

Received 24 June 2026, accepted 9 July 2026, date of publication 16 July 2026, date of current version 22 July 2026.

Digital Object Identifier 10.1109/ACCESS.2026.3713969

RESEARCH ARTICLE

An AI-Based Mobile Robot for Autonomous Ultrasound in Legal Medicine

SARAH GRUBE¹, SARAH LATUS¹, DUSTIN MOEBIUS², MARTIN FISCHER¹,
VIDAS RAUDONIS^{3,4}, AXEL HEINEMANN², BENJAMIN ONDRUSCHKA²,
AND ALEXANDER SCHLAEFER^{1,3}, (Senior Member, IEEE)

¹Institute of Medical Technology and Intelligent Systems, Hamburg University of Technology, 21073 Hamburg, Germany

²Institute of Legal Medicine, University Medical Center Hamburg-Eppendorf, 22529 Hamburg, Germany

³SustAInLivWork Center of Excellence, 51423 Kaunas, Lithuania

⁴Faculty of Electrical and Electronics Engineering, Kaunas University of Technology, 51367 Kaunas, Lithuania

Corresponding author: Sarah Grube (sarah.grube@tuhh.de)

This work was supported in part by the Interdisciplinary Competence Center for Interface Research (ICCIR) funded by Hamburg University of Technology (TUHH) and University Medical Center Hamburg-Eppendorf (UKE); in part by the European Union through Horizon Europe Program under Grant 101059903; and in part by the Research Consortium National Autopsy Network (NATON) through the National Network of University Medicine funded by the Federal Ministry of Research, Technology, and Space under Grant 01KX2524.

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Ethics Committee of Hamburg Chamber of Physicians under Application No. 2020-10353-BO-ff.

ABSTRACT Comprehensive documentation of internal tissue structures plays an important role in legal medicine. Although ultrasound is widely used in medical diagnostics as a powerful modality for injury detection and documentation, its integration into legal medicine workflows remains limited. Potential reasons include physical strain and infection risks due to manual probe placement at different locations along the body. Robotic ultrasound systems have been proposed to address such challenges. However, current systems are typically restricted to specific target regions and are limited in mobility and autonomy, complicating their clinical integration and fully autonomous, task-level operation. We therefore propose a mobile robotic system for autonomous ultrasound scanning across the body. The system localizes the body, estimates internal target structures, moves to reach the target structure, and performs ultrasound scans. We evaluate the system in three stages. First, we assess the accuracy and robustness of target localization and scan path planning across different body types. Second, we conduct experiments with a body phantom to evaluate the repeatability of covering target structures. Third, we demonstrate the system in a real legal medicine setting. Our experiments demonstrate that the system can approach the body, define the scan path, and acquire reproducible ultrasound images of target structures. These results demonstrate the feasibility of fully autonomous ultrasound imaging across multiple anatomical regions using a mobile robotic system. Beyond legal medicine, the proposed mobile system offers perspectives for autonomous imaging in broader clinical settings such as telemedicine and infectious disease settings, supporting remote diagnostics while reducing infection risks.

INDEX TERMS Autonomous robotic agents, autonomous systems, forensic imaging, legal medicine, medical imaging, medical robotics, robotic systems and software, robotic ultrasound, ultrasound imaging.

I. INTRODUCTION

Imaging is an important diagnostic tool in modern medicine, providing valuable information for different applications such as disease assessment, patient monitoring, and treatment

The associate editor coordinating the review of this manuscript and approving it for publication was Yingzhe Wang¹.

planning. While physical examination remains the initial step in clinical diagnostics, imaging is routinely employed to visualize internal structures and pathological changes that are not directly accessible from the outside.

Non-invasive imaging techniques have also gained interest in legal medicine. In this field, comprehensive post mortem documentation across multiple body regions is essential for

analyzing disease pathogenesis and determining the cause of death [1], [2]. Typically, the documentation relies on invasive autopsy. However, autopsy alone often lacks the comprehensive, internal visualization required to document deep-tissue trauma or pathological changes. Consequently, legal medicine is increasingly employing imaging modalities to achieve a more comprehensive documentation of the internal body structures. Currently, the combination of invasive autopsy and non-invasive Computer Tomography (CT) is the gold standard for comprehensive documentation. While autopsy provides direct tissue inspection, CT offers an imaging modality to visualize internal structures across the whole body. However, both autopsy and CT documentation are subject to limitations. Autopsies are frequently restricted or avoided in cases of High-Consequence Infectious Diseases (HCIDs), such as Ebola or SARS-CoV-2, to minimize infection risks for the clinicians [3], [4], [5]. CT, in turn, is limited by its high cost and the need for specialized infrastructure, which restricts its accessibility in many legal medicine institutes, particularly for on-site investigations and in less developed regions. These safety risks and limited CT accessibility result in incomplete post mortem documentation, leading to a lack of systematic internal data [6], which is, however, essential for understanding disease pathogenesis and developing public health responses. This challenge was recently underscored by the SARS-CoV-2 pandemic.

To bridge this gap, US imaging provides a portable, cost-effective, and widely available alternative for internal documentation. However, US imaging is still not widely used in legal medicine. This may be related to specific challenges, e.g., time-consuming applications, operator dependency, and physical strain [7], [8], [9]. Beyond ergonomics, the requirement for direct physical contact poses substantial infection risks. These risks were clearly evident during the SARS-CoV-2 pandemic. While lung US is an established modality for detecting and monitoring COVID-19 manifestations, it requires close-range contact with the examined subject, introducing significant infection risks for sonographers [10], [11], [12].

To address these limitations of US, there has been renewed interest in robotic US scanning systems (RUSS) [13], [14], [15]. These systems are well-suited for repetitive, high-precision tasks and promise to offer consistent image acquisition, improved reproducibility, and the capability for volumetric imaging of large areas. Moreover, by decoupling the sonographer from the US probe, RUSS can reduce the physical strain on the sonographer and enable safe remote examinations in environments with a high risk of infection.

Current research primarily follows two distinct approaches. On the one hand, robotic teleoperation systems, such as the HERCULES platform [16], enable remote examinations via a haptic interface, decoupling the sonographer from the subject's location [13], [17]. Although teleoperation is effective in reducing infection risks and physical strain, it remains an

assistive approach that requires trained sonographers to be continuously available. Hence, teleoperation does not reduce the significant temporal workload of the scanning process. This creates a bottleneck in settings where sonographers availability is already a primary constraint.

On the other hand, autonomous RUSS aim to reduce this workload by automating the scanning process itself. These systems typically employ small collaborative robotic arms in stationary setups, where both the robot and the body remain in pre-aligned, fixed positions [13]. Unlike teleoperation, no continuous human supervision is required, thereby enabling workload redistribution. Achieving this level of autonomy requires the independent execution of target localization, scan path planning, and probe manipulation. While approaches for estimating the location of the target structure based on the body surface have been proposed [18], [19], [20], the scan regions are still commonly defined manually by a clinician [21] or derived from landmark detections [22], [23], [24], [25].

The main drawback of existing RUSS approaches is their limited flexibility, as they are usually realized as a stationary setup. This constraints their clinical applicability due to the limited reach of a small robotic arm. However, a comprehensive examination in legal medicine involves anatomical targets across the whole body, which are difficult to image from a single fixed position. Instead, a sonographer typically moves around the couch to reach feasible US windows for each target.

We propose mimicking this behavior by an intelligent and mobile robotic agent. Hence, we integrate a mobile base into the RUSS, providing additional degrees of freedom (DoF) that extend the operational workspace to the entire examination room. This transition from a stationary tool to a mobile autonomous agent allows the system to reposition itself autonomously relative to the target structure. By dynamically adjusting its base position, the agent maintains the maneuverability required to access optimal US windows across the entire body necessary for a comprehensive legal medicine examination. Such an autonomous mobile agent has the potential to address the challenges of legal medicine. It allows for comprehensive, full body US scans without human intervention, thereby reducing the physical and temporal workload for clinicians [9]. Furthermore, by operating independently, the mobile agent provides a contact-free solution that reduces the infection risks associated with manual post mortem procedures.

In this work, we describe the proposed MOBILE Robotic UltraSound System (MORUSS), an autonomous agent designed for comprehensive US imaging in unknown medical environments. The MORUSS is composed of a small 6 DoF robotic arm mounted on a mobile base, similar to [26], and utilizes a depth camera to perceive the surrounding environment. Unlike stationary systems, the MORUSS utilizes its mobile base to autonomously approach and image different target structures. To illustrate requirements for such a robotic

agent, consider a scenario where the MORUSS is placed at the entrance of a room and is tasked with providing US images of the liver. Fulfilling this task requires the agent to solve three spatial challenges: (1) exploring the environment to identify and localize the body within the room, (2) inferring the position of the liver based on the external body surface, (3) repositioning its base to ensure target reachability and executing the autonomous US scan.

We evaluate our prototypical MORUSS in a simulation, as well as in our laboratory and in a real legal medicine setting. We first evaluate the impact of base mobility on target reachability by benchmarking the mobile platform against a stationary robotic setup. Second, we analyze the robustness of our target localization and base positioning algorithm through a systematic simulation study with full body CT scans. Third, we evaluate the system's performance in a laboratory using a body phantom. We obtain quantitative results regarding the repeatability and accuracy of target-specific US scans. Finally, we demonstrate the system in a legal medicine setting and illustrate its capability to autonomously acquire US images of different organs within a real human body. A clinician validates the visibility of the target structure in the acquired US images by manual segmentation.

