



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

RoCoSDF: Row-Column Scanned Neural Signed Distance Fields for Freehand 3D Ultrasound Imaging Shape Reconstruction

Hongbo Chen¹, Yuchong Gao¹, Shuhang Zhang¹, Jiangjie Wu¹,
Yuexin Ma¹, and Rui Zheng^{1,2} (✉)

¹ School of Information Science and Technology, ShanghaiTech University, Shanghai, China

zhengrui@shanghaitech.edu.cn

² Shanghai Engineering Research Center of Intelligent Vision and Imaging, ShanghaiTech University, Shanghai, China

Abstract. The reconstruction of high-quality shape geometry is crucial for developing freehand 3D ultrasound imaging. However, the shape reconstruction of multi-view ultrasound data remains challenging due to the elevation distortion caused by thick transducer probes. In this paper, we present a novel learning-based framework RoCoSDF, which can effectively generate an implicit surface through continuous shape representations derived from row-column scanned datasets. In RoCoSDF, we encode the datasets from different views into the corresponding neural signed distance function (SDF) and then operate all SDFs in a normalized 3D space to restore the actual surface contour. Without requiring pre-training on large-scale ground truth shapes, our approach can synthesize a smooth and continuous signed distance field from multi-view SDFs to implicitly represent the actual geometry. Furthermore, two regularizers are introduced to facilitate shape refinement by constraining the SDF near the surface. The experiments on twelve shapes data acquired by two ultrasound transducer probes validate that RoCoSDF can effectively reconstruct accurate geometric shapes from multi-view ultrasound data, which outperforms current reconstruction methods. Code is available at <https://github.com/chenhbo/RoCoSDF>.

Keywords: Multi-view reconstruction · Neural shape representation · Freehand 3D ultrasound

1 Introduction

Freehand 3D Ultrasound (US) imaging has gained considerable attention in clinical diagnostics due to its flexibility, portability and large field-of-view imaging [1,2,3,4]. Most current freehand imaging systems usually collect 2D images and 3D poses by scanning the object in a single-view direction (Fig. 1), utilizing a tracking device attached to the ultrasound transducer (UT) [2]. Accurate reconstruction of geometric shape or surface is essential for freehand 3D US




Fig. 1. Two typical data acquisition manners for freehand 3D US imaging using one ultrasound transducer (UT). (a) and (b) Single-view scanning and shape reconstruction in row and column directions from UNSR [6]. (c) Proposed row-column scan and multi-view shape reconstruction. The shape contour (dark) is optimized from row-column scan. In practical use for handheld scans, our row-column scan is not necessarily orthogonal, as indicated by the dashed blue lines and red lines in the left circle.

imaging in clinical settings, particularly for US hard tissue imaging because the US signal is unable to penetrate the boundaries of hard tissues [5,6]. However, single-view scanning provides restricted perspectives, often resulting in insufficient anatomical information and distorted structures in the scanning direction, as illustrated in Fig. 1.

Multi-view US scanning can integrate multiple viewpoints, offering comprehensive spatial understanding and complex anatomical shapes. However, the view-dependent nature of US imaging introduces additional challenges for 3D reconstruction, since imaging a target from multiple scanning directions can produce different reflected intensities. Current existing multi-view reconstruction methods mainly focus on two strategies for freehand or robotic scanning: orientation-based intensity compounding [7,8,1] and neural radiance field (NeRF) based view synthesis [9,10]. Although orientation-based methods can partially address the view-dependent issue, their results are limited by the spatial resolution and discrete pixel connectivity problem [2]. On the other hand, the neural radiance field (NeRF)-based method, as an emerging technology, offers the advantage of being resolution-independent in representation. However, it struggles to accurately capture the geometric shapes [11,12]. Furthermore, the development of these methods in freehand scanning is more challenging because slight motion interference may lead to mismatches in multi-view fusion.

