Laplacian Trajectory Editing for Robotic Ultrasound Systems
Adapting Scan Trajectories to Patient Motion

Koelmans, Toine; Mol, N.; Prendergast, J. Micah

DOI
10.1109/Humanoids65713.2025.11203100

Publication date
2025

Document Version
Final published version

Published in
Proceedings of the IEEE-RAS 24th International Conference on Humanoid Robots, Humanoids 2025

Citation (APA)
Koelmans, T., Mol, N., & Prendergast, J. M. (2025). Laplacian Trajectory Editing for Robotic Ultrasound Systems: Adapting Scan Trajectories to Patient Motion. In *Proceedings of the IEEE-RAS 24th International Conference on Humanoid Robots, Humanoids 2025* (pp. 397-404). (IEEE-RAS International Conference on Humanoid Robots). IEEE. https://doi.org/10.1109/Humanoids65713.2025.11203100

2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids)
September 30 – October 2, 2025 | COEX, Seoul, Korea

Laplacian Trajectory Editing for Robotic Ultrasound Systems: Adapting Scan Trajectories to Patient Motion

Toine Koelmans[1], Nicky Mol[1], J. Micah Prendergast[1]

Abstract— **Robotic Ultrasound Systems (RUSS) provide a promising solution to reduce operator dependency, alleviate physical strain, and meet the growing demand for ultrasound procedures. However, their clinical applicability remains limited by their inability to adapt to dynamic patient movements and tissue deformations during scans. This work introduces a novel framework that leverages Laplacian Trajectory Editing (LTE) for real-time adaptation of scan trajectories in response to both rigid and non-rigid patient movements. it integrates a RGB-D camera to capture surface point clouds, which are processed to estimate displacements between consecutive frames. These displacements define anchor points for LTE-based trajectory adaptations, ensuring smooth motion while preserving local trajectory properties. This approach is validated through experiments spanning rigid phantom movements, generalization across differently shaped phantoms, and non-rigid human arm motion. Adaptation accuracy is quantified by comparing adapted trajectories to a ground-truth reference, with root mean squared errors averaging 0.026 ± 0.012 m in non-rigid scenarios. Real-time trajectory adaptation is achieved, with an average LTE adaptation processing time of 373 ms per trial. Furthermore, our implementation achieved low tracking errors across all conditions while maintaining a high success rate in diverse movement scenarios. These results demonstrate the feasibility of LTE for real-time trajectory adaptation in ultrasound scanning, offering a pathway to more autonomous and clinically viable RUSS implementations.**

I. INTRODUCTION

Musculoskeletal ultrasound (MSK-US) is a non-invasive means for visualizing tissue that relies on patient joint articulation to visualize underlying tissue and reveal musculoskeletal abnormalities [1]. This requisite patient movement, while diagnostically crucial, poses a significant challenge for maintaining probe alignment and imaging consistency, necessitating continuous real-time trajectory adaptation [2].

Manual MSK-US procedures, integral to modern healthcare, are highly operator dependent, with image quality limited by the skills of the sonographer as they carefully adjust and maneuver the US probe during patient motion. Furthermore, MSK-US and US procedures in general suffer from sonographer shortages, ergonomic injuries and cognitive overload [3], [4]. Robotic Ultrasound Systems (RUSS) may offer a promising solution to these problems, however, to date, research in RUSS has predominantly focused on imaging solutions that disregard the complexity of physical human-robot interaction (pHRI) during such a procedure [5]–[7]. For robotic MSK-US procedures, however, the challenge entailed in how the robot should control its trajectory in response to patient motions and interaction forces is a crucial factor in ensuring high-quality imaging.

