

Implicit Neural Representations for Breathing-compensated Volume Reconstruction in Robotic Ultrasound

Yordanka Velikova^{1*}, Mohammad Farid Azampour^{1,3,5*}, Walter Simson², Marco Esposito⁴, and Nassir Navab¹

Abstract—Ultrasound (US) imaging is widely used in diagnosing and staging abdominal diseases due to its lack of non-ionizing radiation and prevalent availability. However, significant inter-operator variability and inconsistent image acquisition hinder the widespread adoption of extensive screening programs. Robotic ultrasound systems have emerged as a promising solution, offering standardized acquisition protocols and the possibility of automated acquisition. Additionally, these systems enable access to 3D data via robotic tracking, enhancing volumetric reconstruction for improved ultrasound interpretation and precise disease diagnosis.

However, the interpretability of 3D US reconstruction of abdominal images can be affected by the patient’s breathing motion. This study introduces a method to compensate for breathing motion in 3D US compounding by leveraging implicit neural representations. Our approach employs a robotic ultrasound system for automated screenings. To demonstrate the method’s effectiveness, we evaluate our proposed method for the diagnosis and monitoring of abdominal aorta aneurysms as a representative use case.

Our experiments demonstrate that our proposed pipeline facilitates robust automated robotic acquisition, mitigating artifacts from breathing motion, and yields smoother 3D reconstructions for enhanced screening and medical diagnosis.

I. INTRODUCTION

Medical ultrasound (US) is widely employed in clinical settings to diagnose conditions of internal tissues and organs. Its real-time imaging offers clinicians instant feedback. In comparison to other imaging methods like X-ray or CT, US presents several advantages: it’s radiation-free, portable, and cost-effective [1]. This, combined with its high soft-tissue contrast, ensures both safety and efficiency in patient care [2]. Ultrasound has become an essential tool for screening and monitoring of the abdominal region [3], [4]. In particular, for individuals aged 50 and above, routine US screening of aorta is recommended, as they might be at risk of developing aneurysms while being asymptomatic. Abdominal Aortic Aneurysm (AAA) is a condition marked by an enlargement in the aorta, which can lead to serious health implications if left untreated [5]. If an enlargement is detected, the




Fig. 1. Overview of the proposed system: a) we use a robotic arm manipulator with a convex ultrasound probe attached to the end-effector to scan the patient, b) Live B-mode, c) Intermediate image, d) Aorta segmentation, e) 3D Reconstruction.

frequency of follow-up screens increases based on its growth rate, ensuring timely interventions [6]. Consistent ultrasound screenings have shown a significant reduction of premature death from AAA in men aged 65 and older [7]. The primary goal of these screenings is to accurately measure and assess the aneurysm’s size and dimensions.

However, the repeatability of image acquisitions can vary largely between clinicians [8] as image quality is strongly based on clinician experience and factors such as probe placement and contact during scanning can reduce inter-acquisition consistency. These constraints highlight the necessity to automate US acquisitions. Recent works have exploited the use of robotic systems for manipulating the US probe autonomously, thus providing standardized imaging results and reducing operator dependency [9], [10]. Solutions for automating US acquisitions have been a focus of study for more than two decades, ranging from the creation of customized end-effectors for guiding US probes to the introduction of complete imaging and navigation systems [11].

Recent work demonstrates an automated workflow for screening of tubular structures using a robotic US system [12]. It is based on real-time image feedback and employs active perception to control robotic motion. They estimate the centerline and radius of a tubular structure and perform

*Shared contribution.

¹ Computer Aided Medical Procedures, Technical University of Munich, Munich, Germany email: dani.velikova@tum.de

² Department of Radiology, Stanford University School of Medicine, Stanford, CA, USA

³ Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran

⁴ ImFusion GmbH, Munich, Germany

⁵ Munich Center for Machine Learning

Github Repository: <https://github.com/danivelikova/INR-3DUS>

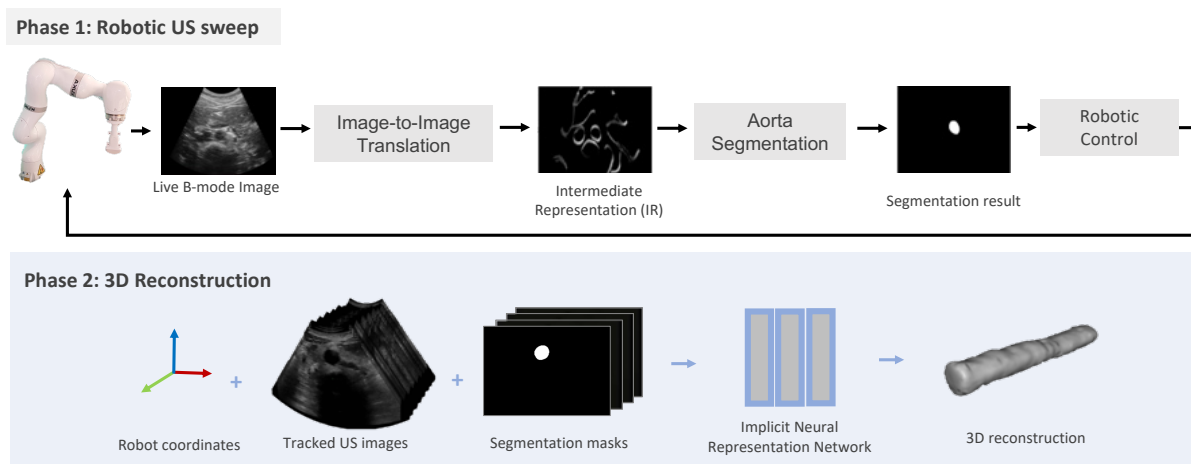


Fig. 2. Overview of the pipeline. Phase 1: Robotic ultrasound acquisition with real-time image segmentation; the robot’s trajectory is adjusted based on the segmentation. Phase 2: The acquired sweep and segmentation train the INR model, which is sampled to produce a dense aorta point cloud. Post-processing of this cloud yields the final mesh.

robotic control in real-time in response to the incoming US images. In addition, they extract the 3D point cloud of the segmented tube and calculate its centerline. However, the method lacks generalizability on in-vivo images.

Furthermore, employing robotic ultrasound enables 3D reconstruction of the aorta. Access to three-dimensional representation can help clinicians compare the progression of aneurysms. Without explicit access to 3D information, physicians need to mentally reconstruct the 3D anatomy, which is a demanding and challenging task, especially for less experienced ones. Having a 3D reconstruction of the aorta volume can give more insights to clinicians and help them to evaluate the diameter of the aneurysm not only in the axial 2D plane but also in any arbitrary tilted plane.

However, creating a consistent US compounding faces hurdles like patient breathing motion. Virga et al. [13] proposed solutions to automate trajectory planning and address motion from breathing. Despite the strides, motion artifacts remain unresolved. In robotics, breathing compensation remains a prominent topic, with gating control techniques and user-in-the-loop systems ensuring stable acquisitions [14], [15].

Lately, the rise of implicit neural representation (INR) stands out, especially for rendering continuous 3D volume representations. In medical imaging, INRs encapsulate anatomical data compactly while preserving resolution. Typically, a Multilayer Perceptron (MLP) in INRs infers pixel or voxel intensity based on their position. Adding to the advancements in this domain, Sitzmann et al.’s SIREN [16], which uses periodic functions over traditional activations, has achieved impressive results in capturing scene nuances. Khan et al. [17] utilized SIREN to obtain a segmentation mask that converges quickly and, by definition, supports super-resolution. Gu et al. [18] employed SIREN for US imaging by integrating the SIREN-based representation with meta-learning. This approach facilitates learning an INR for a new patient using a set of sparse inputs. Continuing the application of SIREN in US imaging, Song et al. [19] aimed

to enhance the segmentation of the carotid artery in US images. They hypothesized that given an US sweep and the segmentation output from a trained 2D U-net, the MLP could infer both pixel intensity and label. Their findings suggest that an INR trained in this manner can produce smoother label maps, preserving the curvature of the carotid artery, even with tracking data errors. In our work, we also delve into INR’s capabilities to enhance our robotic US acquisition and segmentation. In a recent work, exploiting the power of INRs combined with a physics-based rendering module, Wysocki et al. [20] showed how Ultra-NeRF can generate highly accurate B-mode ultrasound images through learning US-specific parameter maps which account for view-dependent variations inherent to ultrasound imaging.