With the advance of neural implicit function, 3D continuous shape representation has recently emerged as a useful technique for improving the geometric appearance in computer vision and medical imaging [13,14,15,16]. These representations are commonly parameterized as a neural network that map 3D coordinates to implicit values, such as signed distance function (SDF) [13,17,18]. Inspired by [17], a self-supervised learning method has recently been reported

for freehand 3D US neural surface reconstruction (UNSR) [6]. It trains a Multi-Layer Perceptrons (MLPs) network to learn neural SDF from input US volumetric masks without the requirements of ground truth signed distance fields, point cloud normals, and occupancy fields. However, the limitations of single-view imaging hinder UNSR from learning precise structure. Fig. 1(a) and (b) illustrate that the overall vertebra shape is elongated in the row and column direction by such a method.

To address these limitations, we propose RoCoSDF: **Row-Column** scanned neural **Signed Distance Function** fields, a novel framework based on neural implicit functions for shape reconstruction of multi-view freehand 3D ultrasound imaging. As shown in Fig. 1(c), RoCoSDF aims to synthesize an SDF field to implicitly represent the actual structures without ground truth shape supervision by utilizing two typical scan views in row and column directions. Specifically, given the row-column segmented points from a shape, we first learn two SDFs for the row-scan and column-scan using a self-supervised strategy in a normalized space. Then, constructive solid geometry (CSG) is employed to extract a signed distance field from the two SDFs, which serves as the initial representation to be refined further. We subsequently sample the query points and SDF values in the SDF field for shape refinement using a supervised strategy. Particularly, we introduce two regularizers to enhance the learning of a more accurate neural SDF field. Our proposed method is quantitatively and qualitatively validated on twelve vertebrae data scanned by two UTs. The results demonstrate superior shape fidelity over existing approaches.

2 Methodology

2.1 Overview

An overview of the proposed framework for multi-view neural shape reconstruction from row-column scan is illustrated in Fig. 2. Our coarse-to-fine reconstruction pipeline primarily comprises four steps. In Step (a), separate training is conducted to predict the row and column SDFs from input row-column datasets. Step (b) aims to fuse the predicted SDFs to obtain a signed distance field, and step (c) involves sampling the query points and SDF value in the distance field for shape refinement. Finally, step (d) extracts the 3D mesh from the optimized SDF using the Marching Cube algorithm [19].

2.2 RoCoSDF: Multi-view Neural Shape Reconstruction

Row-Column Neural SDFs Prediction A neural network f_θ , such as MLP, can be trained as an implicit function to map any 3D point coordinate, $\mathbf{x} = [x, y, z] \in R^3$, to its corresponding SDF value of $s \in R$. Here, θ represents the learnable parameters of the network. The object surface \mathcal{S} is implicitly represented by the zero-level-set of neural SDFs, $f_\theta(\cdot) = 0$. Each 3D shape can be individually parameterized by a $f_\theta(\cdot)$.

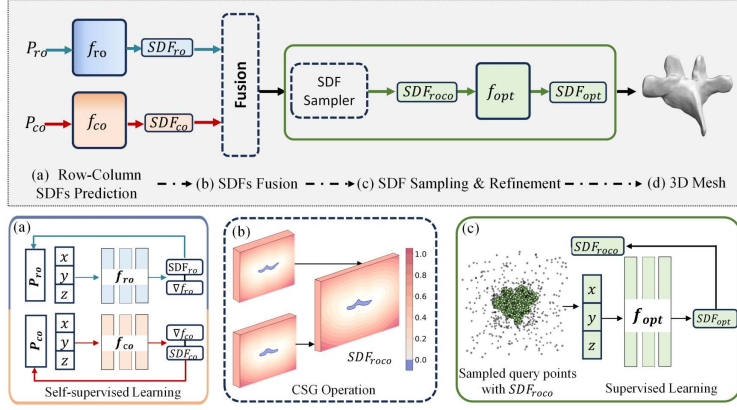


Fig. 2. An overview of the proposed framework for shape reconstruction. (a) Row-Column neural SDFs prediction from point cloud P_{ro} and P_{co} . (b) SDFs fusion using constructive solid geometry (CSG). (c) SDF sampling and refinement.

$$f_{\theta}(\mathbf{x}) = s, \begin{cases} d(\mathbf{x}, \mathcal{S}), & \text{if } \mathbf{x} \text{ is outside of } \mathcal{S} \\ 0, & \text{if } \mathbf{x} \text{ is on the } \mathcal{S} \\ -d(\mathbf{x}, \mathcal{S}), & \text{if } \mathbf{x} \text{ is inside of } \mathcal{S} \end{cases} \quad (1)$$

where $d(\mathbf{x}, \mathcal{S})$ is the positive distance from \mathbf{x} to surface \mathcal{S} . The sign before $d(\mathbf{x}, \mathcal{S})$ is positive (negative) when \mathbf{x} reaches the surface from outside (inside) of object.