II. METHODS

The MORUSS framework is designed to autonomously acquire US images of a selected target structure from within the body. It operates without prior knowledge of the subject-specific anatomy and is robust to inter-subject variability. Central to this approach is the ability to localize the position of the target structure, to autonomously position the mobile base, and to image the target structure. The subsequent sections detail these steps. Section II-A outlines the experimental setup and general workflow, while Section II-B describes the methodology for positioning the mobile base relative to an unknown target. Section II-C describes the algorithmic implementation required for workflow execution. Finally, Section II-D describes the hardware configuration.

A. SYSTEM OVERVIEW AND WORKFLOW

The proposed MORUSS is designed to function as an autonomous agent in legal medicine environments. We consider a scenario in which the body lies on a couch in a room, and the MORUSS is positioned at an arbitrary initial location within the room. The clinician only names the target structure. Thereafter, the MORUSS autonomously approaches the body and acquires images. To enable autonomous US scanning of the target structure, the MORUSS incorporates a mobile base (RB-Theron, Robotnik), a small lightweight robotic arm (UR3, Universal Robots) equipped with a US probe, and a color and depth (RGB-D) camera (Azure Kinect, Microsoft) for navigation, as shown in Fig. 1. A US imaging system (Cicada, Cephasonics) with a linear US probe (CPLA12810A, Cephasonics) and a curved US probe (CPC12835R5P, Cephasonics) are integrated. The RGB-D camera is mounted on the fourth joint of




FIGURE 1. An overview of the MORUSS. (a) Overall setup containing a mobile base and a robotic arm with an RGB-D camera and a US probe attached. The coordinate frames and the corresponding transformations are depicted as well. (b) The proposed autonomous US scanning workflow contains three main steps, i.e., exploration of the room, localization of the target structure, and US scanning of the target structure.

the UR3. The US probe is attached to the robot's end-effector through a custom adapter with an integrated force sensor (SRI M3703A, Sunrise Instruments) to realize force-controlled US scanning.

In order to acquire US images of the target structure without prior knowledge of the target's spatial location, we propose a workflow that transitions from global room perception to target localization within the body and US scanning. The workflow is illustrated in Fig. 1b and consists of three main steps: (1) exploration of the room, (2) localization of the target structure and base repositioning, and (3) explorative US scanning of the target structure. The Supplementary Material provides visualizations of the workflow. During the initial exploration step (1), the MORUSS generates a global 3D map of the room and a 3D scan of the body using the RGB-D camera. This allows the system to identify the body's position within the room and in step (2) to plan a safe navigation path toward it. However, from this data, only the external body surface is known; the location of the internal target structure remains unknown. Therefore, the MORUSS estimates the target's location based on the acquired surface geometry using an anatomical atlas model (2). Additionally, it determines a target-specific base pose to ensure that the target structure is within the reachable imaging workspace of the robotic arm. Once the MORUSS navigated to the determined target-specific base pose, the MORUSS acquires a local RGB-D scan to refine its relative

TABLE 1. Comparison of RUSS approaches in the literature and the proposed MORUSS with respect to system mobility, target structures, experimental validation, initial ROI localization, and scan path planning. Experimental validation indicates the anatomical model or subject used for evaluation. Scan ROI localization refers to the strategy used to determine the body surface region for US acquisition, including manual selection, color-based segmentation, or anatomical landmark detection. Scan path planning summarizes whether the system employs predefined trajectories, a single scan path, or multiple scan paths covering the ROI. Online feedback indicates whether the scan path is adapted online during scanning based on sensor feedback, including force sensing or US image analysis (e.g., image quality assessment or anatomical structure segmentation).

Paper	Mobility	Target structure	Validation	Initial ROI localization	Scan path	Online feedback
Jiang et al. (2021) [27]	Fixed setup	Tubular structures	Vascular phantom	Manual	Single path	Force; US image
Bao et al. (2023) [28]	Fixed setup	Abdominal structure	Abdominal phantom	Manual	Single path	Force
Huang et al. (2024) [29]	Fixed setup	Carotid artery	Human carotid artery	Manual	Single path	Force; US image
Wang et al. (2024) [30]	Fixed setup	Carotid artery	Human carotid artery	Manual	Single path	Force; US image
Sun et al. (2025) [21]	Fixed setup	Musculoskeletal structures	Musculoskeletal phantom	Manual	Single path	Force
Zheng et al. (2025) [31]	Fixed setup	Abdominal structure	Abdominal phantom; human abdominal aorta and common iliac artery	Manual	Multiple paths	Force
Welleweerd et al. (2021) [32]	Fixed setup	Breast	Breast phantom	Manual	Multiple paths	US image
Suligoj et al. (2021) [33]	Fixed setup	Chest, neck and jugular vein	Torso phantom	Manual	Multiple paths	Force
Osburg et al. (2024) [34]	Fixed setup	Superficial femoral artery	Human superficial femoral artery	Manual	Single path	Force; US image
Huang et al. (2018, 2019) [35], [36]	Fixed setup	–	Thyroid, lumbar and breast phantoms; human forearm	Color-based segmentation	Single path	Force
Wang et al. (2023) [37]	Fixed setup	Female breast	Breast phantom	Color-based segmentation	Multiple paths	Force
Lin et al. (2025) [25]	Fixed setup	Thyroid	Thyroid and tumor phantoms; human thyroid	Depth-based segmentation	Predefined path	Force; US image
Mustafa et al. (2013) [22]	Fixed setup	Liver	Human liver	Anatomical landmark	Single path	Force; US image
Su et al. (2024) [23]	Fixed setup	Thyroid	Human thyroid	Anatomical landmark	Single path	Force; US image
Tang et al. (2024) [24]	Fixed setup	Heart	Upper-body and heart phantoms	Anatomical landmark	Predefined path	Force; US image
Tan et al. (2022) [38]	Fixed setup	Breast	Breast phantom	Anatomical landmark	Multiple paths	Force
Wu et al. (2025) [39]	Fixed setup	Kidney	Upper-body phantom	Anatomical landmark	Single path	None
Ma et al. (2021) [40]	Fixed setup	Lung	Upper torso phantom	Anatomical landmark	–	Force
Huang et al. (2021) [41]	Fixed setup	Carotid artery	Artery phantom; human carotid artery	Anatomical landmark	Single path	Force; US image
Jiang et al. (2022) [20]	Fixed setup	Radial artery	Human radial artery	Atlas-based target localization within the forearm	Single path	Force
MORUSS	Mobile base	Whole body	Whole body CTs; whole body phantom including liver, heart and leg vein; full corpses including left/right kidney and liver	Atlas-based target localization within the body	Multiple paths	Force

pose with respect to the body and to acquire detailed body surface geometry for US scan path planning (3). In the final US scanning step, the system plans and performs a US scan based on the local body surface geometry and the estimated location of the target structure. The MORUSS records the robot pose for each US image to reconstruct a US volume covering the target structure. The 3D surface scan and the US images can both be expressed with respect to a single world frame, using the spatial transformations derived in Section II-D. This provides a comprehensive 3D representation of both the external body surface and the internal target structures for subsequent clinical analysis.

Table 1 compares the proposed MORUSS to existing RUSS in terms of system mobility, target structures, experimental validation, region-of-interest (ROI) localization, and scan path planning. These systems typically use fixed robotic setups designed to image a specific anatomical target structure. Experimental validation is usually performed using partial anatomical phantoms or human subjects. Note that the latter often limits the acquisition of ground truth images. Furthermore, the table shows that the ROI, i.e., the area on the body surface that should be covered during US scanning, is typically defined manually or based on color features or anatomical landmarks, e.g., the umbilicus

or mammary papillae. While landmarks and body features are generally derived from RGB-D surface scans, some approaches estimate the skeleton in order to define the ROI. Different approaches for scanning the ROI have been considered. A majority of systems realize a single scan along the surface and some systems employ predefined paths or multiple paths covering the ROI by a more complex scan pattern. Most systems use online feedback during scanning to maintain sufficient contact between the probe and the tissue, ensuring adequate US image quality. These systems typically rely on force sensing or the analysis of US images (e.g., image quality assessment or anatomical structure segmentation). The table also illustrates that the MORUSS extends existing approaches in a number of ways, particularly with respect to the mobility of the base and the related need to identify the respective ROI for targets in different regions of the body, i.e., enabling autonomous navigation and US scanning across the whole body.

B. PRE-COMPUTED BASE POSITIONING

One challenge in autonomous US is the lack of prior spatial information regarding the internal anatomy. Since internal structures are not observable with the RGB-D camera, the MORUSS must infer the anatomical context from the subject's surface geometry. This localization of the target structure is not only essential for determining where to perform the US scan, but also for selecting a base pose such that the target structure lies within the imaging workspace of the robotic arm. Hence, to address these challenges, the MORUSS utilizes a standardized anatomical atlas model [42] as a subject-agnostic spatial reference.