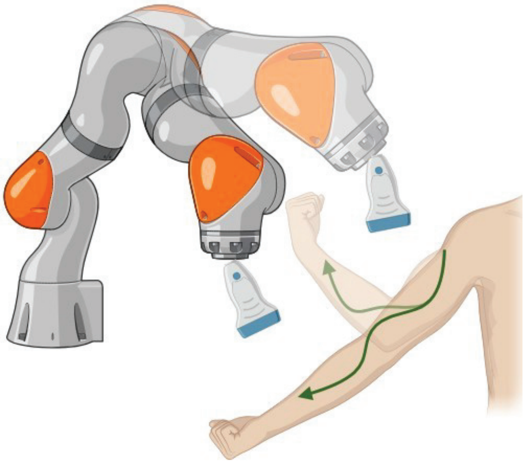

Fig. 1: Example of a Robotic Ultrasound System (RUSS) performing musculoskeletal ultrasound (MSK-US), where continuous patient movement is required for proper imaging. The patient bends their arm while the robot performs the scan, adapting its trajectory in real-time to maintain probe alignment.

Previous work on autonomous robotic ultrasound includes Reinforcement Learning (RL), Learning from Demonstration (LfD) and geometry-based approaches. While RL and LfD based approaches have shown success in adapting to US tasks, they typically require large datasets, lack interpretability and are not able to respond to patient motion. In contrast, geometry-based approaches offer human interpretable heuristics that facilitate more predictable behavior and safety limits, but like their learning-based counterparts, are not robust to geometric changes that result from tissue deformation and patient motion.

Recognizing that dynamic adaptation to patient motion and tissue deformation is critical for robotic MSK-US systems, several studies have attempted to address this. Efforts include adjusting scan paths based on estimates of tissue deformation [8], predefined trajectory mapping to vascular structures as patients move in real-time [9], compensating movement by comparing RGB-D images to CT point clouds [10], and integrating force control, deep-learning-based segmentation

---

This work was supported by the BrightSky project, funded by the R&D Mobiliteitsfonds from the Netherlands Enterprise Agency (RVO) and commissioned by the Ministry of Economic Affairs and Climate Policy.

[1] Department of Cognitive Robotics, Faculty of Mechanical Engineering, Delft University of Technology, Delft, The Netherlands (e-mail: j.m.prendergast@tudelft.nl)

2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids) | 979-8-3315-9869-3/25/$31.00 ©2025 IEEE | DOI: 10.1109/HUMANOIDS65713.2025.11203100

979-8-3315-9869-3/25/$31.00 ©2025 IEEE

and 3D reconstruction for MSK-US [11]. However, these specialized solutions often address only specific types of motion or lack the capability for generalized, real-time adaptation to non-rigid surface changes as observed during dynamic patient movement. Consequently, robust and real-time adaptation to tissue deformation and non-rigid patient movement remains a critical challenge in the clinical adoption of RUSS for use in MSK-US.

To address this gap, we turn to trajectory adaptation techniques established in the broader field of robotics. Methods based on movement primitives like Dynamics Movement Primitives (DMPs), Probabilistic Movement Primitives (ProMPs) and Kernelized Movement Primatives (KMPs) [12]–[14] have been used for encoding and adapting demonstrated trajectories, often from a single example, but are primarily designed for rigid transformations. Other approaches, such as Gaussian Mixture Models (GMMs) and Inverse Reinforcement Learning (IRL) [15], [16] typically require multiple demonstrations. In contrast, Laplacian Trajectory Editing (LTE) [17] offers a compelling alternative suited to the challenge of real-time adaption to non-rigid surface changes: it allows adaptation of trajectories to non-rigid deformations, providing smooth trajectories that preserve local properties, while needing only a single demonstration.

In this work we propose and validate a novel trajectory adaptation framework for RUSS that leverages LTE, informed by real-time RGB-D surface tracking, to achieve dynamic trajectory adaptation to both rigid and non-rigid patient motion after a single demonstration. We validate this approach through proof-of-concept experiments involving anatomical phantoms and non-rigid human arm motion, assessing its real-time performance and ability to generalize to varying patient motion and morphology.

II. METHODS

Figure 2 illustrates the proposed approach. The process begins with an initial kinesthetic demonstration to capture a baseline scan trajectory (Fig. 2a, section II-A). Concurrently, the surface of the scanned region is extracted using RGB-D camera data (Fig. 2b, section II-B). Surface displacement resulting from patient movement, are estimated using point cloud registration (Fig. 2c, section II-C). These estimated displacements serve as input for LTE, which modifies the original trajectory (Fig. 2d, section II-D).

