

Gaze-Guided Robotic Vascular Ultrasound Leveraging Human Intention Estimation

Yuan Bi^{1,2}, Yang Su¹, Nassir Navab^{1,2}, *Fellow, IEEE*, and Zhongliang Jiang^{1,2}

Abstract—Medical ultrasound has been widely used to examine vascular structure in modern clinical practice. However, traditional ultrasound examination often faces challenges related to inter- and intra-operator variation. The robotic ultrasound system (RUSS) appears as a potential solution for such challenges because of its superiority in stability and reproducibility. Given the complex anatomy of human vasculature, multiple vessels often appear in ultrasound images, or a single vessel bifurcates into branches, complicating the examination process. To tackle this challenge, this work presents a gaze-guided RUSS for vascular applications. A gaze tracker captures the eye movements of the operator. The extracted gaze signal guides the RUSS to follow the correct vessel when it bifurcates. Additionally, a gaze-guided segmentation network is proposed to enhance segmentation robustness by exploiting gaze information. However, gaze signals are often noisy, requiring interpretation to accurately discern the operator’s true intentions. To this end, this study proposes a stabilization module to process raw gaze data. The inferred attention heatmap is utilized as a region proposal to aid segmentation and serve as a trigger signal when the operator needs to adjust the scanning target, such as when a bifurcation appears. To ensure appropriate contact between the probe and surface during scanning, an automatic ultrasound confidence-based orientation correction method is developed. In experiments, we demonstrated the efficiency of the proposed gaze-guided segmentation pipeline by comparing it with other methods. Besides, the performance of the proposed gaze-guided RUSS was also validated as a whole on a realistic arm phantom with an uneven surface.

Index Terms—Robotic ultrasound, Gaze tracker, Gaze-guided system, Ultrasound segmentation

I. INTRODUCTION

Medical ultrasound, valued for its non-invasiveness, portability, real-time performance, and affordability, is widely used in clinical practice for screening and intra-operative guidance. However, traditional free-hand ultrasound suffers from inter- and intra-operator variances. Image quality depends on acquisition parameters like contact forces, angles, and probe positioning [1], [2], making it highly operator-dependent and reducing result reproducibility [3]. To tackle such a dilemma, robotic ultrasound systems (RUSS) offer a promising solution to address these challenges [4]–[6].

Manuscript received: September 10, 2024; Revised: December 7, 2024; Accepted: January 15, 2025.

This paper was recommended for publication by Editor Jessica Burgner-Kahrs upon evaluation of the Associate Editor and Reviewers’ comments.

This work was partially supported by SINO-German Mobility Project (M-0221). (Corresponding Author: Zhongliang Jiang, zl.jiang@tum.de)

¹Y. Bi, Y. Su, N. Navab, and Z. Jiang are with the Chair for Computer-Aided Medical Procedures and Augmented Reality, Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany

²Y. Bi, N. Navab, and Z. Jiang are with the Munich Center for Machine Learning, Arcisstraße 21, 80333 München, Germany

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

Owing to the fast development in the field of robotics, robots have shown their superiority over humans in terms of stability and precision. Working together with human experts as colleagues, the interactive RUSS can largely release the burden of cumbersome acquisition from sonographers and greatly enhance the reproducibility of the imaging process [7]–[9]. RUSSs are used for autonomous vascular ultrasound screening [10], [11], but the complexity of vessel distributions and branches often results in multiple vessels or bifurcations in images. Explicit human guidance is essential for online robotic adjustments to better visualize target vessels. This guidance also aids segmentation networks as region proposals, improving vessel segmentation, particularly for small limb vessels (2.70 ± 0.15 mm) [12], which are challenging to extract accurately from B-mode images.

In order to tackle the aforementioned challenges, this work proposes a gaze-guided RUSS. Gaze tracker, as an intuitive and efficient tool for human-machine interaction, has been widely applied to various medical robotic systems [13]–[15]. This study presents the first exploration of integrating gaze tracking into RUSSs, allowing the system to incorporate real-time human supervision to improve the intelligence and robustness of the RUSS during complex scanning tasks. Especially in intra-operative ultrasound imaging, the hands of surgeons are often occupied with essential surgical tools such as needles or catheters, making manual maneuvering of the ultrasound probe challenging. Implementing a gaze tracker as a part of the RUSS helps the operators control the ultrasound scanning without disrupting the surgical workflow. On one hand, the eye tracking information can indicate the vessel of interest during the scanning. On the other hand, the gaze of the operator can also be applied as network attention to assist the vessel segmentation, particularly for challenging ones with limited size. Based on the segmentation results, the robot is controlled to adapt its maneuvers so that the vessel of interest is kept in the middle of the ultrasound image. To further improve the ultrasound image quality during scanning, a confidence-driven control for linear ultrasound probes is proposed. The main contributions of this article are listed as follows:¹²

- A gaze-guided segmentation network is proposed, which effectively integrates attention heatmaps derived from gaze patterns into segmentation. This improves segmentation performance by enabling the system to focus on the relevant regions in ultrasound images and providing the ability to segment vessels based on the operator’s gaze.

¹Codes: https://github.com/yuan-12138/Gaze_RUSS

²Video: <https://www.youtube.com/watch?v=Zul2-fqy0FE>

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- A human intention estimation module is introduced to predict the operator’s intention during scanning by combining gaze patterns with segmentation result history. This module stabilizes noisy gaze signals, providing consistent and reliable guidance to the segmentation network.
- A gaze-guided RUSS is developed and validated on a realistic arm phantom, showcasing its ability to perform complex scans. The system dynamically switched focus between vessels based on the operator’s eye tracking, while the confidence-driven control method continuously adjusted probe orientation to optimize image quality.