1) ANATOMICAL ATLAS MODEL

The atlas model provides the average 3D position of different target structures within a standardized human body (Fig. 2). This allows a reasonably accurate initial guess of the target position within the actual body. The atlas is defined as

$$\text{atlas } \mathcal{A} = \{\text{atlas } \mathcal{O}_i \in \mathbb{R}^3 \mid i = 1, \dots, N\},$$

where $\text{atlas } \mathcal{O}_i$ denotes the point cloud representation of the i -th target structure in the atlas frame. Each target structure $\text{atlas } \mathcal{O}_i \subset \mathbb{R}^3$ is modeled as a rigid object.

2) TARGET-SPECIFIC BASE POSITIONING

The robotic arm of MORUSS has a compact workspace which requires base repositioning to ensure that the US windows for a target-structure $\mathcal{O}_i$ are within the kinematic reach of the robotic arm. Furthermore, US scanning requires kinematic maneuverability to allow for scanning along a continuous trajectory on the body surface. Consequently, to ensure sufficient target coverage, we select a base pose that enables target reachability and sufficient kinematic maneuverability to follow the body surface during the US scan, while accounting for real-world navigation inaccuracies.

However, evaluating the reachability and kinematic maneuverability for every potential base pose on the

operating floor is computationally expensive. We therefore augment the anatomical atlas model with an offline-computed look-up table that contains, for every target structure $\mathcal{O}_i$, a selected feasible base pose $\text{atlas } \mathbf{b}_i^* = [x_i^*, y_i^*, \theta_i^*]^T \in \mathbb{R}^3$ in the atlas frame, where x_i^* and y_i^* denote the planar position and θ_i^* the orientation. During the scanning workflow, the target-specific base pose $\text{atlas } \mathbf{b}_i^*$ is retrieved from the atlas and transformed into the subject coordinate frame.

In order to determine the look-up table, we perform a configuration space analysis (CSA) based on the atlas model. We discretize the potential base poses $\mathcal{B}$ on the operating floor at 50 mm resolution. For every $\mathbf{b} \in \mathcal{B}$ we evaluate each pose by solving a multi-criteria optimization problem using a lexicographic strategy. We first retain only those poses from which the target structure is sufficiently reachable. Among the remaining feasible poses, we subsequently optimize kinematic manipulability and spatial robustness to positioning inaccuracies.

Target Coverage $C_i(\mathbf{b})$: For each base pose $\mathbf{b} \in \mathcal{B}$ and target structure $\mathcal{O}_i$, we compute the target coverage $C_i(\mathbf{b})$, as the fraction of the target volume intersecting the US imaging workspace $W(\mathbf{b})$, where $W(\mathbf{b})$ denotes the union of US windows from all reachable body surface points. Note, that due to physical limitations of US imaging, full coverage may not be achievable for every target structure. To address this, we introduce a coverage threshold parameter $\gamma \in (0, 1]$ and define a subset of feasible base poses

$$\mathcal{B}_{\text{valid}} = \{\mathbf{b} \in \mathcal{B} \mid C_i(\mathbf{b}) \geq \gamma \cdot \max_{\mathbf{b} \in \mathcal{B}} C_i(\mathbf{b})\}.$$

In this work, we set $\gamma = 0.95$, which was found to provide a good trade-off between target coverage, robustness, and flexibility for additional optimization criteria. In general, the subset $\mathcal{B}_{\text{valid}}$ may form multiple disconnected regions in the configuration space, from which we only retain base poses from the largest connected component, denoted as $\tilde{\mathcal{B}}_{\text{valid}}$.

Manipulability $M(\mathbf{b})$ and Spatial Robustness $R_i(\mathbf{b})$: Within the feasible set $\tilde{\mathcal{B}}_{\text{valid}}$, we evaluate each pose with respect to manipulability and spatial robustness. To ensure the robotic arm can follow the body surface, we evaluate its kinematic manipulability using the mean manipulability index [43], averaged over all body surface points within the scanning area. Following [15], [44], we assume an orthogonal probe placement relative to the skin surface during the scanning process.

The spatial robustness score introduces a safety margin that compensates for positioning errors while maintaining sufficient target coverage. It quantifies how close a base pose $\mathbf{b}$ lies to the geometric center of the region formed by $\tilde{\mathcal{B}}_{\text{valid}}$. $R_i(\mathbf{b})$ is defined for all $\mathbf{b} \in \tilde{\mathcal{B}}_{\text{valid}}$ as the minimum Euclidean distance to the boundary of this region. Maximizing $R(\mathbf{b})$ yields poses located near the geometric center, providing sufficient reachability of the target structure despite translational positioning errors.

Finally, we select a feasible base pose $\text{atlas } \mathbf{b}_i^*$ by solving a weighted scalarization of the normalized manipulability and

robustness scores over the subset $\tilde{\mathcal{B}}_{\text{valid}}$:

$$\text{atlas } \mathbf{b}_i^* = \arg \max_{\mathbf{b} \in \tilde{\mathcal{B}}_{\text{valid}}} \left(w_m \hat{M}(\mathbf{b}) + w_r \hat{R}_i(\mathbf{b}) \right)$$

where w_m and w_r are weighting factors that balance the normalized manipulability and robustness scores, $\hat{M}(\mathbf{b})$ and $\hat{R}_i(\mathbf{b})$, respectively.

C. WORKFLOW EXECUTION

In this section, we detail the implementation of the proposed workflow, which enables autonomous US scanning of different subjects. The workflow utilizes the pre-defined atlas model as a reference model to estimate the spatial relationship between the external surface geometry of the body and the internal anatomical targets. This mapping provides the necessary spatial information to autonomously plan a US scan that covers the target structure.

1) EXPLORATION OF THE ROOM

The exploration step acquires a global 3D map of the room, defines a world frame, and localizes both the body and the MORUSS within the world frame. To this end, the MORUSS first performs an explorative scan of the room in its local base frame F_{base} , capturing RGB-D point clouds from different arm poses. This provides a local environment map ${}^{\text{base}}\mathcal{E}$ of the room relative to the MORUSS, where the left superscript indicates that the point cloud is represented in the MORUSS base frame F_{base} . The MORUSS then detects the couch in ${}^{\text{base}}\mathcal{E}$ as a ground-parallel plane [26] and defines the world frame at its center, with axes aligned to the principal directions of the couch. We describe the pose of the MORUSS in the world frame with ${}^{\text{world}}T_{\text{base}}$, where ${}^{\text{world}}T_{\text{base}}$ is the transformation from the MORUSS base frame to the world frame. All points in ${}^{\text{base}}\mathcal{E}$ above the couch are assigned to the body surface ${}^{\text{base}}\mathcal{S}_{\text{body}}$, represented in the base frame F_{base} . We define the body frame at the center of the detected body surface, aligning its axes with the anatomical directions: the x -axis along the transversal direction, the y -axis along the longitudinal direction, and the z -axis along the sagittal direction. This yields the pose of the body relative to the MORUSS as ${}^{\text{base}}T_{\text{body}}$. Based on these steps, the poses of the MORUSS and the body in the world frame are now known and represented by the transformation matrices ${}^{\text{world}}T_{\text{base}}$ and ${}^{\text{world}}T_{\text{body}}$, respectively.

2) LOCALIZATION OF THE TARGET STRUCTURE AND BASE POSITIONING

The MORUSS localizes the target structure in the world frame by registering the atlas model to the acquired body surface ${}^{\text{world}}\mathcal{S}_{\text{body}}$. To this end, the system first scales the atlas model to match the body size and then applies an ICP-based surface registration. This yields the transformation ${}^{\text{world}}T_{\text{atlas}} \in \mathbb{R}^{4 \times 4}$. Hence, the subject-specific locations of the target structures in the world frame can be estimated as

$${}^{\text{world}}\mathcal{O}_i = \{ {}^{\text{world}}T_{\text{atlas}} \mathbf{x} \mid \mathbf{x} \in {}^{\text{atlas}}\mathcal{O}_i \}.$$




FIGURE 2. Workflow for scan path planning. (a) The location of the target structure is estimated by registering the atlas model to the acquired body surface. (b) The ROI for planning the US scan path is selected. First, the set of body surface points is selected, from which the target structure can be reached with US, considering the imaging depth Δ_{depth} (cross-section A-A). Next, a bounding box (blue rectangle) is fitted to these surface points to define the ROI. (c) The scan path is generated such that the ROI is covered by N parallel scan trajectories with the scan direction d_{scan} , spaced by a distance Δ_{scan} . For each slicing plane, surface points within a tolerance $\pm \Delta_{\text{slice}}$ are selected, and a cubic B-spline is fitted to generate a scan path.

Next, the MORUSS retrieves the pre-computed base pose ${}^{\text{atlas}}\mathbf{b}_i^*$ for the target structure from the look-up table in the atlas model and transforms it into the world frame by applying

$${}^{\text{world}}\mathbf{b}_i^* = {}^{\text{world}}T_{\text{atlas}} {}^{\text{atlas}}\mathbf{b}_i^*.$$

The MORUSS then moves to the resulting pose ${}^{\text{world}}\mathbf{b}_i^*$.

3) US SCANNING OF THE TARGET STRUCTURE

In the final step the MORUSS plans a US scan path that enables imaging of the target structure, as visualized in Fig. 2. For this, the MORUSS requires the local surface geometry and the location of the target structure, both in the world frame. To this end, the MORUSS performs an RGB-D scan of the local body surface to refine its pose with respect to the body and to obtain the local surface geometry. This pose refinement also avoids the accumulation of localization errors caused by repeated movements of the mobile base. As the target's location is determined from the atlas model, the MORUSS must account for potential spatial uncertainties.