A. Initial Trajectory Demonstration

To ensure clinical relevance, initial trajectories were based on established ultrasound protocols and obtained through kinesthetic demonstrations, where an operator manually guided the robot's end-effector (holding the mock ultrasonic probe) along the desired path. For experiments involving the **abdominal phantom**, the Obstetric Volume Sweep Imaging (VSI) protocol [18] was used, providing a structured and repeatable sweep from the pelvic region towards the upper abdomen. For the **non-rigid arm experiments**, a trajectory representative of longitudinal MSK-US scans of the upper limb [1] was used, involving a continuous scan from the shoulder to the wrist on a static arm; this served as the baseline path for subsequent dynamic adaptation trials. During both the VSI and MSK-US kinesthetic demonstrations, the end-effector pose was recorded at 200 Hz to obtain the trajectory:

$$\mathcal{T} = \{\mathbf{T}_i\}_{i=1}^{n}, \tag{1}$$

where each pose T_i consists of:

$$\mathbf{T}_i = \begin{bmatrix} \mathbf{p}_i \\ \mathbf{q}_i \end{bmatrix}, \tag{2}$$

with $\mathbf{p}_i = (x_i, y_i, z_i)^\top \in \mathbb{R}^3$ representing the position of the US probe, and $\mathbf{q}_i = (q_{w,i}, q_{x,i}, q_{y,i}, q_{z,i})^\top \in \mathbb{H}$ as the orientation given as a unit quaternion. For trajectory adaptation using LTE, we define the set of position waypoints in Cartesian space as:

$$\mathcal{T}_p = \{\mathbf{p}_i\}_{i=1}^{n}, \quad \mathbf{p}_i \in \mathbb{R}^3, \tag{3}$$

and the set of orientation waypoints in Cartesian space as:

$$\mathcal{T}_o = \{\mathbf{q}_i\}_{i=1}^{n}, \quad \mathbf{q}_i \in \mathbb{H}, \tag{4}$$

where $\mathcal{T}_p$ represents the positional data and $\mathcal{T}_o$ the orientational data of the trajectory.

B. Surface Extraction

To estimate surface displacement for trajectory adaptation, relevant geometric information defining the target surface was extracted from the RGB-D stream. For **phantom-based experiments**, HSV segmentation isolated the phantom using predefined color thresholds. The resulting 2D ROI was projected onto the depth image to generate a dense 3D point cloud which was filtered to remove the table plane.

For **non-rigid arm tracking**, color segmentation was unreliable due to motion and appearance variation. Instead, anatomical keypoints including the shoulder, elbow, wrist and thumb were detected using MediaPipe [19] and projected onto 3D, resulting in the surface representation used for displacement estimation and trajectory adaptation.

Regardless of the extraction method, all data was transformed into the robot's base frame, providing an intuitive global reference frame. Camera extrinsic information needed for this transformation was determined via a visual calibration using a static ArUco marker placed at a known position on the robot's base.

C. Surface Displacement Estimation

To adapt the trajectory to surface changes the relative displacement between the current and previous observation needed to be estimated. This estimation was triggered when the observed distance exceeded a predefined threshold. For **phantom-based** scenarios, registration was triggered if the Chamfer distance between point clouds $\geq$ 0.05 m. In **non-rigid arm tracking**, registration was triggered if the sum of pairwise Euclidean distances between corresponding keypoints was $\geq$ 0.05 m.

