omitting internal structures [11]. The comparative results are detailed in Table 2.

Method	time(s) ↓	L1CD ↓	L2CD ↓
DeepSDF [9]	**94**	19.40	0.206
SIREN [10]	379	15.40	0.171
CAP-UDF [12]	1080	**9.48**	0.030
Ours (MC1 [12])	319	**9.48**	**0.028**
Ours (MC2 [6])		9.49	**0.028**

Table 2. Training time, L1, and L2 mean Chamfer distances ($\times 10^3$) for the closed ShapeNet cars data set. DeepSDF and SIREN were trained on signed distances, while CAP-UDF and our method were trained on unsigned distances.

3.2.2 Open surfaces

Regarding open surfaces, a full quantitative analysis is presented in Table 1. Qualitative results can be seen in Fig. 2. On the one hand, our method demonstrates a significant improvement in efficiency, consuming up to an order of magnitude less computational time than CAP-UDF, while also showing enhanced performance across all three data sets. This advantage is particularly evident in the DeepFashion and Multi-Garment data sets, where CAP-UDF tends to inaccurately close openings (like those at the ends of sleeves in clothing).

In addition to these experiments, we compare our method with CAP-UDF in the context of sphere tracing rendering. This comparison is presented in Fig. 3, where we showcase two rendering examples. Since learned unsigned distance fields in CAP-UDF do not grow linearly away from the surface, the marching steps in sphere tracing often fail to accurately intersect the surface, leading to undesirable visual artifacts. Even when hitting the surface, gradients might be undefined, leading to noisy shaded images. In contrast, our proposed distance function does not exhibit these problems, offering a more robust framework for direct rendering. Furthermore, our ability to compute normals using the maximum curvature field is useful for shading purposes, enhancing the rendering quality without undergoing an intermediate 3D mesh reconstruction step.

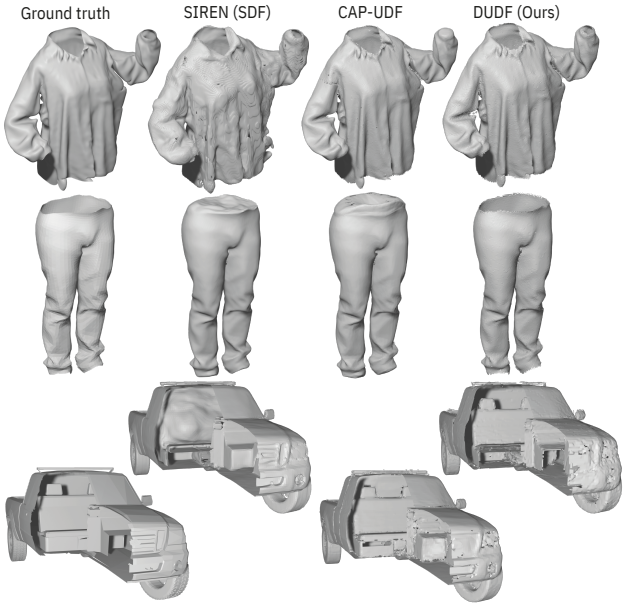

Figure 2. Comparisons on DeepFashion [13] (top row), Multi-Garment [1] (middle row), and ShapeNet cars [2] (bottom row) data sets. DUDF preserves fine details and accurately represents complex geometries without closing holes, outperforming SIREN (SDF), which tends to smooth and round models, and CAP-UDF, which captures sharp features but often closes open surfaces. Reconstructions for CAP-UDF and DUDF performed with MC1 [12].

3.3. Ablation study

We further study the impact of individual loss components on the network's accuracy through an ablation study (Table 3). The methodology involves selectively deactivat-

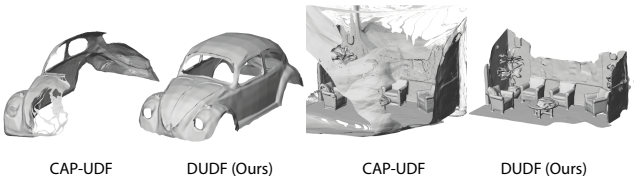

Figure 3. Rendering examples using sphere tracing. CAP-UDF struggles with non-linear growth of their unsigned distance fields, causing the sphere tracing marching step to miss the surface. Conversely, our method demonstrates precision in direct rendering scenarios.

ing the terms related to each in our loss function to isolate their contributions. We adopted the experimental framework from the prior section, training on 30 meshes from the DeepFashion dataset [13]. To elucidate the role of boundary conditions in enhancing reconstruction quality, we incorporated a normal consistency metric into our analysis. Additionally, we contrasted the effects of approximating the ground truth distance function $d_{\mathcal{S}}$ (with its distinct *Eikonal* problem), using the same network architecture and sampling scheme, with *sine* and *ReLU* activations. The findings from this study show that such an approach leads to suboptimal outcomes, thereby underscoring the effectiveness of $t_{\mathcal{S}}$ and the *Eikonal* problem proposed.