To emphasize realism, we highlight that the proposed gaze-guided RUSS does not rely on highly stable gaze input from clinicians. The system can still operate with the intermittent absence of gaze input by using a zero-value map as the gaze attention heatmap (see Sec. III-A). When necessary it can seamlessly integrate the operator’s input, such as when adjusting the scanning focus for complex scenarios like bifurcations.

II. RELATED WORK

A. Robotic Ultrasound Systems

Ultrasound image quality depends heavily on acquisition parameters such as force, orientation, and surface contact. To address this, Pierrot *et al.* [16] implemented a force control schema in a teleoperative RUSS to maintain constant scanning force. Jiang *et al.* [17] introduced an orientation optimization algorithm based on contact force for better visualization. To ensure that the ultrasound probe is in good contact with the surface, Chatelain *et al.* [18] integrated ultrasound confidence map [19] into the control loop and presented confidence-driven control. However, the proposed control algorithm was designed for convex probe. In the sight of such limitation, Jiang *et al.* [20] proposed an orientation correction method based on the confidence map for the linear probes. The experimental results clearly demonstrated its effectiveness, but the theoretical justification is missing.

Recent advancements in RUSS for vascular applications have gained significant attention. Jiang *et al.* [10] proposed an autonomous scanning framework for peripheral vascular diseases using real-time vessel segmentation. Bi *et al.* [21] applied reinforcement learning to autonomously navigate the probe to the longitudinal view of the carotid artery. Huang *et al.* [11] developed a system to autonomously perform carotid scans by imitating clinical protocols, while Goel *et al.* [22] introduced a Bayesian Optimization-based path planning framework for femoral artery screening.

Most existing RUSSs for vascular scanning do not address the challenge of handling multiple vessels and often require additional tools to capture the doctor’s intention for automatic maneuvering. Guidance can be implemented teleoperatively, with an expert operating remotely while the robot on-site follows [23], or through virtual reality, providing virtual guidance to the human operator [24]. Another approach involves pre-planned scanning paths, where the robot provides virtual fixtures to enable reproducible ultrasound scanning for follow-up validations [25]. However, among these strategies, gaze signal—an intuitive and hands-free interaction method—remains

unexplored. This approach is particularly advantageous in intra-operative scenarios, as it allows surgeons to perform tasks without altering workflow.

B. Gaze-guided Medical Robotic Systems

Initial attempts to integrate gaze information into medical robotic systems have been made in various scenarios, especially for laparoscopic surgeries. Noonan *et al.* [26] used eye tracking to control an articulated robotic laparoscope for stable, hands-free visualization. Fujii *et al.* [27] enhanced the gaze-guided laparoscope maneuverability by using a robotic arm instead of a rigid fixture. Clancy *et al.* [28] combined gaze tracking with a liquid lens in the da Vinci system for automatic focus adjustment during minimally invasive surgeries. Gaze tracking has also been applied to constrain laparoscopic surgical tools for tissue safety [29] and to improve collaboration in multi-robot surgeries by visualizing fixation points [30]. Li *et al.* [31] further optimized gaze tracking accuracy by compensating for head movements. Beyond laparoscopic applications, Guo *et al.* [15] utilized gaze to control needle insertion in CT-guided interventions, while Kogkas *et al.* [32] introduced a gaze-guided robotic scrub nurse to deliver surgical tools. To the best of our knowledge, initial attempt to integrate the gaze tracker into RUSS has not yet occurred.

C. Gaze-guided Medical Image Analysis

Unlike optical images, medical images are often challenging to interpret, requiring solid biological knowledge for accurate diagnostics. To enhance the robustness of medical image analysis networks, human experts’ gaze signals are frequently integrated as guidance. Cai *et al.* [33] proposed Temporal SonoEyeNet, which predicts sonographers’ visual attention and ultrasound standard planes, demonstrating that understanding visual attention complements standard plane detection. Similarly, Wang *et al.* [34] used class activation maps [35] to align network attention with human gaze. In another application, Men *et al.* [36] leveraged gaze tracking to improve probe movement prediction for obstetric standard plane navigation, while Alsharid *et al.* [37] showed its benefits in ultrasound video captioning. Most of these works aim to align networks with human attention to enhance reasoning. Beyond implicit guidance, gaze signals also serve as an intuitive tool for human-machine interaction. For instance, Khosravan *et al.* [38] developed a collaborative lesion segmentation system that analyzes and filters doctors’ gaze data to improve segmentation and reduce false positives.

III. METHODS

As shown in Fig. 1, the proposed gaze-guided RUSS consists of two main components: image processing and robotic control. The human operator’s gaze signal is fed into an attention-estimation module, which uses sequential segmentation results and gaze tracking history to predict the operator’s actual attention. It is important to distinguish between the gaze heatmap, directly captured by the eye tracker, and the attention heatmap, which reflects the operator’s true

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 1. Overview of the Proposed Gaze-Guided Interactive RUSS: The human gaze signal, captured by a gaze tracker, is combined with segmentation history in an intention estimation module to infer the operator’s preference, especially when multiple vessels are visible. The resulting attention heatmap guides the gaze-guided segmentation network to produce accurate vessel segmentation masks. These results are integrated into the robotic control loop to keep the target vessel centered in the ultrasound image. A confidence-based orientation correction optimizes probe contact with curved surfaces, improving image quality. Red boxes in ultrasound images highlight shadowed areas caused by improper probe contact.

focus. Sonographers often glance at other areas of the ultrasound image to maintain contextual awareness of surrounding anatomical structures, rather than focusing solely on the vessel of interest. Differentiating these incidental glances from an intentional focus shift is critical to ensure the RUSS follows the correct vessel and avoids unnecessary target switches. The estimated attention heatmap guides the gaze-guided segmentation network to extract the vessel of interest from the ultrasound image. The resulting segmentation directs the robot to center the vessel in the ultrasound image. To maintain good probe contact with the scanning surface, a confidence-driven orientation correction is implemented, adjusting the probe’s rotation around its y-axis based on the ultrasound confidence map to enhance image quality. During acquisition, impedance control [10] is employed to maintain consistent contact force.