A query point \mathbf{x} can be projected to its nearest point \mathbf{x}' on the surface \mathcal{S} along or against the network gradient ($\nabla f_{\theta}(\mathbf{x})$) according to the predicted signed distance.

$$\mathbf{x}' = \mathbf{x} - f_{\theta}(\mathbf{x}) \cdot \nabla f_{\theta}(\mathbf{x}) \quad (2)$$

Here, we utilize two MLPs neural networks to predict the row-column neural SDFs for row-scan and column-scan, respectively.

$$\begin{aligned} SDF_{ro}(\mathbf{x}) &= f_{ro}(\mathbf{x}) \\ SDF_{co}(\mathbf{x}) &= f_{co}(\mathbf{x}) \end{aligned} \quad (3)$$

where f_{ro} and f_{co} are the row-scan SDF decoder and the column-scan SDF decoder, respectively

In step (a), the SDF_{ro} and SDF_{co} map the 3D query points generated from dual-view point clouds P_{ro} and P_{co} to a normalized signed distance field using MLPs by a series of self-supervised loss functions with regularizers.

Signed Distance Fields Fusion Constructive Solid Geometry (CSG) commonly includes three boolean operations: union, intersection and difference, which are computing on the implicit functions [20,18]. With row and column

signed distance fields, an SDF can be coarsely fused through CSG operation by operating the intersection between SDF_{ro} and SDF_{co} .

$$\text{Intersection } f_{ro} \cap f_{co} : SDF_{roco} = \max(SDF_{ro}, SDF_{co}) \quad (4)$$

After the fusion, the actual object shape can be naively reconstructed using SDF_{roco} . However, the direct CSG operation on row-column SDFs will lead to a rough local surface and unexpected errors since CSG is commonly designed for primitive shape elements.

SDF Sampling and Refinement To refine SDF_{roco} , we further train a f_{opt} to optimize the fused distance fields in a supervised way. An SDF sampler is designed to directly sample the query points and SDF values in a 3D cube of SDF_{roco} . We randomly generate the query points within the cube space ranging from -1 to 1. Each query point is fed into the SDF_{roco} to obtain the corresponding SDF value. Following [13], we then sample more aggressively near the zero-level-set of the SDF_{roco} using a Gaussian function, $G(SDF_{roco}(\mathbf{x}), \sigma)$. This strategy facilitates the network learning more detailed SDF near the object surface.

2.3 Model Training and Loss Functions

We train the network using two learning strategies: 1) self-supervised learning for f_{ro} and f_{co} in step (a) since there are no ground truth SDFs available during this stage and 2) supervised learning for f_{opt} in step (c) by using SDF_{roco} as pseudo ground truth.

$$\mathcal{L}_{step_a} = \mathcal{L}_{self}, \mathcal{L}_{step_c} = \mathcal{L}_{super} \quad (5)$$

Self-supervised Learning The definition of self-supervised loss is mainly following UNSR [6]. We additionally introduce a non-manifold regularizer to penalize query points that are not on the surface but are close to zero-level-set due to the constraints of \mathcal{L}_{scc} [21].