Once registration was initiated, the source and target point clouds were defined to estimate displacement. The source

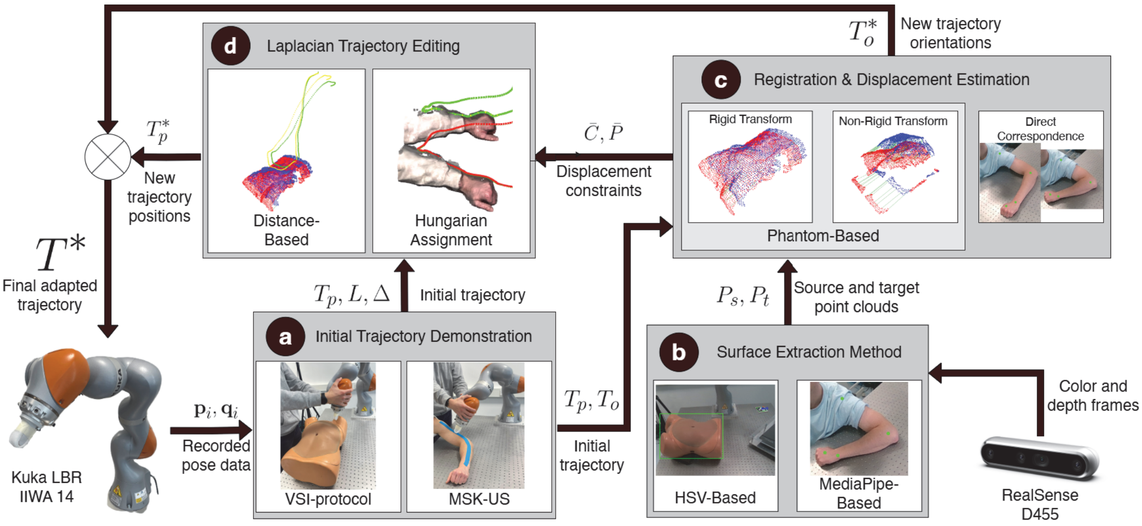

Fig. 2: Overview of the proposed RUSS. An initial kinesthetic demonstration of the scan trajectory is obtained (a), the scanned region's surface is extracted using an RGB-D camera (b), surface displacements are estimated via point cloud registration (c), and the trajectory is adapted to these displacements using LTE (d). This framework enables real-time trajectory adaptation to accommodate rigid and non-rigid surface changes.

point cloud, $\mathcal{P}_s$, consists of M points $\mathbf{p}_i^s = (x_i^s, y_i^s, z_i^s)^\top$, as shown in:

$$\mathcal{P}_s = \{\mathbf{p}_i^s\}_{i=1}^M, \quad \mathbf{p}_i^s = (x_i^s, y_i^s, z_i^s)^\top \in \mathbb{R}^3. \tag{5}$$

Similarly, the target point cloud, $\mathcal{P}_t$, consists of N points $\mathbf{p}_j^t = (x_j^t, y_j^t, z_j^t)^\top$, as given in:

$$\mathcal{P}_t = \{\mathbf{p}_j^t\}_{j=1}^N, \quad \mathbf{p}_j^t = (x_j^t, y_j^t, z_j^t)^\top \in \mathbb{R}^3. \tag{6}$$

For both **rigid and generalization phantom-based adaptation**, a four-step registration process was applied. First, statistical outliers were removed, and voxel-grid downsampling was performed to reduce computational complexity. Second, the centers of the point clouds were aligned. For rigid transformations, Open3D's RANSAC registration method, based on feature matching, was performed. Fast Point Feature Histograms (FPFH) were used as visual features for tracking [20]. In contrast, for **generalization** requiring non-rigid transformations, Coherent Point Drift (CPD) was used [21]. In both cases, Iterative Closest Point (ICP) was applied as a final step to find the correspondences [22]. For the **non-rigid arm motion tracking**, explicit point cloud registration was unnecessary, as the four tracked anatomical landmarks inherently defined surface correspondences. These landmarks provided a direct measure of surface displacement, eliminating the need for additional registration steps.

D. LTE Trajectory Transportation

To adapt the initial trajectory, $\mathcal{T}$, LTE modifies the initial set of positional waypoints, $\mathcal{T}_p$, based on estimated surface displacements, to generate an adapted set of waypoints, $\mathcal{T}_p^*$, preserving smoothness while respecting imposed positional constraints [17].