To address this, the MORUSS implements an explorative scanning strategy. The system plans a continuous scanning trajectory on the body surface to acquire a US volume expected to contain the target structure.

The corresponding region of interest (ROI) for US scan path planning is defined on the body surface as follows. The MORUSS determines the region of the body surface from which the target structure is visible with US, considering the imaging depth of the US probe Δ_{depth} . Next, it fits a bounding box aligned with the desired scan direction $\mathbf{d}_{\text{scan}}$ to this region. The ROI is given by the subset of surface points that lie within this bounding box. The scan direction can be chosen arbitrarily; in the following, we consider transversal scan paths.

To systematically cover the defined ROI, the MORUSS executes a continuous boustrophedon trajectory consisting of parallel scan paths aligned with the scan direction $\mathbf{d}_{\text{scan}}$. To this end, the ROI is partitioned into N parallel slices of thickness Δ_{scan} . For each slice, the subset of body surface points falling within a tolerance Δ_{slice} of the slice midplane is extracted. A cubic B-spline is fitted through these points to generate a continuous 3D scan path, which is discretized into US probe poses, with the probe being placed orthogonally to the body surface.

D. SYSTEM IMPLEMENTATION AND CALIBRATION

We use the Robot Operating System (ROS) [45] for the overall control of the system components, including MoveIt! [46], the Nav2 framework [47] and the *slam_toolbox* [48].

For robotic US scanning, we represent all system components within the world frame, see Fig. 1. We establish the transformations between all system components by calibration. We obtain the transformation ${}^{\text{arm}}T_{\text{flange}}$ from the flange of the robotic arm to the base of the robotic arm from the robotic forward kinematics. We determine the transformation ${}^{\text{arm}}T_{\text{camera}}$ from the RGB-D camera frame to the base frame of the robotic arm via hand-eye calibration using Park's method [49]. We estimate the transformation ${}^{\text{arm}}T_{\text{base}}$ from the mobile base to the UR3 arm via edge-based calibration using a room corner [50]. The RGB-D camera on the UR3 arm and the RB-Theron's built-in depth camera are used to acquire room-edge point clouds, which are then aligned to determine the transformation. We establish the transformation ${}^{\text{flange}}T_{\text{us}}$ from the US image frame to the flange of the robotic arm by Z-wire calibration [51]. We transform a pixel $\mathbf{p}_I$ from the US image matrix $\mathbf{I}$ into the base frame of the robotic arm using

$$\begin{bmatrix} {}^{\text{arm}}\mathbf{p}_{I(c;r)} \\ 1 \end{bmatrix} = {}^{\text{arm}}T_{\text{flange}} {}^{\text{flange}}T_{\text{us}} \left[c\Delta x - \frac{w_{\text{us}}}{2}, 0, r\Delta y, 1 \right]^T \quad (1)$$

where $c, r \in \mathbb{N}_0$ denote the c -th column and r -th row in the US image matrix, w_{us} is the probe width, and Δy and Δx represent the pixels height and width, respectively.

III. EXPERIMENTAL EVALUATION

We evaluate the proposed MORUSS workflow through a simulation-based validation, quantitative phantom studies conducted in a laboratory, and a real-world application in the legal medicine.

A. VALIDATION OF THE WORKFLOW

Since target localization, robot base positioning, and US scan path planning depend on the body anatomy, we evaluate their accuracy and robustness across a wide range of body shapes using systematic simulations of the workflow. These simulations assess how anatomical variations affect the localization accuracy, the robustness of the proposed robot base positioning algorithm and the target coverage achieved by the resulting US scan paths.

Dataset and Simulation Environment: For this purpose, we use 14 full-body CT datasets acquired at the Institute of Legal Medicine at the University Medical Center Hamburg-Eppendorf (Legal Medicine, UKE). The study is approved by the Ethics Committee of the Hamburg Chamber of Physicians and the study protocol includes the informed consent of relatives or legal representatives. We perform the analysis for exemplary target structures. These include abdominal organs (liver, right kidney, and left kidney) as well as musculoskeletal structures (upper leg extremities). While the abdominal target structures show organ-specific reachability, the musculoskeletal scenario is required for diagnostic tasks, such as assessing secondary vascular occlusion or deep muscle rupture following high-energy bone fractures [2]. To estimate the coverage in the upper leg extremity, we use the femur as a surrogate volumetric target, since relevant structures lie between the skin surface and the bone.

Each CT dataset serves as basis for simulating the workflow. We extract the body surface ${}^{\text{world}}S_{\text{CT, body}}$ and the target structures ${}^{\text{world}}\mathcal{O}_{\text{CT},i}$ from the CT data using TotalSegmentator [52]. Note that the initial exploration step is not included in the simulation. Therefore, we use the extracted body surface from the CT as a representative of the camera-acquired surface point cloud ${}^{\text{world}}S_{\text{body}}$ in the real system.

Experimental Procedure: For each dataset, we execute the proposed workflow. First, we register the atlas model to the body surface ${}^{\text{world}}S_{\text{CT, body}}$ to estimate the corresponding location of the target structure ${}^{\text{world}}\mathcal{O}_i$ and the base pose ${}^{\text{world}}\mathbf{b}_i^*$. Second, we plan the US scan path for the estimated target structure ${}^{\text{world}}\mathcal{O}_i$ with the mobile robot base positioned at ${}^{\text{world}}\mathbf{b}_i^*$.

Evaluation Metrics: We evaluate the *target localization* accuracy using the Average Symmetric Surface Distance (ASSD) between the estimated structure $\mathcal{O}_i$ and the CT segmentation $\mathcal{O}_{\text{CT},i}$ used as reference model. To assess the robustness of the planned scan path against these anatomical deviations, we utilize the Encompassment Scaling Factor (ESF). The ESF quantifies whether the planned ROI fully covers the target structure $\mathcal{O}_{\text{CT},i}$. The ESF measures

the uniform scaling required for the ROI bounding box to fully enclose the bounding box of $\mathcal{O}_{CT,i}$. A value of $ESF \leq 1.0$ indicates successful total encompassment, despite spatial variations in the atlas-based estimation.

We use the *target coverage* $C_i(\mathbf{b})$ to evaluate the planned scan path and to analyze the advantage of the mobile base in the MORUSS over stationary setups. We compare the target-specific base pose $\mathbf{b}_i^*$ against two stationary baselines, which are located at the longitudinal midpoint of the left or right side of the couch and referred to as “Fixed-Left” and “Fixed-Right”, respectively. To calculate the target coverage $C_i(\mathbf{b})$, we synthesize a virtual US volume along the planned path. Specifically, we generate a virtual US image at each discrete probe pose. Then, we linearly interpolate between the imaged regions of consecutive poses to approximate continuous scanning. Coverage is calculated as the volumetric overlap between this reconstructed US volume and the segmentation $\mathcal{O}_{CT,i}$ used as reference model.

B. QUANTITATIVE LAB EVALUATION

To evaluate the proposed MORUSS workflow, we conduct experiments in a controlled laboratory setting using a male body phantom containing three embedded target structures, see Fig. 3. We quantitatively evaluate the target coverage, workflow repeatability and US image quality metrics. We place gelatin-filled cavities with 3D inclusions at the anatomical locations of the liver, heart, and left femoral vein. We model the liver as a volumetric, wedge-shaped structure; the heart as an elliptical volume; and the femoral vein as a tubular structure, see Fig. 3a. We preserve the phantom’s body surface using a negative mold, shown in Fig. 3c and 3d. We create a phantom-specific atlas by scanning the phantom surface using a 3D scanner (Artec Eva, Artec 3D) and manually integrating the target structures at the locations known from the phantom’s design, see Fig. 3a.

We then scan each target five times, repositioning the MORUSS within the room before each repetition to evaluate its ability to autonomously localize and approach the target. The robot maintains probe–surface contact using proportional force control with a reference force of 2.5 N, while US gel ensures proper acoustic coupling. We transform the acquired 2D US images into the world frame using the known probe poses and (1) to reconstruct US volumes and segment the target structures.

In a first step, we evaluate whether the US scans reliably cover the whole target structure across repetitions. In a second step, we evaluate the segmentation error together with the global positioning accuracy of the mobile base. As a reference, a consensus segmentation is generated for each target structure by voxel-wise majority voting across the five repetitions. We report the ASSD, the 95th Percentile Hausdorff Distance (HD95) and the Dice Score between each segmentation and its corresponding consensus. First, the metrics are computed directly in the world frame, capturing both segmentation error and global positioning error. Then, the segmentations are aligned to the consensus using ICP to




FIGURE 3. Body phantom used for the laboratory experiments. (a) The atlas model with heart, liver, and femoral vein inclusion colored in red is shown and (b) the corresponding body phantom in the laboratory setting. (c) 3D-printed molds are used to create the gelatin phantoms that mimic organs, such as the heart. (d) The gelatin phantoms are integrated into the body phantom, preserving the body’s surface contour.

remove global localization offsets. The difference between the errors computed before and after ICP alignment provides an estimate of the global localization variability introduced by MORUSS navigation.

We evaluate the acquired US image quality across the repetitions using the Signal-to-Noise-Ratio (SNR) and contrast-to-noise ratio (CNR).