In this work, we present a robotic US system for autonomous aortic scan with breathing compensation for improved 3D US volume reconstruction. The robot control strategy does not require pre-operative registration with other imaging modalities in order to perform the navigation, nor external tracking systems or cameras, but only a single manual initialization step by the medical personnel. Once the probe is placed such that the aorta is visible in the US image, it can follow it to cover the whole region of interest without further intervention. By leveraging a recent method for aorta segmentation incorporating the concept of intermediate representation [21], it can be directly applied to real US images, while being trained only on simulated ones.

Our key contribution is the development of a breathing-compensated US reconstruction method. This technique harnesses an INR model, trained using the gathered tracked US images and their respective segmentation masks. This network not only crafts a patient-tailored continuous 3D function from a subset of the images but also compensates for noisy or missing segmentation frames to generate a smooth, finely detailed 3D aorta reconstruction.

II. METHODOLOGY

The proposed robotic ultrasound screening system is composed of a robotic manipulator and a convex US probe rigidly attached to the robot’s end-effector, see Fig 1. The overall workflow setup of the method consists of two main steps: robotic ultrasound sweep and 3D reconstruction, as shown in Fig. 2, which are explained in detail below.

A. Image acquisition and segmentation

Prior to beginning the acquisition procedure, the probe must be manually placed on top of the patient’s abdomen, usually in the upper part, where the aorta is visible in the US image. Our navigation during the acquisition depends on robust real-time segmentation of the aorta. To achieve this, we utilize a recently proposed pipeline called CACTUSS [21], which is tailored for the task of aorta segmentation in US images. CACTUSS proposes a novel concept of IR between CT and US, and it is trained on a large set of simulated intermediate images from CT labelmaps. IR space is introduced to transfer both simulated and real images to this space by using a domain adaptation network. Then, a segmentation network is trained on the transferred simulated images in the IR space. Training in the IR space allows for direct segmentation of real US images, even though the network has been trained only on simulated ones, and it shows great generalization capabilities across patients.

As depicted in Fig. 2 each US B-mode image is first passed through the pre-trained image-to-image translation network, which translates it to the IR image space. The translated image is then passed to the second network, which gives the segmentation result. Finally, the prediction output is used to control the robot’s movement in real-time, enabling visual servoing and autonomous US acquisition.

B. Navigation pipeline

Along the anteroposterior axis, the position of the end-effector is controlled using a constant force, ensuring that the probe maintains a consistent connection to the skin for optimal image quality. For longitudinal and transverse axes, a visual-based control pipeline is used to update the position of the robot’s end-effector. During the scan, a sequence of two-dimensional B-mode frames is collected from the US machine, and the pose is updated in real-time based on the output of the segmentation algorithm.

To avoid navigating to a wrong position as a result of incorrect segmentation output, we ensure that the aorta is always maintained in the center of the image. We keep track of the center of the previous segmentations, and the robot moves along the transverse axis to ensure that the aorta is always centered. Consequently, a step on the longitudinal axis is performed, and the robot moves until no aorta is visible, allowing a complete sweep to be obtained. To enable our control pipeline, it is necessary that the images are transformed into the coordinate frame of the robot. For every image frame, the full chain of transformations from robot base to image frame is calculated (see Fig. 3):

$${}^W T_I = {}^W T_E \cdot {}^E T_P \cdot {}^P T_I \quad (1)$$

where ${}^P T_I$ maps from the origin of the image to the probe by converting image pixels to millimeters, ${}^E T_P$ is the transformation from the tip of the probe to the end effector, which is known from the 3D model of the mount, ${}^W T_E$ is the transformation from the robot’s end-effector to the robot base and is calculated using the robot’s kinematic model, which is given by the manufacturer.




Fig. 3. Robot transformations.