A. Gaze-Guided Segmentation Network

During the training of the proposed segmentation network, pseudo attention maps are generated from the label maps of B-mode ultrasound images and paired with ultrasound images to simulate human attention on vessels. This approach addresses the instability and noise in real gaze signals, which could mislead the network, and the difficulty of directly obtaining human attention heatmaps. By incorporating randomness into their generation, these pseudo attention maps create a more diverse and extensive dataset, reducing bias and improving the network’s robustness and generalization. The pseudo attention heatmaps are generated based on the ground truth label maps. The centroid point of the heatmap (X_m, Y_m) is determined by the center point of the vessel label (X_l, Y_l) .

$$(X_m, Y_m) \sim \mathcal{N}((X_l, Y_l), \Sigma_c) \quad (1)$$

$$(x_m, y_m) \sim \mathcal{N}((X_m, Y_m), \Sigma_m) \quad (2)$$

where Σ_c represents the covariance matrix of the Gaussian distribution. Based on the sampled centroid point of heatmap

(X_m, Y_m) , N points (x_m, y_m) are selected and set to one. The image that contains all the sampled points (x_m, y_m) is convolved by a 30×30 kernel with all weights setting to 1. After the convolution, the resulting image is normalised to $[0, 1]$. In this manner, a diffused heatmap H is generated, simulating the density of the eye tracking points. Here $\Sigma_c \in R^{2 \times 2}$ and $\Sigma_m \in R^{2 \times 2}$ are diagonal matrices, with their diagonal values setting to 15 and 25, respectively.

As depicted in Fig. 2 (a), the segmentation network follows a U-shaped design. To enhance its performance, we incorporate attention heatmaps generated by the Human Intention Estimation Module based on human eye gaze signals. These maps represent the stabilized focus of the operator and are used to dynamically guide the network to prioritize the vessel of interest. A transformer attention block [39] is incorporated to combine external attention signals from heatmaps with features extracted from ultrasound images. This adaptation enables the network to dynamically align its focus with the operator’s intentions. As depicted in Fig. 2 (b), the transformer attention block uses attention heatmaps as guidance to emphasize regions of interest. Inspired by [40], subsampling is applied within the transformer block to reduce computational complexity. The transformed features are then concatenated with those from the ultrasound image encoder, facilitating the decoder’s ability to generate precise segmentation masks.

If an ultrasound image contains multiple vessels, they are individually labeled and treated as distinct training samples to generate pseudo attention heatmaps. To reduce reliance on explicit guidance, $n\%$ of the generated attention maps are replaced with uniform maps where all weights are set to zero. This means that for $n\%$ of the training samples, no attention is provided, prompting the network to segment all vessels in the image independently. For this study, n is set to 10%. This approach enhances the network’s adaptability, allowing gaze tracking to act as a performance booster rather than a critical




Fig. 2. (a) The overall design of the proposed gaze-guided segmentation network. (b) The structure of the transformer attention block. (c) The structure of the residual block.

component of the segmentation framework. During inference, pseudo attention heatmaps are replaced by the output of the attention estimation module (see Sec. III-B).

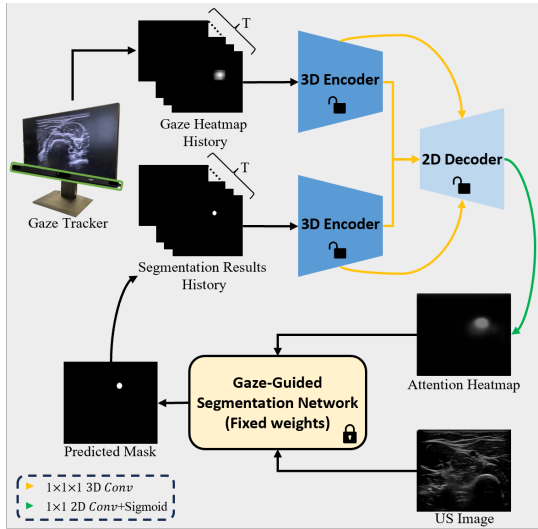


Fig. 3. The design of the human intention estimation module.

B. Stabilized Attention Heatmap Generation

Gaze heatmap is different from the human intention heatmap. Although the attention of the operator remains focused on the target vessel, it is unrealistic for them to keep steering at it throughout the whole scanning process. Instead, the sonographers tend to take glances at the surrounding anatomical structures to gain an overall contextual awareness. Therefore, without additional processing, the raw

gaze heatmaps are not suitable to be directly used as attentions in the proposed segmentation network. In essence, the desired heatmap of the segmentation network should reveal the preference of the human operator among multiple vessels in the ultrasound image. To achieve this, the Human Intention Estimation Module analyzes sequential gaze patterns and segmentation label maps from previous timesteps, as shown in Fig. 3. By combining these inputs, the module addresses the limitations of relying on gaze or segmentation history alone. For instance, if the segmentation history consistently focuses on a single vessel and the gaze heatmap briefly deviates, it is likely a temporal distraction. Conversely, stable deviation in the gaze heatmap suggests an intentional attention switch. This integration allows the network to reliably distinguish between temporal distractions and intentional changes in focus, ensuring robust guidance for segmentation.