Method	time(s) ↓	L1CD ↓	L2CD ↓	NC ↓
Baseline	326	**9.14**	**0.027**	**0.020**
$\lambda_e = 0$	312	9.43	0.028	0.033
$\lambda_g = 0$	150	9.16	**0.027**	0.021
$\lambda_e, \lambda_g = 0$	**109**	9.44	**0.027**	0.035
$\lambda_\mu, \lambda_\sigma = 0$	302	9.24	**0.027**	0.031
$d_{\mathcal{S}}$ (sine)	150	31.2	0.830	0.057
$d_{\mathcal{S}}$ $(ReLU)$	145	47.5	2.500	0.149

Table 3. Quantitative impact of each loss component on the reconstruction accuracy of our network. We report training time, L1, and L2 mean Chamfer distances $\times 10^3$, and Normal Consistency (NC). Last two rows correspond to approximating the ground truth distance function $d_{\mathcal{S}}$ instead of $t_{\mathcal{S}}$.

3.4. Computing mean and Gaussian curvature

A significant advantage of our method over previous approaches is the possibility to compute curvatures. Previous work learned non-differentiable functions, which preclude the direct computation of geometrical properties such as mean and Gaussian curvature, extracted through second order differentiation. In Fig. 4 we show mean and Gaussian curvatures for open and closed surfaces computed with our method.

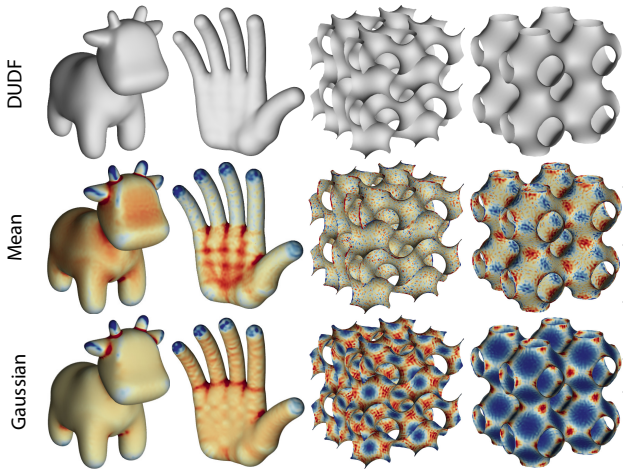

Figure 4. Mean and Gaussian curvatures computed with our method for closed surfaces (left) and open surfaces (right).

4. Limitations and future work

Our method encounters certain limitations. Primarily, by approximating the function $t_{\mathcal{S}}$ (Eq. 1), the *Eikonal* equation requires supervision with ground truth unsigned distances, in contrast to signed distance approaches where it is constant everywhere. The quadratic behavior of the function near the isosurface produces a wider near-zero strip which can lead to slightly inflated representations. Additionally, computing the maximum curvature direction field is computationally expensive, requiring three network passes and the diagonalization of the Hessian matrix.

5. Conclusion

In this work we introduced Differentiable Unsigned Distance Fields (DUDF) with Hyperbolic Scaling, a novel approach that addresses inherent limitations of traditional UDFs in representing open surfaces. By applying a hyperbolic transformation to the distance field, we define a new variant of the *Eikonal* problem, tailored with unique boundary conditions. The conducted experiments provide evidence that our approach can lead to improvements in reconstruction quality and computational performance when compared to several state-of-the-art methods. Moreover, the ability to accurately calculate topological features such as normals and curvatures is a notable benefit of our model, offering additional utility in geometric processing tasks and rendering.

References

[1] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *IEEE International Conference on Computer Vision (ICCV)*. IEEE, oct 2019. 2, 3

[2] 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. 2, 3

[3] Mattéo Clémot and Julie Digne. Neural skeleton: Implicit neural representation away from the surface. *Computers & Graphics*, pages 368–378, 2023. 1, 2

[4] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996. 1

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

[6] Benoit Guillard, Federico Stella, and Pascal Fua. Meshudf: Fast and differentiable meshing of unsigned distance field networks, 2022. 3

[7] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. Ieee, 2011. 1

[8] Tiago Novello, Guilherme Schardong, Luiz Schirmer, Vinicius da Silva, Helio Lopes, and Luiz Velho. Exploring differential geometry in neural implicits, 2022. 1, 2

[9] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 3

[10] 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. 1, 2, 3

[11] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *CoRR*, abs/1905.10711, 2019. 3

[12] Junsheng Zhou, Baorui Ma, Yu-Shen Liu, Yi Fang, and Zhizhong Han. Learning consistency-aware unsigned distance functions progressively from raw point clouds. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3

[13] 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. *CoRR*, 2020. 3, 4