For this work, the LTE implementation of [23] was used. In LTE, we represent $\mathcal{T}_p$ as an undirected graph $\mathcal{G}$, with vertices $\mathcal{V}$ and edges $\mathcal{E}$ defining the connectivity between the vertices. Each point $\mathbf{p}_i$ in $\mathcal{T}_p$ corresponds to the vertex v_i in $\mathcal{V}$. The set of neighboring points of a vertex v_i, is given by

$$\mathcal{N}(v_i) = \{v_j \in \mathcal{V} \mid (v_i, v_j) \in \mathcal{E}\} \tag{7}$$

To capture the local properties of the trajectory, the waypoints in Cartesian space $\mathbf{p}_i$ were converted into Laplacian coordinates $\Delta = L\mathbf{T}_p$, where $\mathbf{T}_p$ is the matrix representation of the waypoint set $\mathcal{T}_p$. Δ is the concatenation of all δ_i which, for the vertex v_i, is given as:

$$\delta_i = \sum_{v_j \in \mathcal{N}(v_i)} w_{ij} \left(\frac{\mathbf{p}_i - \mathbf{p}_j}{\sum_{v_j \in \mathcal{N}(v_i)} w_{ij}} \right), \tag{8}$$

where w_{ij} is the edge weight between vertices v_i and v_j, set to 1 in our approach. To obtain Δ, the Laplacian matrix $L \in \mathbb{R}^{n \times n}$ is used:

$$L_{ij} = \begin{cases} 1, & \text{if } i = j, \\ -\frac{w_{ij}}{\sum_{v_k \in \mathcal{N}(v_i)} w_{ik}}, & \text{if } v_j \in \mathcal{N}(v_i), \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

All Laplacian coordinates are then given by $\Delta = L\mathbf{T}_p$, where $\mathbf{T}_p$ represents the matrix form of the waypoint set $\mathcal{T}_p$. Once the Laplacian coordinates were computed, LTE found the adapted trajectory by solving a least-squares problem.

The matrix representation of the adapted set of waypoints was found by solving the optimization problem:

$$\mathbf{T}_p^* = \arg\min_{\mathbf{T}_p^*} \|L\mathbf{T}_p^* - \Delta\|^2 + \lambda \|\bar{P}\mathbf{T}_p^* - \bar{\mathbf{C}}\|^2, \tag{10}$$

where $\bar{P}$ is a diagonal binary matrix where each entry $\bar{P}_{kk}$ is set to 1 if waypoint $\mathbf{p}_k^*$ in $\mathcal{T}_p^*$ is constrained, and 0 otherwise. $\bar{C}$ contains the target positions for the constrained waypoints in $\mathcal{T}_p^*$. λ is a regularization weight, controlling the trade-off between preserving local shape properties and adhering to the new displacement constraints.