As shown in Fig. 3, in our implementation, $T=64$ (approximately 2.13 seconds) was chosen empirically. This value aligns well with the average human fixation duration (290 ± 155 ms) [41] and avoids the time delay associated with higher T values. Two 3D convolutional encoders are employed to extract features from the gaze signals and segmentation results, respectively. At each scale, the extracted 3D feature maps are transformed into 2D feature maps through $1 \times 1 \times 1$ 3D convolution. The transformed 2D features are then concatenated with upsampled features in decoder to generate the final attention heatmap for the segmentation network.

For training, different gaze patterns are recorded from human operators serving as training data. During training, the well-trained weights of the segmentation network is frozen and the output attention map from the attention estimation module is fed into the segmentation network (see Fig. 2) to perform segmentation task. The resulting DICE loss between

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

the predicted segmentation map and the ground truth label was backpropagated to update the weights of the attention estimation module. By utilizing the fixed segmentation network, the DICE loss guided the attention estimation module to output refined attention maps that align with the segmentation task, effectively training the module to infer human attention.

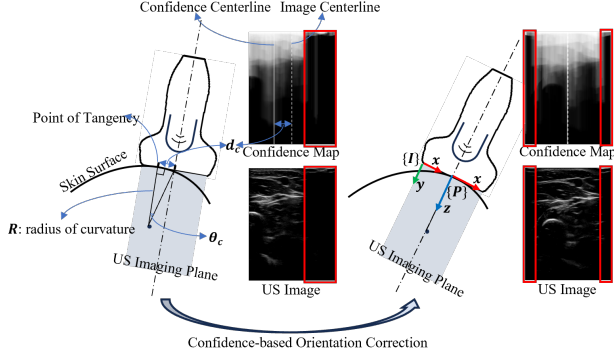


Fig. 4. The illustration of confidence-driven probe orientation correction for linear ultrasound probe. The red boxes in the images indicate the shadow areas caused by improper contact. $\{I\}$ represents the ultrasound imaging coordinate system, while $\{P\}$ is the probe coordinate system.

C. Confidence-based Orientation Correction for Linear Probe

The ultrasound probe's coordinate system is shown in Fig. 1. Translational movement along the z-axis is controlled using impedance control to ensure proper contact force with the scanning surface [10]. The y-axis represents forward motion during vessel scanning, while x-axis translation is guided by segmentation results to center the target vessel in the ultrasound image. In order to optimize the ultrasound image quality, a confidence-based orientation correction method for linear ultrasound probe is proposed.

The confidence map [19] is widely used for assessing ultrasound image quality, assigning per-pixel confidence values based on a simplified sound propagation model. While the original confidence map relies on a computationally intensive random walk algorithm, the scan-line-based confidence map [42] offers significantly improved efficiency. To ensure real-time performance in the proposed confidence-driven control for the linear ultrasound probe, the scan-line-based confidence map is utilized. Let $I(x, y)$ denote the intensity value of a ultrasound image at position (x, y) in the image coordination system. The scan-line-based confidence map is then given by:

$$C(X, Y) = 1 - \left(\int_0^{Y_{max}} f(I(X, y)) dy \right)^{-1} \int_0^Y f(I(X, y)) dy \quad (3)$$

where $f(*)$ is a non-linear mapping function between the radiofrequency signal intensity and the gray scale image intensity I . Here, $f(x) = x^2$ is used.

As shown in Fig. 4, dark regions in the confidence maps (highlighted by red rectangles) indicate improper contact between the ultrasound probe and the skin surface. To enhance image quality, a correction angle (θ_c) is calculated to rotate

the probe around its y-axis, improving the confidence level in the central region of the ultrasound image. The confidence-weighted centerline is determined as follows:

$$x_c = \frac{1}{C_\Omega} \iint_{(x,y) \in \Omega} x C(x, y) y dy dx \quad (4)$$

where $C_\Omega = \iint_{(x,y) \in \Omega} C(x, y) y dy dx$.

The position of the confidence-weighted centerline can be viewed as the tangency point between the linear transducer and the curved surface. Then as shown in Fig. 4, the distance (d_c) between the confidence-weighted centerline and the image centerline is computed.

$$d_c = x_c - X_c \quad (5)$$

where X_c is the position of the image centerline. Then, the correction angle θ_c can be computed using Eq (6):

$$\theta_c = \arctan\left(\frac{d_c}{R}\right) \quad (6)$$

where R is the radius of curvature in the ideal case, this would allow for a one-time adjustment. However, due to the varying curvature of the human body, accurately estimating this parameter online for human tissue is challenging, particularly considering the pressure-induced deformation. To address this, angular correction was performed using visual servoing mode. The parameter d_c ensures final convergence, aligning the confidence-based centerline with the centerline of the ultrasound imaging plane of the linear probe. A small R may cause fluctuations in movement. To balance motion smoothness with angular correction efficiency, R was empirically set to 10 cm based on experimental results.

IV. EXPERIMENTS AND RESULTS

A. Experimental Setup

To validate the proposed gaze-guided segmentation framework, ultrasound images of radial and interosseous arteries were collected using the Siemens Juniper ultrasound machine (ACUSON Juniper, SIEMENS AG) under Institutional Review Board approval from the Technical University of Munich (reference number 2022-87-S-KK). Ten ultrasound sweeps were obtained from five adult male volunteers, covering both arms, resulting in 2421 images. Image labeling was conducted under clinical supervision using ImFusionSuite (ImFusion GmbH). While the dataset size is relatively small, this reflects the practical challenges of collecting labeled data in medical imaging. The primary aim of this study is to validate how gaze signals can facilitate autonomous vessel segmentation, even with limited data availability, a common issue in medical applications. To train the attention estimation module, gaze data was collected from two volunteers instructed to focus on either the radial or interosseous artery during ultrasound sweeps. Each volunteer concentrated on a single vessel while the sweep video played, capturing their ground truth intention. To enhance robustness, volunteers were allowed to occasionally glanced at other image regions to reflect natural gaze variability. A total of 160 gaze recordings were obtained using the Tobii Eye Tracker 5, with 16 recordings per sweep. The networks were trained using the Adam optimizer with a fixed

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE I

SEGMENTATION RESULTS (DICE SCORE) OF DIFFERENT SEGMENTATION NETWORK ARCHITECTURES AND THE META DATA OF EACH VOLUNTEER.