Calibration To map between pixels and millimeters, we use the dimensions of the image in pixels ($W \times H$) and the opening angle of the probe θ . The imaging depth d in mm is typically specified by the user only once in the beginning. The mapping can be calculated by the matrix ${}^W T_I$:

$${}^W T_I = \begin{bmatrix} \frac{w}{W} & 0 & 0 & -\frac{w}{2} \\ 0 & 0 & -1 & 0 \\ 0 & \frac{d}{H} & 0 & -y \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (2)$$

where W and H are the width and height of the image in pixels, w is the corresponding offset in mm along the width of the image, derived by $w = 2R \sin(\theta/2)$, where $R = d + r$ is the radius defined as the sum of the depth of the image d and the radius r from the origin of the probe to its tip. Additionally, a small offset h along the height of the image is calculated from the tip of the probe to the origin of the image: $h = r - \cos(\theta/2)r$.

C. 3D Volume Reconstruction via the INR Model

To utilize INR-based models for volume reconstruction, we first acquire a set of observations from the 3D space, along with their corresponding poses. In our study, the observation refers to the US sweep, while the pose is derived from the robotic tracking stream. The aim of our INR model is not just to encode the anatomical intensity values but also the semantic label of the aorta. Analogous to our observation in the intensity domain, we utilize the segmentation results obtained during the data acquisition phase (refer to section II-A) as our observations from the semantic domain. Similar to [19], our INR model has been adapted to produce both volume intensity and semantic attributes. Through this approach, we ensure the learning of a continuous smooth function in the semantic domain, which solidifies attributes

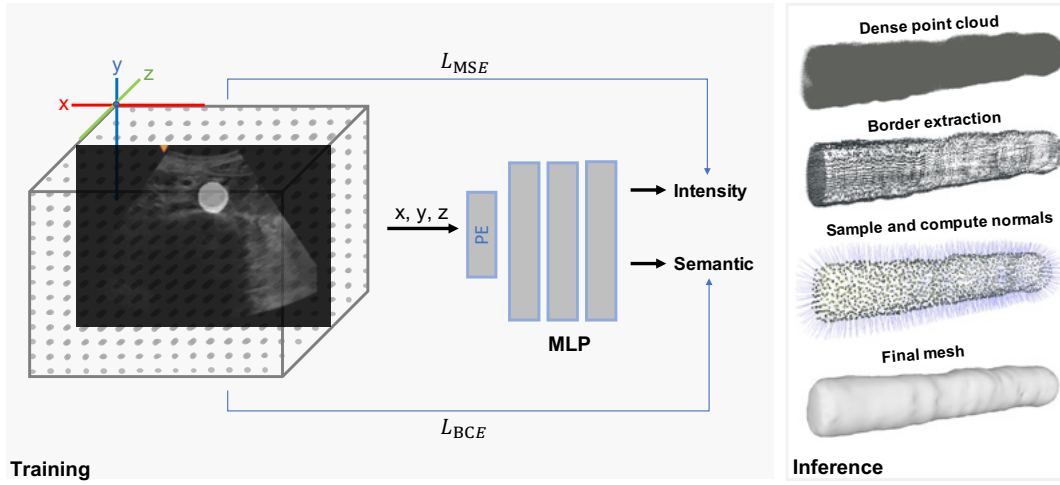


Fig. 4. Schematic of the INR model. The model accepts a normalized 3D voxel position (ranging from -1 to 1) as input and produces two outputs: the intensity and the semantic label of the voxel. For training, we leverage the ultrasound (US) sweep alongside its corresponding segmentation. During inference, we sample across all points, yielding a dense point cloud from the semantic results. This is subsequently post-processed to generate the final mesh.

vital for vessel segmentation. The INR model produces the output:

$$\mathbf{y} = (s, i) \in \mathcal{Y} \subseteq \mathbb{R}^{M+1} \quad (3)$$

where $s \in \mathcal{S} = \mathbb{R}^M$ represents the semantic vector and M denotes the number of semantic classes. $i \in \mathbb{R}$ stands for the volume intensity. An INR model is trained to infer \mathbf{y} for each given position in the 3D domain. In particular, if we define a location vector as:

$$\mathbf{x} = (x, y, z) \in \mathcal{X} \subseteq \mathbb{R}^3, \quad (4)$$

The trained INR is function that maps \mathbf{x} to our desired output:

$$\mathbf{y} = f_{\Theta}(\mathbf{x}) \quad (5)$$

Here, f_{Θ} designates the overall INR model, with Θ representing its trained weights. When trained, f_{Θ} encodes the entire intensity and semantic information as a continuous function that can be sampled at any resolution. The design of our network can be found in Fig. 4.