$$\begin{aligned} \mathcal{L}_{self} &= \mathcal{L}_{sdf} + \lambda_{scc} \mathcal{R}_{nonmfd} \mathcal{L}_{scc} + \lambda_{adl} \mathcal{L}_{adl}, \\ \mathcal{R}_{nonmfd} &= \exp(-\alpha_{nonmfd} |f_{ro/co}(\mathbf{x})|), \mathcal{L}_{sdf} = \|\mathbf{x}' - \hat{\mathbf{x}}\|_2^2, \\ \mathcal{L}_{scc} &= 1 - \cos(\nabla f_{ro/co}(\mathbf{x}), (\mathbf{x}' - \hat{\mathbf{x}}) / \|\mathbf{x}' - \hat{\mathbf{x}}\|_2) \end{aligned} \quad (6)$$

where $\hat{\mathbf{x}} \in P$ is the nearest neighbor points of \mathbf{x} , $\nabla f_{ro/co}$ is the gradient of network f_{ro} and f_{co} , \mathcal{L}_{sdf} is an L2-distance loss to optimize the projected \mathbf{x}' in Eq. 2 reaching its nearest $\hat{\mathbf{x}}$, and \mathcal{L}_{adl} is the adversarial learning strategy in [6].

Supervised Learning The supervised loss is a L1-norm distance between the predicted SDF and SDF_{roco} with a manifold regularizer [13,21].

$$\mathcal{L}_{super} = \|f_{opt}(\mathbf{x}) - SDF_{roco}\|_1 + \lambda_{mfd} \mathcal{R}_{mfd}, \mathcal{R}_{mfd} = \|f_{opt}(\mathbf{x})\|_1 \quad (7)$$

Table 1. Performance comparison of our approach with the baseline on two datasets.

Methods	CD (mm)↓	HD (mm)↓	MAD (mm)↓	RMSE (mm)↓
UT1 Scans				
UNSR [6] (Row)	2.16 ± 0.16	5.21 ± 1.77	1.84 ± 0.14	2.25 ± 0.17
UNSR [6] (Col)	2.11 ± 0.18	5.82 ± 0.80	1.78 ± 0.17	2.25 ± 0.21
RoCoSDF (Ours)	1.75 ± 0.09	4.08 ± 0.74	1.34 ± 0.05	1.70 ± 0.03
UT2 Scans				
UNSR [6] (Row)	2.40 ± 0.62	5.22 ± 2.47	1.97 ± 0.67	2.39 ± 0.86
UNSR [6] (Col)	2.54 ± 0.63	7.53 ± 2.45	2.25 ± 0.69	2.97 ± 0.95
RoCoSDF (Ours)	2.03 ± 0.36	4.87 ± 2.80	1.53 ± 0.47	1.92 ± 0.74

where \mathcal{R}_{mfd} is a manifold regularizer to penalize query points away from the surface for smoothness.

3 Experiments and Results

3.1 Data Acquisition and Preparation

Six computer-aided designed (CAD) vertebra models are used as phantoms from 3D printing, including three typical thoracic vertebrae (T4, T8, T12) and three typical lumbar vertebrae (L1, L3, L5). These CAD models serve as ground truth for evaluation. For the generalizability analysis, two different US transducers (UT1 and UT2) are adopted to collect two datasets using the row-column scan. For one model scanned by one UT, each row-column scan corresponds to 1 shape and 2 scans. Two UTs obtain 12 shapes and 24 scans from 6 models. An average of 540 ± 159 frames from UT1 and 686 ± 244 frames from UT2 are collected. The electromagnetic (EM) positioning sensor is attached on the transducers for locating images in tracking space. The EM positioning sensor and UT are calibrated using the Levenberg-Marquardt algorithm [22]. More device parameters are listed in the supplementary material.

After the data acquisition, the 3D points in the region-of-interest (ROI) are segmented and transformed from the tracking space to a normalized 3D space, $[-1, 1]$. Farthest Point Sampling (FPS) [23] is applied to the row and column 3D point clouds for downsampling to speed up the training. We randomly sample 20000 3D points from each segmented dataset using FPS and normalized them to a 3D cube space to acquire the raw point cloud. Then, for each point in the point clouds, we sample 20 query points following Gaussian distribution with mean 0 and standard deviation [6,17]. The standard deviation is defined as the distance between \mathbf{x} and its 50th nearest neighbor points. Additionally, we randomly sample more query points within the same cube space to ensure the networks learn an SDF everywhere.