Method	Artery	V_1	V_2	V_3	V_4	V_5
UNet [43]	Radial	0.456 ± 0.400	0.562 ± 0.294	0.084 ± 0.158	0.401 ± 0.301	0.292 ± 0.373
	Interosseous	0.359 ± 0.276	0.470 ± 0.297	0.1821 ± 0.256	0.467 ± 0.322	0.389 ± 0.323
Att-UNet [44]	Radial	0.486 ± 0.331	0.611 ± 0.283	0.254 ± 0.298	0.708 ± 0.181	0.385 ± 0.397
	Interosseous	0.436 ± 0.283	0.353 ± 0.224	0.515 ± 0.377	0.481 ± 0.340	0.498 ± 0.278
UTNet [40]	Radial	0.671 ± 0.255	0.710 ± 0.191	0.275 ± 0.249	0.748 ± 0.080	0.263 ± 0.329
	Interosseous	0.338 ± 0.207	0.345 ± 0.314	0.485 ± 0.288	0.367 ± 0.250	0.409 ± 0.302
UNet++ [45]	Radial	0.346 ± 0.332	0.707 ± 0.221	0.215 ± 0.223	0.739 ± 0.090	0.495 ± 0.259
	Interosseous	0.442 ± 0.251	0.421 ± 0.252	0.581 ± 0.316	0.242 ± 0.239	0.549 ± 0.261
Ours + Pseudo Att. Heatmap	Radial	0.794 ± 0.130	0.740 ± 0.199	0.389 ± 0.346	0.880 ± 0.053	0.709 ± 0.263
	Interosseous	0.603 ± 0.223	0.672 ± 0.185	0.741 ± 0.241	0.755 ± 0.165	0.746 ± 0.166
Ours + Raw Gaze Heatmap	Radial	0.500 ± 0.423	0.585 ± 0.351	0.259 ± 0.309	0.602 ± 0.384	0.404 ± 0.336
	Interosseous	0.212 ± 0.235	0.659 ± 0.167	0.563 ± 0.290	0.558 ± 0.327	0.428 ± 0.329
Ours + Est. Att. Heatmap	Radial	0.736 ± 0.295	0.721 ± 0.255	0.335 ± 0.322	0.822 ± 0.090	0.737 ± 0.145
	Interosseous	0.642 ± 0.096	0.684 ± 0.125	0.764 ± 0.105	0.741 ± 0.156	0.736 ± 0.106
Meta Data						
BMI		26.6	24.9	24.5	26.9	21.4
Average Size	Radial	3.0 mm	2.3 mm	1.7 mm	2.8 mm	2.6 mm
	Interosseous	3.3 mm	2.5 mm	2.8 mm	3.1 mm	2.8 mm

learning rate of 1×10^{-5} on a single GPU (Nvidia GeForce RTX 4070).

The proposed gaze-guided robotic ultrasound scanning system comprises a robot arm (KUKA LBR iiwa 7 R800, KUKA Roboter GmbH), a linear ultrasound probe (12L3, Siemens AG) attached to the robot via a 3D-printed probe holder, and a gaze tracker (Tobii Eye Tracker 5, Tobii AB) mounted below a display screen. The gaze tracker is calibrated using the manufacturer-provided software (Tobii Experience, Tobii AB), and real-time gaze signals are visualized as heatmap on the screen for user feedback. Data collection focused exclusively on eye-focusing heatmaps of ultrasound images, ensuring anonymity by avoiding the capture of participants' faces or eyes. Consent was obtained from all volunteers prior to their participation. The robot operates via a Robot Operating System (ROS) interface, and ultrasound images are accessed using a frame grabber (USB Capture HDMI Plus, Magewell Electronics). A commercial arm phantom (BPA304, Blue Phantom GmbH) is used for evaluation.

B. Results of Gaze-guided Segmentation Network

The ultrasound dataset described in Sec. IV-A is used to validate the segmentation performance of the proposed network against other segmentation networks. For individual volunteer performance, five models are trained for each network structure, with four volunteers used for training and the remaining one for testing. The models are indexed as V_1 to V_5 in Table I. Unlike other networks, which require separate weights for radial and interosseous arteries, the proposed network uses a single set of weights to segment both arteries. The target vessel is automatically determined based on gaze tracking results. During training, pseudo attention heatmaps are generated separately for the radial and interosseous artery labels. The Dice Score is used as the evaluation metric. This section evaluates the performance of the gaze-guided segmentation network using pseudo attention heatmaps, while Sec. IV-C examines performance with estimated attention heatmaps.