To preserve the high-frequency information of the volume intensity, we apply Position Encoding (PE) [16], which maps each 3D coordinate into a more expansive dimensional space. The encoding of \mathbf{x} can be achieved with the function $\phi : \mathbb{R} \rightarrow \mathbb{R}^{2L}$:

$$\gamma(p) = (\sin(2^0 \pi p), \cos(2^0 \pi p), \dots, \sin(2^{L-1} \pi p), \cos(2^{L-1} \pi p)) \quad (6)$$

Where p refers to the normalized values (ranging from -1 to 1) of the 3D positions (x, y, z) respectively.

The main component of f_{Θ} is an MLP. In this MLP, the path for achieving the two desired outputs from f_{Θ} , meaning the intensity and the semantic information are shared except for the last linear layer. As in [16], to depict the high-frequency domain, the SIREN activation function is employed. Further, we define the volume intensity loss as \mathcal{L}_i and the semantic loss as \mathcal{L}_s :

$$\mathcal{L}_i = \sum_{x \in \mathcal{X}} \|\hat{i}(x) - i(x)\|^2 \quad (7)$$

$$\mathcal{L}_s = - \sum_{x \in \mathcal{X}} \left[\sum_{m=1}^M s^m(x) \log \hat{s}^m(x) \right]. \quad (8)$$

In the above, \mathcal{X} represents the input location space, and $\hat{i}(x)$ and $i(x)$ denote the predicted and ground truth volume intensities, respectively. Similarly, $\hat{s}^l(x)$ and $s^l(x)$ indicate the semantic probability for class l in the predictions and ground truth, respectively. Combining these losses, the overall training loss can be defined as:

$$\mathcal{L} = \mathcal{L}_i + \mathcal{L}_s \quad (9)$$

During training, our semantic observations are derived from the CACTUSS segmentation pipeline. However, due to breathing motions and inaccuracies in translating to the IR space, some slices might contain erroneous segmentations. Including these in training could compromise the accuracy of our INR model. To mitigate this, for each slice, we retain only the largest connected component, setting all else to zero. Moreover, by monitoring the aorta's radius across previous slices and employing a moving average approach, we obtain an estimated radius for the current slice. If there's a significant discrepancy between this estimate and the actual value, we label the slice as incorrectly segmented.

For each slice determined to be correctly labeled, both semantic and intensity losses are calculated for its respective voxels. For those identified as mislabeled, only the intensity loss is computed. This strategy harnesses valuable geometric information from correct slices while sidestepping potential pitfalls posed by erroneous semantic labels.

Once the INR network is trained, inference is performed as illustrated in Fig. 4 to achieve the final 3D volume reconstruction. We begin by sampling new points to form

a dense point cloud from our predictions. We then identify boundary points by fitting a convex hull to each slice. Using the furthest point sampling method, points are sampled, and their normals are computed. Lastly, we employ the Poisson surface reconstruction [22] to produce the final mesh.

III. SETUP

1) *Hardware*: The system consists of a curvilinear transducer 5C1 ACUSON Juniper US machine (Siemens Healthineers, Germany) which is attached to the end-effector of a robotic manipulator (KUKA LBR iiwa 14 R820, KUKA Roboter GmbH, Germany) using a 3D-printed probe holder. We maintain a constant low force (5N) and stiffness(200 N/m) to allow comfort for the patient during breathing. The robot is controlled via iiwa stack and ROS. To access the images, a frame grabber (Epiphan Video, Canada) is used and visualized in real-time on ImFusion Suite (ImFusion GmbH, Germany). The image depth is set to 10cm. For working with the data, ethics approval was obtained.¹

2) *Training Details*: The INR network consists of an MLP with 6 layers, each of size 256, and employs positional encoding of length 10. The network is trained for each dataset individually. The batch size, in our case, is the number of image slices and is set to 1. We trained for 1000 epochs with a learning rate of 1e-3 and Adam optimizer.