C. REAL-WORLD EVALUATION

We further demonstrate the feasibility of the proposed MORUSS workflow in a legal medicine setting using four corpses (3 male, 1 female). The study is approved by the Ethics Committee of the Hamburg Chamber of Physicians and the study protocol includes the informed consent of relatives or legal representatives. This qualitative evaluation in the legal medicine is intended to demonstrate that the MORUSS can successfully acquire US images of the target structure. Before scanning, we acquire a CT scan of the corpse and segment relevant target structures using TotalSegmentator to serve as a reference model. While the MORUSS autonomously determines the target-specific base pose, we manually navigate the MORUSS towards it with a joystick due to safety and workspace constraints. A clinician evaluates the target visibility within the acquired US images from the MORUSS by manually segmenting the visible structures. Subsequently, we evaluate the US-based segmentations with respect to the CT reference model. In order to use the CT data as reference model, we register the CT volume to the world frame using an ICP alignment of the segmented skin surface ${}^{CT}\mathcal{S}_{CT, body}$, represented in the CT frame, with the RGB-D surface scan ${}^{world}\mathcal{S}_{body}$.

IV. RESULTS AND DISCUSSION

A. VALIDATION OF THE WORKFLOW

1) TARGET LOCALIZATION

Table 2 reports the ASSD between the atlas-estimated target structure $\mathcal{O}_i$ and the segmentation $\mathcal{O}_{CT,i}$ used as

TABLE 2. Evaluation of the target structure localization approach in the simulation analysis. The mean $\pm$ standard deviation of the ASSD between the estimated $\mathcal{O}_i$ and $\mathcal{O}_{CT,i}$ is reported, along with the ESF of the ROI.

Target structure i	ASSD [mm]	ESF [–]
Liver	19.53 $\pm$ 8.39	0.70 $\pm$ 0.10
Right kidney	12.89 $\pm$ 6.52	0.58 $\pm$ 0.21
Left kidney	15.17 $\pm$ 4.69	0.97 $\pm$ 0.70
Right femur	24.56 $\pm$ 9.51	0.75 $\pm$ 0.06
Left femur	24.34 $\pm$ 10.74	0.77 $\pm$ 0.09

TABLE 3. Comparison of the reachable target coverage $C_i(\mathbf{b})$ for fixed ($\mathbf{b}_i^{\text{fixed}, L/R}$) versus target-specific ($\mathbf{b}_i^*$) base poses within the simulation analysis. Results are reported as mean $\pm$ standard deviation.

Target structure i	$C_i(\mathbf{b}_i^{\text{fixed}, R})$ [%]	$C_i(\mathbf{b}_i^{\text{fixed}, L})$ [%]	$C_i(\mathbf{b}_i^*)$ [%]
Liver	80.41 $\pm$ 12.51	10.98 $\pm$ 13.14	88.12 $\pm$ 9.55
Right kidney	98.14 $\pm$ 3.17	1.20 $\pm$ 4.30	98.05 $\pm$ 3.22
Left kidney	2.91 $\pm$ 5.31	98.76 $\pm$ 2.39	99.21 $\pm$ 2.16
Right femur	97.61 $\pm$ 2.96	7.29 $\pm$ 11.49	98.93 $\pm$ 1.09
Left femur	7.48 $\pm$ 14.25	97.63 $\pm$ 3.50	98.70 $\pm$ 1.37

reference model. The registration of the atlas model to the extracted body surfaces $^{\text{world}}\mathcal{S}_{CT, \text{body}}$ achieved successful localization across all targets, with ASSD values ranging from 12.89 mm to 24.56 mm. While these spatial deviations reflect the morphological variance between the atlas and the individual subject's anatomy, the selected scanning region contained the target structure in all cases, as indicated by the ESF values in Table 2, which are below 1.0 for all target structures.

These findings demonstrate that autonomous US imaging can be performed without prior knowledge of the exact target location. The US scanning approach provides a spatial buffer to compensate for anatomical deviations, enabling the system to exploratively scan and search for the target structure. Consequently, the proposed target localization approach is appropriate for subsequent steps in the scanning workflow.

2) TARGET COVERAGE AND BASE POSITIONING

The reachability maps in Fig. 4 visualize the reachable target coverage evaluated over the search space $\mathcal{B}$. The floor region that contains the feasible base poses $\tilde{\mathcal{B}}_{\text{valid}}$ is indicated with a dotted line. Hence, the maps illustrate the spatial distribution of base poses from which the robotic arm can reach a target coverage $> 95\%$. For contralateral targets such as the left and right kidneys, the regions with feasible base poses $\tilde{\mathcal{B}}_{\text{valid}}$ are spatially disjoint. Hence, a single stationary base pose is kinematically insufficient for scanning both targets.

In Table 3 the target coverage $C_i(\mathbf{b})$ is reported, to quantitatively evaluate the fixed baselines against the target-specific base pose. The target-specific base pose $\mathbf{b}_i^*$ significantly outperformed the contralateral baseline ($p < 0.05$) and provided more consistent results across varying body geometries, as indicated by the lower standard deviation. The target-specific base pose reached a mean coverage exceeding 95% for all structures except the liver (88.12%).

Due to the inter-subject variability in the lateral extent of the liver, for some bodies the MORUSS can not cover the whole liver structure with a single base pose. However, the MORUSS can overcome this workspace limitation by adding a second base pose on the opposite side of the couch. This results in liver coverage that exceeds 95%. Note that due to the limited US imaging depth, full coverage is not always achievable.

While the stationary ipsilateral baseline provides comparable reach, the target-specific pose $\mathbf{b}_i^*$ lies more centered within the feasible workspace $\tilde{\mathcal{B}}_{\text{valid}}$, visualized in Fig. 5. This central positioning within $\tilde{\mathcal{B}}_{\text{valid}}$ offers two main advantages. First, it provides robustness against localization inaccuracies and inter-subject anatomical variations; second, it significantly increases the manipulability of the US probe. High manipulability is critical for maintaining consistent probe-to-skin contact across the body surfaces and to facilitate explorative scanning maneuvers such as fanning and tilting. Statistical evaluation confirms higher manipulability for $\mathbf{b}_i^*$ relative to fixed baselines ($p < 0.05$ for most targets). This improvement is illustrated for the liver in Fig. 5, which contrasts the manipulability maps of the different base pose strategies. The CSA is performed once during initial system configuration using the ATLAS model (5 cm floor discretization; 1 cm skin surface resolution), requiring approximately 11.86 h. The runtime could be reduced by using a coarser spatial resolution, at the cost of reduced CSA completeness. However, the CSA characterizes the workspace of the proposed system based on its hardware configuration and is therefore independent of the individual body geometry. Consequently, it only needs to be precomputed once during the initial system configuration. During clinical application, the atlas-based approach facilitates a direct application of the precomputed CSA to the case-specific body geometry, enabling, e.g., real-time estimation of the base pose for US scanning. The simulation experiments further indicate robust target localization performance under typical anatomical variations across different body types. Future work could investigate the performance for more complex cases, such as in obesity or severe trauma cases.

B. QUANTITATIVE LAB EVALUATION

Fig. 6a shows an acquired RGB-D surface scan of the body phantom. In Fig. 6b, the segmented target structure volumes are visualized together with the RGB-D surface scan in the world frame. Table 4 reports the results of the repeated US experiments conducted in the laboratory. All US scans fully covered the whole target structure. In the world frame, the liver shows the largest localization deviations with a mean ASSD of 3.51 ± 2.44 mm and a mean HD95 of 7.78 ± 4.76 mm. These values represent the cumulative effect of segmentation inaccuracies and global localization deviations. To estimate the localization variability introduced by MORUSS navigation, we compute the difference between




FIGURE 4. Body surface of an exemplary CT scan used in the simulation analysis, with the planar xy -projection of the target structure $\mathcal{O}_i$ indicated by a blue rectangle. The floor heatmap visualizes the target coverage $C_i(\mathbf{b})$ evaluated over the search space $\mathcal{B}$. The color of each pixel quantifies the percentage of the target volume that can be reached from that specific base pose $\mathbf{b}$, ranging from 0% (black) to 100% (white). The dotted black boundary indicates the high-coverage subset $\mathcal{B}_{\text{valid}} \subseteq \mathcal{B}$, where $C_i(\mathbf{b}) \geq 95\%$. This mapping illustrates the spatial distribution of feasible base poses relative to the anatomy.




FIGURE 5. The body surface of an exemplary CT scan used in the simulation analysis is shown with the liver as target structure (dotted black rectangle). Two distinct heatmaps are visualized, with higher values corresponding to better performance. First, the grayscale floor heatmap visualizes the target coverage $C_i(\mathbf{b})$ evaluated over the search space $\mathcal{B}$. The color of each pixel quantifies the percentage of the target volume that can be reached from that specific base pose $\mathbf{b}$. Lighter regions indicate a higher percentage of the target volume reachable from that specific base pose. Second, the resulting kinematic capability to place the US probe onto the skin is projected onto the body surface for a selected base pose (red cross). This color-coded heatmap represents the local normalized manipulability $\hat{M}(\mathbf{b})$ of the US probe for each point on the skin surface. The comparison between the fixed baseline poses and the atlas-derived target-specific pose illustrates the influence of base positioning on manipulability.

the world frame error and the error obtained after ICP alignment. The largest variance was observed for the liver, with a difference of 2.55 mm in ASSD and 4.86 mm in HD95. These results indicate that the spatial localization accuracy of

the overall workflow lies in the millimeter range. The mean SNR and CNR are reported in Table 4.