The segmentation results are shown in Table I. The proposed network outperforms all the SOTA segmentation networks. We can observe that without the guidance of gaze signal, the segmentation of interosseous arteries can hardly achieve a dice score of 0.55. Conversely, segmentation performance on radial arteries occasionally surpasses 0.70 under similar conditions. Such phenomenon is due to the fact that interosseous artery is located deeper in the arm, while radial artery is more superficial to the skin surface. Based on the propagation principle of ultrasound waves, the deeper part of the ultrasound image usually have more severe attenuation effect than the upper part of the image. Therefore, the visibility and clearness of interosseous artery is in general worse than the radial artery. An exception is the case of V_3 , for V_3 , the lower segmentation performance across all methods can be attributed to the extraordinarily small size of the radial artery (1.7 mm), which is a significant outlier compared to the average adult male radial artery size (2.68 ± 0.24 [12]). The difficulty of segmenting such an extreme outlier highlights the limitations of the current dataset. In contrast, the significant improvement in V_5 's segmentation performance with gaze guidance is due to the thinner subcutaneous tissue layer resulting from a lower BMI as shown in Table I. The proposed gaze-guided method successfully leveraged the gaze signal to focus on the target region, despite structural differences in the data. Besides, for the segmentation of radial artery, in some cases, UTNet and UNet++ perform almost equally well as our proposed method, especially in the case of V_2 , where UTNet achieved slightly worse mean value with smaller variance. This is because the radial artery is generally easy to identify from the ultrasound images. Moreover, we need to be aware that UTNet is trained to specifically focus on one artery at a time. Whereas the proposed gaze-guided segmentation network is trained to segment both vessels individually at a time. In contrast, for the segmentation of interosseous artery, the proposed gaze guided network clearly outperforms all the other network structures. The explicit guidance from human operator, in this case, can not only help the network to determine the segmentation target but also provide vital guidance to the network to locate

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

challenging structures.

C. Results with Attention Estimation Module

In this section, we combine the attention estimation module with the segmentation module to validate their joint performance. To train the attention estimation module, 160 gaze recordings were captured (16 per ultrasound sweep). For half of the recordings, volunteers were instructed to focus on the radial artery, while for the other half, they were asked to focus on the interosseous artery. For each gaze recording, volunteers were instructed to focus on only one of the two vessels for the entire duration of the ultrasound video. Of these, 120 recordings were used for training and 40 (four per sweep, two per vessel) for testing. To prevent the network from overly relying on perfect segmentation results, imperfect segmentation masks were generated using unprocessed gaze heatmaps. Training samples included a mix of perfect labels, imperfect segmentation masks, and blank images. This approach avoids the estimation module relying solely on segmentation history and ensures proper utilization of gaze heatmaps. As noted in Sec. III-B, during training, the gaze-guided segmentation network weights are fixed, and only the weights of the estimation module are updated.

The lower half of Table I presents the segmentation results of the gaze-guided network with and without the attention estimation module. Results without the estimation module were obtained by directly feeding raw gaze heatmaps into the network, following the same testing strategy as Sec. IV-B. The trained gaze-guided segmentation network from Sec. IV-B was used to train the corresponding estimation module, and the matching test gaze recordings were employed for evaluation. The absence of the estimation module significantly hinders segmentation performance, largely due to the noise and instability of raw gaze signals, compounded by gaze tracker calibration errors. The segmentation results with the estimation module closely match those achieved using pseudo attention heatmaps. In most cases, the pseudo attention heatmap delivers the best performance. However, in some instances, the estimated attention heatmap outperforms the pseudo attention heatmap, with differences within 0.04, likely due to data imbalance. The execution time of segmentation and attention estimation modules is 36.1 ± 0.5 ms.

D. Gaze-Guided Robotic Ultrasound Scanning




Fig. 5. (a) The experimental setup of the gaze-guided RUSS. (b) Reconstruction results of one representative scanning.

As shown in Fig. 5(a), experiments were conducted on an arm phantom to validate the proposed gaze-guided RUSS. The start and end points of the scans were manually selected, with an average scanning length of 12.9 ± 0.7 cm. An eye tracker mounted below the screen captured the operator's gaze signals. The segmentation network was trained on phantom data, while the attention estimation module was adapted from the model trained in Sec. IV-C. Fig. 5(b) displays the ultrasound reconstruction of the two vessels in the arm phantom. The blue regions represent reconstruction results from the proposed network's segmentation masks, overlaid on the ground truth vessel model in green. Both vessels were segmented successfully as the operator's attention switched from vessel 1 to vessel 2, demonstrating smooth and effective vessel switching.

TABLE II
PERFORMANCES W/ AND W/O CONFIDENCE-BASED ORIENTATION CORRECTION (AVERAGED d_c).

	w/o Confidence-based Orientation Correction	w/ Confidence-based Orientation Correction
Sweep 1	1.42 ± 1.43 mm	7.21 ± 0.77 mm
Sweep 2	1.70 ± 2.03 mm	7.11 ± 0.92 mm
Sweep 3	1.92 ± 1.63 mm	8.95 ± 3.18 mm

Experiments were also conducted to validate the control performance of the confidence-based orientation correction. We conducted 3 sweeps without the confidence-based control and 3 sweeps with the confidence-based control activated. Table II shows the averaged d_c from Eq. 5, which calculates how much the confidence centerline is deviated from the image centerline. A high value of d_c indicates the existence of a big portion of shadow area in the image. The proposed confidence-based method can efficiently optimize the ultrasound confidence by adjusting the in-plane orientation of the ultrasound probe.

V. CONCLUSION AND DISCUSSIONS

In this work, a gaze-guided RUSS is proposed. In order to integrate the human gaze signal into the control loop of RUSS, we proposed a gaze-guided segmentation network. Since the human gaze signal is noisy and unstable, a human intention estimation module is proposed to reveal the essence human attention from the gaze signal. Experiments have been conducted on the ultrasound images of radial and interosseous arteries. The proposed network demonstrated superior segmentation performance compared to SOTA methods. We acknowledge the limited dataset size and plan to expand it with more diverse anatomical structures and patient data to enhance generalizability. Additionally, future work will address the generalization challenge by exploring novel network architecture designs [46]. These efforts aim to further validate and improve the system's adaptability. Finally, the whole system was tested on an arm phantom with robots. The robot is controlled to centralise the vessel of interest based on the operator's preference interpreted from the gaze signal. To make sure the robot can establish proper contact with the scanning surface, a confidence-based orientation adjustment method is proposed. The phantom test validated the proficiency of the

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

whole system. In the future, we would further validate the proposed system on real human tissue.