IV. EVALUATION

To achieve an optimal 3D reconstruction free from breathing-induced distortions, we take into account images captured during the exhale phase only. Breathing movements can introduce errors in segmentation masks and generate noise. The stability of images during exhale allows for more accurate segmentation prediction. We conducted two primary experiments: breath-hold and normal breathing, which are detailed in the subsequent sections. In both experiments, we compared our novel INR pipeline with conventional techniques. The breath-hold experiment, being user-controlled to capture images only during the exhale phase, acts as a reference point for our model evaluation. Fig. 5 illustrates both modes. Additionally, we compute the Laplacian and report the overall roughness of the final 3D volume surfaces.

A. Breath-hold with button

In this experiment, a volunteer is instructed to press and hold a button after a deep exhale. As this action is performed, the robot moves, and the US images and segmentation masks are stacked with their 3D location based on the robotic tracking data. In Fig. 5, the left side shows the trajectory of the images in 3D, where a straight trajectory is observed during the button hold. Upon releasing the button, the robot halts, allowing the volunteer to breathe freely. This pause is marked by the peaks in the trajectory, and no images are captured during this time.

We train our INR model using the acquired 2D B-modes, their corresponding segmentations, and 3D coordinates. For

this experiment, every third image from the scan is used to build the training set. We compare the results of our model to a conventional linear interpolation method, implemented in ImFusion Suite².

B. Free breathing

The second experiment is performed during free breathing, without any restrictions on the user. The right side of Fig. 5 displays the breathing trajectory. Once the sweep is completed, the segmentations are again compiled in 3D. The top of Fig. 7 portrays the 3D volume right after the acquisition without any breathing adjustments, showing distortions due to respiratory movements. To derive a high-quality 3D reconstruction and compare it with the breath-hold experiment, we subsequently extract images only from the exhale phase (red). Using the 3D trajectory from the robot, we identify all local minima, selecting only the images within a narrow range surrounding each minima. These images are used as training data for the INR network. The results from the network are again compared with the conventional linear interpolation technique.

C. Second volunteer

We demonstrate the generalizability of our algorithm using a US sweep acquired from a second volunteer. It's important to note that INR-based methods are trained specifically for each patient, necessitating retraining from scratch for every new individual. Since the INR is patient-specific, the generalization of our full pipeline depends on the capabilities of the underlying segmentation method. Our segmentation method (CACTUSS) has been tested on multiple subjects, and, as demonstrated in the original paper, shown to be generalized to unseen patients. By testing on a second volunteer, we confirm that while the INR network needs individualized training, the full pipeline itself is broadly applicable. The qualitative results for this volunteer, under both breath-hold and free-breathing scenarios, are presented in Fig. 6 and 7 respectively, and quantitative results are reported in Table I.

V. RESULTS AND DISCUSSION

Fig. 6 illustrates the 3D volumes from the breath-hold experiment on both volunteers using both conventional and INR methods. Our technique provides a notably smoother 3D reconstruction than traditional approaches, which often fails when interpolating larger gaps. By learning a patient-specific continuous function, our method smoothly interpolates between frames, resulting in an enhanced 3D reconstruction.

In Fig. 7, the results from the free-breathing experiment are presented. The top two volumes show the 3D aorta volume without any breathing compensation and the parts from exhale phase only are shown in red. The two lowermost 3D reconstructions compare traditional linear interpolation with the result from our network-driven approach. Once again, our method yields a better reconstruction and smoothly fills in the gaps. In comparison, the linear interpolation approach

¹Approval was obtained from the ethics committee of the Technical University of Munich.

²ImFusion GmbH, Munich, Germany

TABLE I

COMPARISON OF MESH SMOOTHNESS UNDER TWO BREATHING MODES: "BREATH-HOLD BUTTON" AND "FREE BREATHING". RESULTS INDICATE THAT OUR PIPELINE PRODUCES SMOOTH OUTPUTS FOR BOTH MODES. FURTHER, UTILIZING THE INR WITH FILTERED INPUT ENHANCES THE FINAL OUTPUT'S SMOOTHNESS.

Method	Breath-Hold Button		Free breathing	
	no INR	with INR	no INR	with INR
Laplacian average V1	0.627	0.296	0.517	0.251
Laplacian median V1	0.452	0.298	0.396	0.247
Laplacian average V2	0.639	0.432	0.641	0.424
Laplacian median V2	0.460	0.430	0.506	0.416




Fig. 5. Left:breath-hold mode, right: free breathing mode.

struggles with larger gaps, resulting in pronounced step edges in the final reconstruction.