3.2 Evaluation Metrics

Our approach is compared with the baseline method UNSR, the recent first method for the application of neural shape representation to freehand 3D US

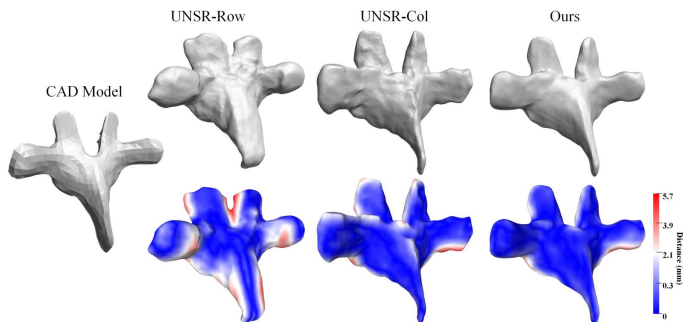


Fig. 3. Visualization of thoracic vertebra T4. The top and bottom are the 3D meshes and colored error maps from CAD model, UNSR (row-scan), UNSR (column-scan) and RoCoSDF (ours).

imaging [6]. Four evaluation metrics are used to assess the reconstruction quality: Chamfer Distance (CD), Hausdorff Distance (HD), Mean Absolute Distance (MAD) and Root Mean Square Error (RMSE). The distances are computed between the points randomly sampled from reconstructed mesh and the corresponding CAD models.

We directly visually compare our approach with the traditional orientation-based reconstruction algorithms due to the poor geometry appearance of traditional algorithms. The posed images from row-column scans are compounded using Pixel-Based Methods (PBM) [2,7]. The reconstructed volume is visualized through volume rendering. We show the results in the supplementary material.

3.3 Implementation Details

Three neural networks, f_{row} , f_{col} and f_{opt} , consist of 8 MLPs and 256 hidden channels with the skip connection at the fourth layer. A Relu activation function is attached after each layer except for the last layer. We train f_{row} , f_{col} and f_{opt} for 1.0×10^4 iterations using the Adam optimizer. The learning rate is set to 0.001 with a cosine decay schedule. The batch size, λ_{scc} and λ_{adl} are set to 5000, 0.01 and 0.01, respectively. The α_{nonmfd} and λ_{mfd} are set to 100 and 0.6. Our network is implemented using Pytorch and trained on a single NVIDIA RTX 3090 GPU with 24 GB memory. After the training, we set the mesh resolution to 256^3 and the threshold of Marching Cube to 0 to extract the zero-level-set of surface boundaries using SDF_{opt} .

3.4 Qualitative and Quantitative Results

As listed in Table 1, RoCoSDF achieves leading performance with UNSR across all four evaluation metrics on both two datasets. Our approach demonstrates a 27% and 24% improvement over the UNSR row scan in UT1 scans in terms of MAD and RMSE, and 32% and 35% improvement over the UNSR column scan

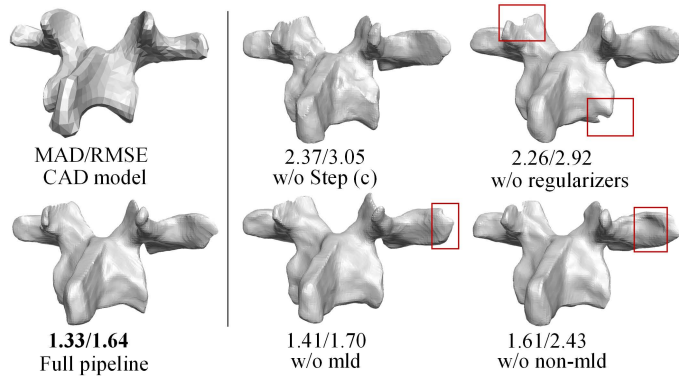


Fig. 4. Ablation study on the effectiveness of step (c) and two regularizers.

in UT2 scans in terms of MAD and RMSE. More specifically, we observe that the quality of single-view reconstruction results from the UT2 scans dataset is inferior compared to those obtained from the UT1 scan dataset. This discrepancy in reconstruction performance can be attributed to the lower elevation resolution of the UT2 probe. Nonetheless, our approach retains the capability to accurately recover the real structure from the row-column scan data, for example, a substantial average decrease of 19% and 35% in RMSE than UNSR-row and UNSR-column. Statistical significance is established with p-value < 0.01 against UNSR across all metrics.