REFERENCES

- [1] J. Tan, J. Li, Y. Li, B. Li, Y. Leng, Y. Rong, and C. Fu, "Autonomous trajectory planning for ultrasound-guided real-time tracking of suspicious breast tumor targets," *IEEE Trans. Autom. Sci. Eng.*, 2023.
- [2] Q. Huang, J. Lan, and X. Li, "Robotic arm based automatic ultrasound scanning for three-dimensional imaging," *IEEE Trans. Ind. Inform.*, vol. 15, no. 2, pp. 1173–1182, 2018.
- [3] J. Tan, B. Li, Y. Li, B. Li, X. Chen, J. Wu, B. Luo, and et al., "A flexible and fully autonomous breast ultrasound scanning system," *IEEE Trans. Autom. Sci. Eng.*, vol. 20, no. 3, pp. 1920–1933, 2022.
- [4] Z. Jiang, S. E. Salcudean, and et al., "Robotic ultrasound imaging: State-of-the-art and future perspectives," *Med. Image Anal.*, p. 102878, 2023.
- [5] M. Akbari, J. Carriere, S. Sloboda, T. Meyer, N. Usmani, S. Husain, and M. Tavakoli, "Robot-assisted breast ultrasound scanning using geometrical analysis of the seroma and image segmentation," in *2021 IROS*, 2021, pp. 3784–3791.
- [6] Y. Bi, Z. Jiang, F. Duellmer, D. Huang, and N. Navab, "Machine learning in robotic ultrasound imaging: Challenges and perspectives," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2024.
- [7] K. Li, Y. Xu, and M. Q.-H. Meng, "An overview of systems and techniques for autonomous robotic ultrasound acquisitions," *IEEE Trans. Med. Robot. Bionics*, vol. 3, no. 2, pp. 510–524, 2021.
- [8] F. von Haxthausen, S. Böttger, D. Wulff, J. Hagenah, and et al., "Medical robotics for ultrasound imaging: current systems and future trends," *Current Robotics Reports*, vol. 2, no. 1, pp. 55–71, 2021.
- [9] Q. Huang, J. Zhou, and Z. Li, "Review of robot-assisted medical ultrasound imaging systems: Technology and clinical applications," *Neurocomputing*, vol. 559, p. 126790, 2023.
- [10] Z. Jiang, Z. Li, M. Grimm, M. Zhou, M. Esposito, W. Wein, and et al., "Autonomous robotic screening of tubular structures based only on real-time ultrasound imaging feedback," *IEEE Trans. Ind. Electron.*, 2021.
- [11] Q. Huang, B. Gao, and M. Wang, "Robot-assisted autonomous ultrasound imaging for carotid artery," *IEEE Trans. Instrum. Meas.*, 2024.
- [12] W. Wahood, S. Ghozy, A. Al-Abdulghani, and D. F. Kallmes, "Radial artery diameter: a comprehensive systematic review of anatomy," *Journal of neurointerventional surgery*, vol. 14, no. 12, pp. 1274–1278, 2022.
- [13] D. P. Noonan, G. P. Mylonas, A. Darzi, and G.-Z. Yang, "Gaze contingent articulated robot control for robot assisted minimally invasive surgery," in *IROS*, 2008, pp. 1186–1191.
- [14] I. Tong, O. Mohareri, S. Tatasurya, C. Hennessey, and et al., "A retrofit eye gaze tracker for the da vinci and its integration in task execution using the da vinci research kit," in *IROS*, 2015, pp. 2043–2050.
- [15] J. Guo, Y. Liu, Q. Qiu, J. Huang, and et al., "A novel robotic guidance system with eye-gaze tracking control for needle-based interventions," *IEEE Trans. Cogn. Dev. Syst.*, vol. 13, no. 1, pp. 179–188, 2019.
- [16] F. Pierrot, E. Dombre, E. Dégoulange, L. Urbain, P. Caron, S. Boudet, and et al., "Hippocrate: A safe robot arm for medical applications with force feedback," *Med. Image Anal.*, vol. 3, no. 3, pp. 285–300, 1999.
- [17] Z. Jiang, M. Grimm, M. Zhou, Y. Hu, J. Esteban, and N. Navab, "Automatic force-based probe positioning for precise robotic ultrasound acquisition," *IEEE Trans. Ind. Electron.*, vol. 68, no. 11, pp. 11200–11211, 2020.
- [18] P. Chatelain, A. Krupa, and N. Navab, "Confidence-driven control of an ultrasound probe," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1410–1424, 2017.
- [19] A. Karamalis, W. Wein, T. Klein, and N. Navab, "Ultrasound confidence maps using random walks," *Med. Image Anal.*, vol. 16, no. 6, pp. 1101–1112, 2012.
- [20] Z. Jiang, N. Danis, Y. Bi, M. Zhou, M. Kroenke, T. Wendler, and N. Navab, "Precise repositioning of robotic ultrasound: Improving registration-based motion compensation using ultrasound confidence optimization," *IEEE Trans. Instrum. Meas.*, vol. 71, pp. 1–11, 2022.
- [21] Y. Bi, Z. Jiang, Y. Gao, T. Wendler, and et al., "Vesnet-rl: Simulation-based reinforcement learning for real-world us probe navigation," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 6638–6645, 2022.
- [22] R. Goel, F. Abhimanyu, K. Patel, J. Galeotti, and H. Choset, "Autonomous ultrasound scanning using bayesian optimization and hybrid force control," in *ICRA*, 2022, pp. 8396–8402.
- [23] Y. Fu, W. Lin, X. Yu, J. J. Rodríguez-Andina, and H. Gao, "Robot-assisted teleoperation ultrasound system based on fusion of augmented reality and predictive force," *IEEE Trans. Ind. Electron.*, vol. 70, no. 7, pp. 7449–7456, 2022.