Table I provides a quantitative assessment of mesh smoothness under the two breathing modes: breath-hold button and free breathing. The metrics used are the average and median magnitudes of the Laplacian vectors over the entire mesh, which serve as indicators of the overall surface roughness or smoothness. Notably, for both breathing modes, the output of the INR model has lower Laplacian roughness values for both volunteers, indicating smoother meshes. Specifically, under the breath-hold mode of the first volunteer, the average roughness decreased from 0.627 without INR to 0.296 with INR, and similarly, for the free-breathing mode, from 0.517 to 0.251. The results for the second volunteer follow a similar trend. Moreover, both the average and median roughness values are similar, which suggests a more uniform curvature distribution across the mesh. These findings highlight the efficacy of our INR pipeline in enhancing the smoothness of the reconstructed meshes, irrespective of the breathing mode.

From our evaluation, it's evident that prioritizing images from the exhale phase significantly enhances the quality of 3D reconstructions, mitigating the distortions and noise introduced by breathing movements. Our innovative INR pipeline, when compared to traditional methods, consistently showcased superior performance in both the breath-hold and free-breathing experiments. The breath-hold experiment, which closely mirrors an idle state, served as an effective benchmark, reinforcing the efficacy of our approach. Particularly in the free-breathing experiment, our method adeptly handled the challenges posed by respiratory movements, producing smoother reconstructions and effectively interpolating missing gaps. In contrast, conventional linear




Fig. 6. 3D volumes from two volunteers from conventional (top) and INR method (bottom) from the Breath-Hold Button experiment. Conventional methods fail in interpolating larger gaps, while our INR-based approach offers smoother interpolation between frames.




Fig. 7. Top image: the 3D volumes acquired after the free breathing experiment from two volunteers, showing distortions due to respiratory movements. Red: images only from the exhale phase, used for training of the INR. Bottom image: 3D volume generated with conventional method, 3D volume reconstruction result from our network.

interpolation techniques exhibited limitations, especially in handling larger gaps, leading to pronounced step edges in the reconstructions. These findings underscore the potential of our INR pipeline in advancing the field of 3D ultrasound imaging, especially in scenarios where breathing-induced distortions are a concern. Future work aims to enhance the pipeline's adaptability to diverse anatomical variations and its compatibility with various ultrasound probes and machines, broadening its applicability in clinical settings.

VI. CONCLUSIONS

The advancements in ultrasound imaging and its advantages over other imaging modalities have led to its wide usage in diagnostics and screening. However, challenges such as inter-operator variability and breathing-induced motion limit its applications. In this work, we introduced a novel approach leveraging INRs to effectively counteract the distortions caused by breathing when constructing a 3D volume. Furthermore, we combined this method with a robotic ultrasound system, allowing for standardized acquisitions. The screening of the abdominal aorta is targeted as a use case to verify the efficacy of the approach. The resultant smoother 3D reconstructions, as showcased by our experiments, can help clinicians visualize and inspect the progression of a disease and perform diameter measures more easily. The findings from our experiments underscore the potential of our approach to enhance clarity and improve diagnostic accuracy using 3D ultrasound images.