The laboratory experiments confirmed complete coverage of the target structure in the US volumes across all targets and

TABLE 4. Quantitative analysis of the US scans acquired in the laboratory experiments, using ASSD, HD95, and Dice scores relative to the consensus segmentations. The three target structures liver, heart, and leg vein are evaluated within the world frame and after ICP alignment. Additionally, the image quality is analyzed by SNR and CNR. All values represent the mean $\pm$ standard deviation across $N = 5$ experimental trials.

Metric	Liver	Heart	Leg vein
<i>Initial (World Frame)</i>			
ASSD [mm]	3.51 ± 2.44	2.96 ± 1.59	1.55 ± 0.95
HD95 [mm]	7.78 ± 4.76	7.16 ± 3.48	3.68 ± 1.95
Dice Score [-]	0.75 ± 0.17	0.60 ± 0.20	0.46 ± 0.28
<i>After Registration (ICP)</i>			
ASSD [mm]	0.96 ± 0.51	0.77 ± 0.64	0.63 ± 0.35
HD95 [mm]	2.92 ± 1.36	2.59 ± 2.38	1.68 ± 0.69
Dice Score [-]	0.93 ± 0.04	0.90 ± 0.08	0.76 ± 0.14
<i>US Image Metrics</i>			
SNR [dB]	16.34 ± 0.73	17.42 ± 0.30	15.08 ± 1.38
CNR [-]	1.78 ± 0.03	1.09 ± 0.03	0.65 ± 0.06




FIGURE 6. Visualization of the acquired body phantom data. One representative scan per target in the world frame is shown. (a) RGB-D body surface scan of the body phantom with femoral vein (A), liver (B), and heart phantom (C). (b) For each target, the segmented inclusions (red) from the acquired US scan volume are represented within the world frame F_{world} . (c) A reconstructed US volume of the leg vein phantom corresponding to position (A) is visualized using MATLAB volume rendering with a linear alpha map and grayscale colormap. (d) The same reconstructed US volume is shown with the segmented leg vein overlaid in orange.

varying starting positions of the MORUSS, demonstrating the repeatability of the scanning workflow. The observed localization errors of approximately 5 mm in laboratory experiments were compensated by a local RGB-D surface acquisition prior to US scanning. The observed Dice Score variability, particularly for small structures such as the femoral vein, reflected segmentation sensitivity rather than incomplete imaging. These results demonstrate that the MORUSS can reliably perform autonomous US scans of different target structures.

The US image quality was sufficient for volumetric reconstruction and structure segmentation. However, future




FIGURE 7. Visualization of legal medicine data from an example corpse. All modalities are registered in the world frame. (a) The registered and segmented CT data serve as a spatial reference and are used to validate that the US scans captured the intended anatomical target. (b) We visualize the RGB-D scan and the CT volume using a checkerboard overlay to illustrate their spatial alignment. Note that we blurred areas for anonymity. (c) Example US images of the liver and (d) a vessel.

iterations should focus on optimizing the US image parameters to fulfill diagnostic standards.

C. REAL-WORLD EVALUATION

Fig. 7 shows a checkerboard-style multi-modality visualization of the corpse in 3D Slicer, including the segmented CT target structures and the RGB-D surface point cloud. The RGB-D scan captures the body surface in detail, including features such as a bruise from cardiopulmonary resuscitation. The corresponding CT and US images confirm spatial alignment and anatomical correspondence between modalities. Although diagnostic image quality was not a primary endpoint of this study, visual review by two experienced clinicians confirmed that relevant anatomical structures and landmarks were identifiable and clinically interpretable, including organ boundaries, major vascular structures, and relevant soft-tissue interfaces. The resulting image quality was considered comparable to that routinely




FIGURE 8. US and CT data acquired in the legal medicine. (Left column) For each corpse, the 3D CT scan is cropped at the plane of a representative US image to visualize the internal anatomy. (Right column) The corresponding 2D view of the CT cross-section along the US imaging plane is shown with the US image overlaid. All modalities are registered in the world frame.

encountered during conventional post mortem US examinations with manual probe placement on the body surface. This demonstrates the capability of the MORUSS to perform US acquisition within a legal medicine framework, where all modalities (RGB-D, CT, US) are consistently aligned within the world frame to provide a multimodal visualization for the clinician.

Fig. 8 presents representative US images for all corpses, overlaid with the corresponding CT slice and the segmented target structure from the clinician. Generally, the US images successfully visualize the intended target structures. The spatial validation with the CT reference model confirms anatomical correctness. In Fig. 8d, the kidney appears partially truncated in the US image due to acoustic shadowing from bone structures. Note, that in future work these effects could be further mitigated by advanced scan path strategies that consider adaptive probe orientation to optimize acoustic windows and improve image quality [53]. Beyond that, there is further potential to enhance performance using machine learning methods, in particular for US image quality assessment [27], [54], [55] and localization of anatomical structures in point clouds [19], which could further refine the workflow.

The legal medicine experiments were conducted on four corpses, representing typical anatomical body types in legal medicine. Cases involving more complex body types, such as severe obesity, were not considered within the scope of this study. Nevertheless, these cases represent a relevant direction for future work, as they may introduce additional challenges and potentially limit US imaging performance due to acoustic penetration constraints. However, based on expert assessment by the co-authors active in legal medicine, most routine forensic examinations involve body types within the normal range and are not substantially affected by extreme obesity or severe trauma. Therefore, the selected corpses are considered representative of common forensic practice. Hence, despite the limited sample size, the experiments demonstrated the robustness of the proposed system within the scope of this study.

Beyond technical development, ethical and regulatory aspects are important for the future integration of autonomous robotic US into clinical workflows. Recently, these aspects have gained interest in research [13], [56]. This also includes operational requirements such as interaction with clinical personnel, advanced obstacle detection, and standardized calibration procedures that should be regularly performed within the workflow to ensure sustained system precision. One challenge in this context is that forensic environments are designed for human operation and are often narrow and cluttered. This poses additional challenges for reliable autonomous navigation and safety-aware system behavior.

Overall, the results demonstrate the feasibility of robust autonomous US scanning with the proposed AI-supported framework. As outlined in Table 1, the MORUSS differs from other approaches by integrating a mobile platform and atlas-based internal anatomy localization, enabling fully autonomous US scanning across the entire body tested in

a post mortem legal medicine setting. The legal medicine experiments verify the system's real-world applicability, particularly its ability to reliably capture clinically relevant structures. Furthermore, the MORUSS does not require direct contact between clinician and corpse, which can be beneficial in legal medicine scenarios with, e.g., highly infectious corpses.

V. CONCLUSION

In this work, we introduced a novel mobile robotic US system. In contrast to typical approaches for robotic US in clinical settings, the proposed system can move autonomously to reach target structures across the whole body. Moreover, the system can autonomously localize a body, estimate the anatomical target structure within, and generate a US scan path covering the respective structure. This allows for an unprecedented level of autonomy, enabling US image acquisition for targets along the full body and with minimal human contact and intervention. In the context of legal medicine, this level of autonomy could reduce clinician workload and help lower infection risks.

To this end, we evaluated different aspects of the system. First, we showed that suitable positions of the system's mobile base can be determined based on a surface scan and an atlas model. Second, our lab evaluation illustrated that the system can robustly reach the planned target regions with sufficient precision to ensure full coverage of the respective structures in the US images. Finally, we demonstrated that the approach is feasible in a real legal medicine setting.

Beyond legal medicine, mobile robotic US may also be beneficial in other clinical settings, such as telemedicine. The proposed approach does not require a fixed setup with respect to the body, and motion planning and US scanning can be realized without human intervention. Hence, the system represents an important step towards fully autonomous US imaging.

ACKNOWLEDGMENT

We would like to acknowledge that the RB-Theron was donated by ROBAST Robotic Assistant GmbH, Hamburg, Germany.

REFERENCES

- [1] D. Möbius, A. Fitzek, N. Hammer, A. Heinemann, A. Ron, J. Schädler, J. Zwirner, and B. Ondruschka, "Ultrasound in legal medicine—A missed opportunity or simply too late? A narrative review of ultrasonic applications in forensic contexts," *Int. J. Legal Med.*, vol. 135, no. 6, pp. 2363–2383, Nov. 2021.
- [2] F. Dedouit, M. Ducloyer, J. Elifritz, N. L. Adolphi, G. W. Yi-Li, S. Decker, J. Ford, Y. Kolev, and M. Thali, "The current state of forensic imaging—clinical forensic imaging," *Int. J. Legal Med.*, vol. 139, no. 4, pp. 1639–1646, 2025.
- [3] J. L. Burton, "Health and safety at necropsy," *J. Clin. Pathol.*, vol. 56, no. 4, pp. 254–260, Apr. 2003.
- [4] J. M. Brandner, P. Boor, L. Borchering, C. Edler, S. Gerber, A. Heinemann, J. Hilsenbeck, A. Kasajima, L. Lohner, B. Märkl, J. Pablik, A. S. Schröder, J. Slotta-Huspenina, L. Sommer, J.-P. Spermhake, S. von Stillfried, and S. Dintner, "Contamination of personal protective equipment during COVID-19 autopsies," *Virchows Archiv*, vol. 480, no. 3, pp. 519–528, Mar. 2022.