We present a qualitative result of thoracic vertebra T4 in Fig. 3 and more results in supplementary material. The reconstructed mesh and the corresponding CAD model are superimposed to visualize the error map with color coding. The results from the single-view scanning by UNSR fail to reconstruct the accurate shape structure regardless of row-scan or column-scan, with errors mainly distributed along the scanning direction. Compared to them, our method can extract latent shape information from each view to complementarily recover accurate and smooth surfaces, as shown in the colorized error map at the bottom of Fig. 3.

We qualitatively demonstrate an ablation study on lumbar vertebra L3 to explore the effectiveness of step (c) and two regularizers. As shown in Fig. 4, the generated meshes from direct CSG operation without step (c) yield obvious editing traces and noises. The non-manifold regularization facilitates the network to learn a more compact shape and thin structures, while the manifold regularization refines the surface details.

4 Conclusion

We present RoCoSDF, a novel neural-SDF-based framework for multi-view free-hand 3D US shape reconstruction from row-column scanned data. Our approach

solves the challenges of view-dependent and pixel connectivity issues in multi-view 3D US by directly operating on 3D signed distance fields. In addition, we design a coarse-to-fine optimization strategy to enhance the shape appearance with surface regularizers, eliminating the need for additional ground truth shape supervision. Evaluation results on two UTs demonstrate the high fidelity and generalizability of our method.

The proposed RoCoSDF can effectively address the elevation thickness problem in freehand 3D US imaging as well as be easily extensible for robotic 3D US imaging. This new design also holds promise to reconstruct precise shape from incomplete or noisy in-vivo structures by exploring other fusion strategies (such as union and interpolation) or denoising modules with specific regularizations [18,24].

Acknowledgments. This work was supported by the Natural Science Foundation of China (No. 12074258).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Hennersperger, C., Baust, M., Mateus, D., Navab, N.: Computational Sonography. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 459–466. *Lecture Notes in Computer Science*, Springer International Publishing, Cham (2015). https://doi.org/10.1007/978-3-319-24571-3_55
2. Mohamed, F., Siang, C.V.: A Survey on 3D Ultrasound Reconstruction Techniques. *IntechOpen* (Apr 2019). <https://doi.org/10.5772/intechopen.81628>
3. Luo, M., Yang, X., Wang, H., Dou, H., Hu, X., Huang, Y., Ravikumar, N., Xu, S., Zhang, Y., Xiong, Y., Xue, W., Frangi, A.F., Ni, D., Sun, L.: RecON: Online learning for sensorless freehand 3D ultrasound reconstruction. *Medical Image Analysis* **87**, 102810 (Jul 2023). <https://doi.org/10.1016/j.media.2023.102810>
4. Luo, M., Yang, X., Yan, Z., Li, J., Zhang, Y., Chen, J., Hu, X., Qian, J., Cheng, J., Ni, D.: Multi-IMU with Online Self-consistency for Freehand 3D Ultrasound Reconstruction. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 342–351. *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-43907-0_33
5. Huang, Q., Luo, H., Yang, C., Li, J., Deng, Q., Liu, P., Fu, M., Li, L., Li, X.: Anatomical prior based vertebra modelling for reappearance of human spines. *Neurocomputing* **500**, 750–760 (Aug 2022). <https://doi.org/10.1016/j.neucom.2022.05.033>
6. Chen, H., Kumaralingam, L., Lou, E.H.M., Punithakumar, K., Li, J., Pham, T.T., Le, L.H., Zheng, R.: Neural Implicit Surface Reconstruction of Freehand 3D Ultrasound Volume with Geometric Constraints (Jan 2024). <https://doi.org/10.48550/arXiv.2401.05915>