REFERENCES

- [1] Vivien Gibbs, David Cole, and Antonio Sassano. *Ultrasound physics and technology: how, why and when*. Elsevier Health Sciences, 2011.
- [2] Thomas L Szabo. *Diagnostic ultrasound imaging: inside out*. Academic press, 2004.
- [3] W. Schäberle, L. Leyerer, Wilma Schierling, and Karin Pfister. Ultrasound diagnostics of the abdominal aorta: English version. *Gefäßchirurgie*, 20:22–27, 01 2015.
- [4] L Beales, Stephen Wolstenhulme, Joseph Evans, Robert West, and David Scott. Reproducibility of ultrasound measurement of the abdominal aorta. *The British journal of surgery*, 98:1517–25, 11 2011.
- [5] Brant W Ullery, Richard L Hallett, and Dominik Fleischmann. Epidemiology and contemporary management of abdominal aortic aneurysms. *Abdominal Radiology*, 43:1032–1043, 2018.
- [6] TC Hartshorne, CN McCollum, JJ Earnshaw, J Morris, and A Nasim. Ultrasound measurement of aortic diameter in a national screening programme. *European Journal of Vascular and Endovascular Surgery*, 42(2):195–199, 2011.
- [7] David Brewster, Jack Cronenwett, John Hallett, K Johnston, William Krupski, and Jon Matsumura. Guidelines for the treatment of abdominal aortic aneurysms: Report of a subcommittee of the joint council of the american association for vascular surgery and society for vascular surgery. *Journal of vascular surgery : official publication, the Society for Vascular Surgery [and] International Society for Cardiovascular Surgery, North American Chapter*, 37:1106–17, 06 2003.
- [8] Keyu Li, Yangxin Xu, and Max Q-H Meng. An overview of systems and techniques for autonomous robotic ultrasound acquisitions. *IEEE Transactions on Medical Robotics and Bionics*, 3(2):510–524, 2021.
- [9] Pierre Chatelain, Alexandre Krupa, and Nassir Navab. Confidence-driven control of an ultrasound probe: Target-specific acoustic window optimization. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3441–3446. IEEE, 2016.
- [10] Christoph Hennemperger, Bernhard Fuerst, Salvatore Virga, Oliver Zettinig, Benjamin Frisch, Thomas Neff, and Nassir Navab. Towards mri-based autonomous robotic us acquisitions: A first feasibility study. *IEEE Transactions on Medical Imaging*, 36(2):538–548, 2017.
- [11] Zhongliang Jiang, Septimiu E Salcudean, and Nassir Navab. Robotic ultrasound imaging: State-of-the-art and future perspectives. *Medical Image Analysis*, page 102878, 2023.
- [12] Zhongliang Jiang, Zhenyu Li, Matthias Grimm, Mingchuan Zhou, Marco Esposito, Wolfgang Wein, Walter Stechele, Thomas Wendler, and Nassir Navab. Autonomous robotic screening of tubular structures based only on real-time ultrasound imaging feedback. *IEEE Transactions on Industrial Electronics*, 69:7064–7075, 2020.
- [13] Salvatore Virga, Oliver Zettinig, Marco Esposito, Karin Pfister, Benjamin Frisch, Thomas Neff, Nassir Navab, and Christoph Hennemperger. Automatic force-compliant robotic ultrasound screening of abdominal aortic aneurysms. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 508–513. IEEE, 2016.
- [14] K.K. Bredahl, Nikolaj Eldrup, C Meyer, Jonas Eiberg, and Henrik Sillesen. Reproducibility of ecg-gated ultrasound diameter assessment of small abdominal aortic aneurysms. *European journal of vascular and endovascular surgery : the official journal of the European Society for Vascular Surgery*, 45, 01 2013.
- [15] Caroline Vienne and Alexandre Krupa. Intensity-based ultrasound visual servoing: Modeling and validation with 2-d and 3-d probes. *Robotics, IEEE Transactions on*, 29:1003–1015, 08 2013.
- [16] 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.
- [17] Muhammad Osama Khan and Yi Fang. Implicit neural representations for medical imaging segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 433–443. Springer, 2022.
- [18] Ang Nan Gu, Purang Abolmaesumi, Christina Luong, and Kwang Moo Yi. Representing 3d ultrasound with neural fields. In *Medical Imaging with Deep Learning*, 2022.
- [19] Sheng Song, Yunqian Huang, Jiawen Li, Man Chen, and Rui Zheng. Development of implicit representation method for freehand 3d ultrasound image reconstruction of carotid vessel. In *2022 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4. IEEE, 2022.
- [20] Magdalena Wysocki, Mohammad Farid Azampour, Christine Eilers, Benjamin Busam, Mehrdad Salehi, and Nassir Navab. Ultra-nerf: neural radiance fields for ultrasound imaging. In *Medical Imaging with Deep Learning*, pages 382–401. PMLR, 2024.
- [21] Yordanka Velikova, Walter Simson, Mehrdad Salehi, Mohammad Farid Azampour, Philipp Paprottka, and Nassir Navab. Cactus: Common anatomical ct-us space for us examinations. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pages 492–501, Cham, 2022. Springer Nature Switzerland.
- [22] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, page 0, 2006.