- [5] L. Geoffroy, L. Tuchtan, M.-D. Piercecchi-Marti, and C. Delteil, "Post-mortem transmission risk of infectious disease: A systematic review," *Legal Med.*, vol. 71, Nov. 2024, Art. no. 102530.
- [6] W. F. Mawalla, "The contribution of autopsy in COVID-19 pandemic: Missed opportunity in sub-Saharan Africa," *Forensic Sci. Int., Rep.*, vol. 2, Dec. 2020, Art. no. 100136.
- [7] S. Murphey, "Work related musculoskeletal disorders in sonography," *J. Diagnostic Med. Sonography*, vol. 33, no. 5, pp. 354–369, Sep. 2017.
- [8] Z. Zangiabadi, F. Makki, H. Marzban, F. Salehinejad, A. Sahebi, and S. Tahernejad, "Musculoskeletal disorders among sonographers: A systematic review and meta-analysis," *BMC Health Services Res.*, vol. 24, no. 1, p. 1233, Oct. 2024.
- [9] K. J. Hollitt, S. Milanese, M. Joseph, and R. Perry, "Can automation and artificial intelligence reduce echocardiography scan time and ultrasound system interaction?" *Echo Res. Pract.*, vol. 12, no. 1, p. 11, Jun. 2025.
- [10] J. S. Abramowicz, J. M. Basseal, C. Brezinka, A. Dall'Asta, J. Deng, G. Harrison, J. C. S. Lee, A. Lim, K. Maršal, P. Miloro, L. C. Poon, K. Å. Salvesen, R. Sande, G. T. Haar, S. C. Westerway, M. X. Xie, and C. Lees, "ISUOG safety committee position statement on use of personal protective equipment and hazard mitigation in relation to SARS-CoV-2 for practitioners undertaking obstetric and gynecological ultrasound," *Ultrasound Obstetrics Gynecology*, vol. 55, no. 6, pp. 886–891, Jun. 2020.
- [11] Z. Y. Zu, M. D. Jiang, P. P. Xu, W. Chen, Q. Q. Ni, G. M. Lu, and L. J. Zhang, "Coronavirus disease 2019 (COVID-19): A perspective from China," *Radiology*, vol. 296, no. 2, pp. E15–E25, Aug. 2020.
- [12] M. Akbari, J. Carriere, T. Meyer, R. Sloboda, S. Husain, N. Usmani, and M. Tavakoli, "Robotic ultrasound scanning with real-time image-based force adjustment: Quick response for enabling physical distancing during the COVID-19 pandemic," *Frontiers Robot. AI*, vol. 8, Mar. 2021, Art. no. 645424.
- [13] Z. Jiang, S. E. Salcudean, and N. Navab, "Robotic ultrasound imaging: State-of-the-art and future perspectives," *Med. Image Anal.*, vol. 89, Oct. 2023, Art. no. 102878.
- [14] K. Munir, A. F. Al-Battal, A. Alsheghri, H. Becher, M. Noga, and K. Punithakumar, "A survey of autonomous robotic ultrasound scanning systems," *IEEE Access*, vol. 13, pp. 103178–103197, 2025.
- [15] F. von Haxthausen, S. Böttger, D. Wulff, J. Hagenah, V. García-Vázquez, and S. Ipsen, "Medical robotics for ultrasound imaging: Current systems and future trends," *Current Robot. Rep.*, vol. 2, no. 1, pp. 55–71, Mar. 2021.
- [16] P. M. Kebria, S. Nahavandi, N. Kutaiba, N. Koutrouza, N. Yang, H. Asadi, and G. Guest, "HERCULES: Haptically-enabled remotely controlled ultrasound examination system," *IEEE Access*, vol. 13, pp. 75585–75598, 2025.
- [17] P. M. Kebria, S. Nahavandi, and R. D. Howe, "Learning human-robot interactions in perturbed teleoperations," in *Proc. IEEE 21st Int. Conf. Autom. Sci. Eng. (CASE)*, Aug. 2025, pp. 1987–1993.
- [18] M. Komaritzan, S. Wenninger, and M. Botsch, "Inside humans: Creating a simple layered anatomical model from human surface scans," *Frontiers Virtual Reality*, vol. 2, Jul. 2021, Art. no. 694244.
- [19] P. Henrich and F. Mathis-Ullrich, "LOOC: Localizing organs using occupancy networks and body surface depth images," *IEEE Access*, vol. 13, pp. 36930–36938, 2025.
- [20] Z. Jiang, Y. Gao, L. Xie, and N. Navab, "Towards autonomous atlas-based ultrasound acquisitions in presence of articulated motion," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 7423–7430, Jul. 2022.
- [21] D. Sun, A. Cappellari, B. Lan, M. Abayazid, S. Stramigioli, and K. Niu, "Automatic robotic ultrasound for 3D musculoskeletal reconstruction: A comprehensive framework," *Technologies*, vol. 13, no. 2, p. 70, Feb. 2025.
- [22] A. S. B. Mustafa, T. Ishii, Y. Matsunaga, R. Nakadate, H. Ishii, K. Ogawa, A. Saito, M. Sugawara, K. Niki, and A. Takanishi, "Development of robotic system for autonomous liver screening using ultrasound scanning device," in *Proc. IEEE Int. Conf. Robot. Biomimetics (ROBIO)*, Dec. 2013, pp. 804–809.
- [23] K. Su, J. Liu, X. Ren, Y. Huo, G. Du, W. Zhao, X. Wang, B. Liang, D. Li, and P. X. Liu, "A fully autonomous robotic ultrasound system for thyroid scanning," *Nature Commun.*, vol. 15, no. 1, p. 4004, May 2024.
- [24] X. Tang, H. Wang, J. Luo, J. Jiang, F. Nian, L. Qi, L. Sang, and Z. Gan, "Autonomous ultrasound scanning robotic system based on human posture recognition and image servo control: An application for cardiac imaging," *Frontiers Robot. AI*, vol. 11, May 2024, Art. no. 1383732.
- [25] X.-X. Lin, M.-D. Li, S.-M. Ruan, W.-P. Ke, H.-R. Zhang, H. Huang, S.-H. Wu, M.-Q. Cheng, W.-J. Tong, H.-T. Hu, D.-N. He, R.-F. Lu, Y.-D. Lin, M. Kuang, M.-D. Lu, L.-D. Chen, Q.-H. Huang, and W. Wang, "Autonomous robotic ultrasound scanning system: A key to enhancing image analysis reproducibility and observer consistency in ultrasound imaging," *Frontiers Robot. AI*, vol. 12, Feb. 2025, Art. no. 1527686.
- [26] S. Grube, S. Latus, M. Fischer, V. Raudonis, A. Heinemann, B. Ondruschka, and A. Schlaefler, "A mobile robotic approach to autonomous surface scanning in legal medicine," *Int. J. Comput. Assist. Radiol. Surgery*, vol. 21, no. 3, pp. 575–584, Sep. 2025.
- [27] Z. Jiang, Z. Li, M. Grimm, M. Zhou, M. Esposito, W. Wein, W. Stechele, T. Wendler, and N. Navab, "Autonomous robotic screening of tubular structures based only on real-time ultrasound imaging feedback," *IEEE Trans. Ind. Electron.*, vol. 69, no. 7, pp. 7064–7075, Jul. 2022.
- [28] X. Bao, S. Wang, L. Zheng, R. J. Housden, J. V. Hajnal, and K. Rhode, "A novel ultrasound robot with force/torque measurement and control for safe and efficient scanning," *IEEE Trans. Instrum. Meas.*, vol. 72, pp. 1–12, 2023.
- [29] Q. Huang, B. Gao, and M. Wang, "Robot-assisted autonomous ultrasound imaging for carotid artery," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–9, 2024.
- [30] Z. Wang, Y. Han, B. Zhao, H. Xie, L. Yao, B. Li, M. Q.-H. Meng, and Y. Hu, "Autonomous robotic system for carotid artery ultrasound scanning with visual servo navigation," *IEEE Trans. Med. Robot. Bionics*, vol. 6, no. 4, pp. 1436–1447, Nov. 2024.
- [31] Y. Zheng, H. Ning, E. Rangarajan, A. Merali, A. Geale, L. Lindenroth, Z. Xu, W. Wang, P. Kruse, S. Morris, L. Ye, X. Fu, K. Rhode, and R. J. Housden, "Design of a cost-effective ultrasound force sensor and force control system for robotic extra-body ultrasound imaging," *Sensors*, vol. 25, no. 2, p. 468, Jan. 2025.
- [32] M. K. Welleweerd, A. G. de Groot, V. Groenhuis, F. J. Siepel, and S. Stramigioli, "Out-of-plane corrections for autonomous robotic breast ultrasound acquisitions," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2021, pp. 12515–12521.
- [33] F. Suligoj, C. M. Heunis, J. Sikorski, and S. Misra, "RobUSt—An autonomous robotic ultrasound system for medical imaging," *IEEE Access*, vol. 9, pp. 67456–67465, 2021.
- [34] J. Osburg, A. Scheibert, M. Horn, R. Pater, and F. Ernst, "Automatic robotic Doppler sonography of leg arteries," *Int. J. Comput. Assist. Radiol. Surgery*, vol. 19, no. 10, pp. 1965–1974, Jul. 2024.
- [35] Q. Huang, B. Wu, J. Lan, and X. Li, "Fully automatic three-dimensional ultrasound imaging based on conventional B-scan," *IEEE Trans. Biomed. Circuits Syst.*, vol. 12, no. 2, pp. 426–436, Apr. 2018.
- [36] Q. Huang, J. Lan, and X. Li, "Robotic arm based automatic ultrasound scanning for three-dimensional imaging," *IEEE Trans. Ind. Informat.*, vol. 15, no. 2, pp. 1173–1182, Feb. 2019.
- [37] Z. Wang, B. Zhao, P. Zhang, L. Yao, Q. Wang, B. Li, M. Q.-H. Meng, and Y. Hu, "Full-coverage path planning and stable interaction control for automated robotic breast ultrasound scanning," *IEEE Trans. Ind. Electron.*, vol. 70, no. 7, pp. 7051–7061, Jul. 2023.
- [38] J. Tan, Y. Li, B. Li, Y. Leng, J. Peng, J. Wu, B. Luo, X. Chen, Y. Rong, and C. Fu, "Automatic generation of autonomous ultrasound scanning trajectory based on 3-D point cloud," *IEEE Trans. Med. Robot. Bionics*, vol. 4, no. 4, pp. 976–990, Nov. 2022.
- [39] C. Wu, H. Lan, J. Ma, X. Li, and J. Wen, "Point cloud-guided ultrasound robotic scanning path planning for the kidney based on anatomical positioning," *Comput. Biol. Med.*, vol. 192, Jun. 2025, Art. no. 110191.
- [40] X. Ma, Z. Zhang, and H. K. Zhang, "Autonomous scanning target localization for robotic lung ultrasound imaging," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2021, pp. 9467–9474.
- [41] Y. Huang, W. Xiao, C. Wang, H. Liu, R. Huang, and Z. Sun, "Towards fully autonomous ultrasound scanning robot with imitation learning based on clinical protocols," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 3671–3678, Apr. 2021.
- [42] H. Schlehlein, B. W. Herr II, E. M. Quardokus, A. Bueckle, and K. Börner, "HuBMAP CCF 3D Reference Object Library," 2022. Accessed: Dec. 15, 2022. [Online]. Available: <https://humanatlas.io/3d-reference-library?version=11&organ=All%20organs>
- [43] T. Yoshikawa, "Manipulability of robotic mechanisms," *Int. J. Robot. Res.*, vol. 4, no. 2, pp. 3–9, Jun. 1985.
- [44] 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, Nov. 2021.
- [45] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Y. Ng, "ROS: An open-source robot operating system," in *Proc. ICRA Workshop Open Source Softw.*, 2009.
- [46] D. Coleman, I. A. Şucan, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: A moveit! Case study," *J. Softw. Eng. Robot.*, vol. 5, no. 1, pp. 3–16, 2014.