7. Park, C.K., Trumppour, T., Gyacskov, I., Bax, J.S., Tessier, D., Gardi, L., Ico, M., Fenster, A.: Improving three-dimensional automated breast ultrasound resolution with orthogonal images. In: Bottenus, N., Boehm, C. (eds.) *Medical Imaging 2023: Ultrasonic Imaging and Tomography*. p. 3. SPIE, San Diego, United States (Apr 2023). <https://doi.org/10.1117/12.2653141>
8. Wright, R., Gomez, A., Zimmer, V.A., Toussaint, N., Khanal, B., Matthew, J., Skelton, E., Kainz, B., Rueckert, D., Hajnal, J.V., Schnabel, J.A.: Fast fetal head compounding from multi-view 3D ultrasound. *Medical Image Analysis* p. 102793 (Mar 2023). <https://doi.org/10.1016/j.media.2023.102793>
9. Wysocki, M., Azampour, M.F., Eilers, C., Busam, B., Salehi, M., Navab, N.: Ultra-NeRF: Neural Radiance Fields for Ultrasound Imaging (Jan 2023). <https://doi.org/10.48550/arXiv.2301.10520>
10. Gaits, F., Mellado, N., Basarab, A.: Ultrasound volume reconstruction from 2D Freehand acquisitions using neural implicit representations. In: 21st IEEE International Symposium on Biomedical Imaging (ISBI 2024). p. à paraître. IEEE Signal Processing Society and IEEE Engineering in Medicine and Biology Society, Athènes, Greece (May 2024)
11. Zha, R., Cheng, X., Li, H., Harandi, M., Ge, Z.: EndoSurf: Neural Surface Reconstruction of Deformable Tissues with Stereo Endoscope Videos. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 13–23. Lecture Notes in Computer Science, Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-43996-4_2
12. Batlle, V.M., Montiel, J.M.M., Fua, P., Tardós, J.D.: LightNeuS: Neural Surface Reconstruction in Endoscopy Using Illumination Decline. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 502–512. Lecture Notes in Computer Science, Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-43999-5_48
13. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 165–174 (2019)
14. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS* (2021)
15. Amiranashvili, T., Lüdke, D., Li, H.B., Zachow, S., Menze, B.H.: Learning continuous shape priors from sparse data with neural implicit functions. *Medical Image Analysis* **94**, 103099 (May 2024). <https://doi.org/10.1016/j.media.2024.103099>
16. Wiesner, D., Suk, J., Dummer, S., Nečasová, T., Ulman, V., Svoboda, D., Wolterink, J.M.: Generative modeling of living cells with SO(3)-equivariant implicit neural representations. *Medical Image Analysis* **91**, 102991 (Jan 2024). <https://doi.org/10.1016/j.media.2023.102991>
17. Baorui, M., Zhizhong, H., Yu-Shen, L., Matthias, Z.: Neural-pull: Learning signed distance functions from point clouds by learning to pull space onto surfaces. In: *International Conference on Machine Learning (ICML)* (2021)
18. Marschner, Z., Sellán, S., Liu, H.T.D., Jacobson, A.: Constructive Solid Geometry on Neural Signed Distance Fields. In: *SIGGRAPH Asia 2023 Conference Papers*. pp. 1–12. ACM, Sydney NSW Australia (Dec 2023). <https://doi.org/10.1145/3610548.3618170>

19. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics* **21**(4), 163–169 (Aug 1987). <https://doi.org/10.1145/37402.37422>
20. Sharma, G., Goyal, R., Liu, D., Kalogerakis, E., Maji, S.: CSGNet: Neural Shape Parser for Constructive Solid Geometry. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5515–5523 (2018)
21. Yang, H., Sun, Y., Sundaramoorthi, G., Yezzi, A.: StEik: Stabilizing the Optimization of Neural Signed Distance Functions and Finer Shape Representation (Nov 2023). <https://doi.org/10.48550/arXiv.2305.18414>
22. Levenberg, K.: A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics* **2**(2), 164–168 (1944). <https://doi.org/10.1090/qam/10666>
23. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. pp. 5105–5114. NIPS’17, Curran Associates Inc., Red Hook, NY, USA (Dec 2017)
24. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. In: *Proceedings of Machine Learning and Systems 2020*, pp. 3569–3579 (2020)