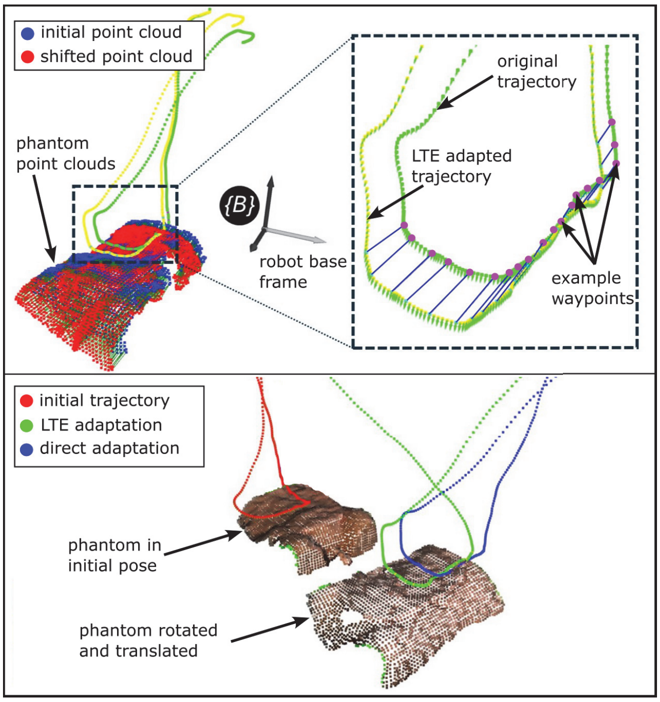

Fig. 3: **Top:** LTE adaptation based on point cloud registration. The original (green) and adapted (yellow) trajectories, derived from the registration between the source and target point clouds of the abdominal phantom, are shown. Inset box shows a visualization of $\bar{P}$ and $\bar{C}$, where magenta points indicate selected waypoints ($\bar{P}$) and blue lines represent the displacement constraints ($\bar{C}$). **Bottom:** example LTE adaptation in which a phantom has been both translated and rotated. The direct adaptation without LTE (blue) is also shown for comparison

The dynamic definition of the displacement constraints, $\bar{C}$ and the selected waypoints $\bar{P}$, based on point cloud input is key to applying LTE (Fig. 3). The selection of which waypoints from the original trajectory, $\mathbf{p}_k \in \mathcal{T}_p$, influence the adaptation process is given by:

$$\mathcal{K} = \{k \mid \mathbf{p}_k \in \mathcal{T}_p\}. \tag{11}$$

For **rigid and generalization phantom-based adaptation**, where point clouds were dense, a nearest-neighbor search using a KDTree was used to identify waypoints in $\mathcal{T}_p$ that were within a predefined distance threshold of 0.03 m from $\mathcal{P}_s$, indicating physical contact with the surface. Not all contact points were used as constraints, as this would overly restrict the adaptation and leave little room to preserve local properties during LTE optimization. Instead, every z^{th} contact point was selected, where z was a tunable parameter. A smaller z enforced stricter positional tracking by adding more constraints, while a larger z provided greater flexibility in preserving the original trajectory's shape. For **non-rigid arm motion tracking**, where the representation consisted of sparse keypoints, the Hungarian algorithm was applied to optimally assign four waypoints from $\mathcal{T}_p$ to the four keypoints in $\mathcal{P}_s$ [24]. The adaptation was then constrained such that the corresponding waypoints in $\mathcal{T}_p^*$ aligned with these key anatomical locations.

Each selected waypoint $\mathbf{p}_k \in \mathcal{T}_p$ defined a corresponding constraint in the adapted trajectory $\mathcal{T}_p^*$. These constraints were indexed in the set $\mathcal{K}$ and for each $k \in \mathcal{K}$, the corresponding entry in $\bar{P}$ was set to 1:

$$\bar{P}_{kk} = 1, \quad \forall k \in \mathcal{K}. \tag{12}$$

The displacement for each constraint was computed based on the difference between corresponding points in $\mathcal{P}_s$ and $\mathcal{P}_t$:

$$\mathbf{d}_k = \mathbf{p}_m^t - \mathbf{p}_\ell^s, \quad \forall k \in \mathcal{K}, \tag{13}$$

where $\mathbf{p}_\ell^s \in \mathcal{P}_s$ and $\mathbf{p}_m^t \in \mathcal{P}_t$ denote corresponding points in the source and target point clouds, respectively. The selected constraints were then applied to the adapted trajectory $\mathcal{T}_p^*$, modifying the corresponding waypoints as

$$\bar{\mathbf{C}}_k = \mathbf{p}_k + \mathbf{d}_k, \quad \forall k \in \mathcal{K}. \tag{14}$$

Additionally, to maintain trajectory continuity and ensure smooth motion transitions, the robot's current position was added as a constraint. This ensured that the current robot position was always a part of $\mathcal{T}_p^*$. Fig. 3 shows a transportation following from the registration procedure.

To ensure optimal ultrasound propagation, the probe was oriented perpendicular to the scanning surface [2]. Surface normals of the target point cloud, $\mathcal{P}_t$, were estimated using the Open3D library. For each waypoint in the adapted trajectory, $\mathcal{T}_p^*$, the nearest point in $\mathcal{P}_t$ within 0.03 m was identified using a KDTree, and its surface normal defined the probe's orientation $\mathbf{q}_i$, directed toward the surface. If no nearby point was found, the original orientation from $\mathcal{T}$ was retained. Combining LTE-modified positions with these normals yielded the final trajectory $\mathcal{T}^*$.