- [24] D. Black, Y. Oloumi Yazdi, A. H. Hadi Hosseinabadi, and S. Salcudean, "Human teleoperation-a haptically enabled mixed reality system for teleultrasound," *Human-Computer Interaction*, vol. 39, no. 5-6, pp. 529–552, 2024.
- [25] D. Huang, C. Yang, M. Zhou, A. Karlas, N. Navab, and Z. Jiang, "Robot-assisted deep venous thrombosis ultrasound examination using virtual fixture," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 381–392, 2024.
- [26] D. P. Noonan, G. P. Mylonas, J. Shang, C. J. Payne, A. Darzi, and G.-Z. Yang, "Gaze contingent control for an articulated mechatronic laparoscope," in *2010 3rd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechanics*, 2010, pp. 759–764.
- [27] K. Fujii, G. Gras, A. Salerno, and G.-Z. Yang, "Gaze gesture based human robot interaction for laparoscopic surgery," *Med. Image Anal.*, vol. 44, pp. 196–214, 2018.
- [28] N. T. Clancy, G. P. Mylonas, G.-Z. Yang, and D. S. Elson, "Gaze-contingent autofocus system for robotic-assisted minimally invasive surgery," in *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2011, pp. 5396–5399.
- [29] G. P. Mylonas, K.-W. Kwok, D. R. James, D. Leff, F. Orihuela-Espina, A. Darzi, and G.-Z. Yang, "Gaze-contingent motor channelling, haptic constraints and associated cognitive demand for robotic mis," *Med. Image Anal.*, vol. 16, no. 3, pp. 612–631, 2012.
- [30] K.-W. Kwok, L.-W. Sun, G. P. Mylonas, D. R. James, F. Orihuela-Espina, and G.-Z. Yang, "Collaborative gaze channelling for improved cooperation during robotic assisted surgery," *Annals of biomedical engineering*, vol. 40, pp. 2156–2167, 2012.
- [31] Z. Li, I. Tong, L. Metcalf, C. Hennessey, and S. E. Salcudean, "Free head movement eye gaze contingent ultrasound interfaces for the da vinci surgical system," *IEEE Robot. Autom. Lett.*, vol. 3, no. 3, pp. 2137–2143, 2018.
- [32] A. Kogkas, A. Ezzat, R. Thakkar, A. Darzi, and G. Mylonas, "Free-view, 3d gaze-guided robotic scrub nurse," in *MICCAI*, 2019, pp. 164–172.
- [33] Y. Cai, R. Droste, H. Sharma, P. Chatelain, L. Drukker, A. T. Papageorghiou, and J. A. Noble, "Spatio-temporal visual attention modelling of standard biometry plane-finding navigation," *Med. Image Anal.*, vol. 65, p. 101762, 2020.
- [34] S. Wang, X. Ouyang, T. Liu, Q. Wang, and D. Shen, "Follow my eye: Using gaze to supervise computer-aided diagnosis," *IEEE Transactions on Medical Imaging*, vol. 41, no. 7, pp. 1688–1698, 2022.
- [35] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proceedings of CVPR*, 2016, pp. 2921–2929.
- [36] Q. Men, C. Teng, L. Drukker, A. T. Papageorghiou, and J. A. Noble, "Gaze-probe joint guidance with multi-task learning in obstetric ultrasound scanning," *Med. Image Anal.*, vol. 90, p. 102981, 2023.
- [37] M. Alsharid, Y. Cai, H. Sharma, L. Drukker, A. T. Papageorghiou, and J. A. Noble, "Gaze-assisted automatic captioning of fetal ultrasound videos using three-way multi-modal deep neural networks," *Med. Image Anal.*, vol. 82, p. 102630, 2022.
- [38] N. Khosravan, H. Celik, B. Turkbey, E. C. Jones, B. Wood, and U. Bagci, "A collaborative computer aided diagnosis (c-cad) system with eye-tracking, sparse attentional model, and deep learning," *Med. Image Anal.*, vol. 51, pp. 101–115, 2019.
- [39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [40] Y. Gao and et al., "Utmet: a hybrid transformer architecture for medical image segmentation," in *MICCAI*, 2021, pp. 61–71.
- [41] J. M. Henderson and W. Choi, "Neural correlates of fixation duration during real-world scene viewing: evidence from fixation-related (fire) fmri," *Journal of cognitive neuroscience*, vol. 27, no. 6, pp. 1137–1145, 2015.
- [42] P. Chatelain, "Quality-driven control of a robotized ultrasound probe," Ph.D. dissertation, Technische Universität München, 2016.
- [43] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*, 2015, pp. 234–241.
- [44] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, and et al., "Attention gated networks: Learning to leverage salient regions in medical images," *Med. Image Anal.*, vol. 53, pp. 197–207, 2019.
- [45] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: Redesigning skip connections to exploit multiscale features in image segmentation," *IEEE Trans. Med. Imaging*, vol. 39, no. 6, pp. 1856–1867, 2019.
- [46] Y. Bi, Z. Jiang, R. Clarenbach, R. Ghotbi, A. Karlas, and N. Navab, "Mi-segnet: Mutual information-based us segmentation for unseen domain generalization," in *MICCAI*. Springer, 2023, pp. 130–140.