- [47] S. Macenski, F. Martín, R. White, and J. G. Clavero. (2020). *The ROS 2 Navigation System (NAV2)*. Accessed: Nov. 10, 2025. [Online]. Available: <https://navigation.ros.org>
- [48] S. Macenski and I. Jambrecic, "SLAM toolbox: SLAM for the dynamic world," *J. Open Source Softw.*, vol. 6, no. 61, p. 2783, May 2021.
- [49] F. C. Park and B. J. Martin, "Robot sensor calibration: Solving $AX=XB$ on the Euclidean group," *IEEE Trans. Robot. Autom.**(1989–June 2004), vol. 10, no. 5, pp. 717–721, Oct. 1994.
- [50] A. Perez-Yus, E. Fernandez-Moral, G. Lopez-Nicolas, J. J. Guerrero, and P. Rives, "Extrinsic calibration of multiple RGB-D cameras from line observations," *IEEE Robot. Autom. Lett.*, vol. 3, no. 1, pp. 273–280, Jan. 2018.
- [51] R. Li, Y. Cai, K. Niu, and E. V. Poorten, "Comparative quantitative analysis of robotic ultrasound image calibration methods," in *Proc. 20th Int. Conf. Adv. Robot. (ICAR)*, Dec. 2021, pp. 511–516.
- [52] J. Wasserthal, H.-C. Breit, M. T. Meyer, M. Pradella, D. Hinck, A. W. Sauter, T. Heye, D. T. Boll, J. Cyriac, S. Yang, M. Bach, and M. Segeroth, "TotalSegmentator: Robust segmentation of 104 anatomic structures in CT images," *Radiology: Artif. Intell.*, vol. 5, no. 5, Sep. 2023, Art. no. 230024.
- [53] Y. Bi, C. Qian, Z. Zhang, N. Navab, and Z. Jiang, "Autonomous path planning for intercostal robotic ultrasound imaging using reinforcement learning," *Sci. Rep.*, vol. 16, no. 1, p. 6356, Feb. 2026.
- [54] Y. Bi, Z. Jiang, F. Duermel, D. Huang, and N. Navab, "Machine learning in robotic ultrasound imaging: Challenges and perspectives," *Annu. Rev. Control, Robot., Auto. Syst.*, vol. 7, no. 1, pp. 335–357, Jul. 2024.
- [55] Y. Song, Z. Zhong, B. Zhao, P. Zhang, Q. Wang, Z. Wang, L. Yao, F. Lv, and Y. Hu, "Medical ultrasound image quality assessment for autonomous robotic screening," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 6290–6296, Jul. 2022.
- [56] M. Nadeem, Y. I. Alzoubi, E. Elbasi, and A. E. Topcu, "Robot-assisted surgery: A comprehensive review of literature, challenges, ethical considerations, and current trends," *IEEE Access*, vol. 13, pp. 215647–215680, 2025.

SARAH GRUBE received the B.Sc. and M.Sc. degrees in mechanical engineering from Hamburg University of Technology, Hamburg, in 2019 and 2022, respectively, where she is currently pursuing the Ph.D. degree with the Institute of Medical Technology and Intelligent Systems. Her research interest includes robotic ultrasound.

SARAH LATUS received the M.Sc. degree in electrical engineering, and the Ph.D. degree from Hamburg University of Technology (TUHH), Germany, in 2025. She is currently a Senior Research Engineer with the Institute of Medical Technology and Intelligent Systems, TUHH. Her research interests include soft tissue characterization using miniaturized imaging probes, particularly ultrasound and optical coherence tomography, combined with advanced image and signal processing.

DUSTIN MOEBIUS received the Medical Doctor degree from the University of Leipzig. He undertook his specialist training in internal medicine in a communal clinic in Itzehoe, Germany. He is currently a Resident with the Institute of Legal Medicine, University Medical Center Hamburg-Eppendorf, Hamburg, Germany, finishing his training in legal medicine. His research interests include implementing clinical approaches and their feasibility in a post-mortem setting.

MARTIN FISCHER received the B.Sc. and M.Sc. degrees in mechanical engineering from the University of Rostock. Following several positions at different universities and research institutions, he is currently with the Institute of Medical Technology and Intelligent Systems, Hamburg University of Technology. In his current role, he is engaged in the design and implementation of experimental test systems and offering technical expertise to support ongoing scientific research.

VIDAS RAUDONIS received the Ph.D. degree in computer science and informatics engineering from Kaunas University of Technology, Lithuania, in 2010. He is currently a Professor with the Department of Automation, Faculty of Electrical and Electronics Engineering, Kaunas University of Technology. He is also a Principal Investigator and is actively involved with the Artificial Intelligence Competence Centre for Sustainable Living and Working (SustAIInLivWork). He previously served as an Expert of Innovations and Research at medical device companies, including ESCO Medical and AT Medical, and was the Founder and the CEO of the Computer Vision Company JSC "Power of Eye." He has published numerous scientific papers in various international journals and conferences. Additionally, he leads the "Smart Automatic Control Systems" Research Group and supervises multiple Ph.D. students in the field of computational intelligence and computer vision. His research interests include computational intelligence, computer vision, human–computer interaction, artificial neural networks, assistive technology, and robotics.

AXEL HEINEMANN is currently a German Forensic Pathologist and the Assistant Medical Director of the Institute of Legal Medicine, University Medical Center Hamburg-Eppendorf (UKE). His research interest includes method development for forensic imaging applications.

BENJAMIN ONDRUSCHKA is currently a German Forensic Pathologist and the Director of the Institute of Legal Medicine, University Medical Center Hamburg-Eppendorf (UKE). He holds an academic leadership role combining casework, research, and teaching. As a Principal Investigator in the National Autopsy Network, he has contributed to the integration of these techniques into modern forensic autopsy practice. His research interests include forensic neurotraumatology, postmortem biochemistry, and advanced post-mortem imaging methods.

ALEXANDER SCHLAEFER (Senior Member, IEEE) is currently a Professor with Hamburg University of Technology, where he heads the Institute of Medical Technology and Intelligent Systems. He is also involved in the Artificial Intelligence Competence Centre for Sustainable Living and Working (SustAIInLivWork) and the Interdisciplinary Competence Center for Interface Research (ICCIR). His research interests include robotics, navigation, and machine learning, particularly for clinical applications.

• • •