III. EXPERIMENTAL VALIDATION

The proposed RUSS was validated through a series of experiments on an **abdominal phantom** for adapting the trajectory to rigid transformations and a **human arm** for adapting the trajectory to non-rigid transformations. The experimental setup consists of a KUKA LBR iiwa 14 R820 (KUKA AG, Augsburg, Germany) robotic manipulator, which is controlled using the KUKA Fast Robot Interface (FRI), equipped with a 3D-printed mock ultrasound probe at its end-effector. An Intel RealSense D455 depth camera (Intel Corporation, Santa Clara, CA, USA), mounted above the robot workspace, provided real-time RGB-D data. For clinical practice wave propagation through tissue requires the robot to be in physical contact with the patient. In order to ensure a safe contact, the robot is controlled at the joint-torque level using a Cartesian Impedance Controller. The interaction force/torque in Cartesian space was defined as:

$$\boldsymbol{f} = \boldsymbol{K}(\boldsymbol{x}_a - \boldsymbol{x}_d) + \boldsymbol{D}\dot{\boldsymbol{x}}_a + \boldsymbol{S}_f \boldsymbol{f}_d, \tag{15}$$

where $\boldsymbol{K}$ and $\boldsymbol{D}$ represent the stiffness and damping matrices $\in \mathbb{R}^{6\times 6}$, $\boldsymbol{x}_a \in \mathbb{R}^6$ and $\boldsymbol{x}_d \in \mathbb{R}^6$ being the actual and desired robot end-effector pose in Cartesian space, $\boldsymbol{S}_f \in \mathbb{R}^{6\times 6}$ a diagonal matrix used for selecting the axis in which the desired force is applied and $\boldsymbol{f}_d \in \mathbb{R}^6$ the desired force, both in Cartesian space. We used a critically damped design $\boldsymbol{D} = 2\zeta\sqrt{\boldsymbol{K}}$, with damping ratio $\zeta = 0.7$. The robot and tool were gravity-compensated. Translational stiffness values of the stiffness matrix $\boldsymbol{K}$ were set to 300.0 N/m and 10.0 Nm/rad for rotation. A desired force $\boldsymbol{f}_d$ was set to 2.0 N along the z-axis, to ensure consistent contact but only when the probe was within 0.1 m of the surface's z-height–extracted from depth images–to prevent exertion of force when no contact was needed. The reference pose $\boldsymbol{x}_d$ was determined by the output of the proposed approach (Fig. 2). For LTE, the regularization weight λ was set to 1 and the anchor waypoint subsampling factor z was set to 6, in order to balance adhering to positional constraints and preserving local properties

A. Validation Experiments

The performance of the proposed approach was assessed across three distinct experiments, designed to evaluate its trajectory adaptation capabilities for increasing complexity. Anatomical phantoms were used for the first two experiments, followed by a human-arm actively moving to represent the more complex non-rigid case of MSK-US. For all experiments, a static replay served as the baseline for comparison. Position and acceleration data from the robot were recorded to analyze tracking accuracy and motion smoothness.

1) Rigid Phantom Motion: To evaluate adaptation to rigid transformations, an abdominal phantom was manually moved after an initial trajectory demonstration. Five trials were conducted for each motion type (translation, rotation, tilt).

2) Generalization Across Phantom Morphologies: To evaluate the ability to transfer a trajectory demonstrated on one phantom to another with a different anatomical shape, the male phantom was replaced with the female phantom after the initial demonstration. Again five trials were conducted.

3) Non-Rigid Arm Motion: To simulate the dynamic patient movement, reflecting MSK-US use cases, experiments were conducted with a human participant actively flexing and extending at the elbow joint which resting the forearm on a table. To establish a ground-truth trajectory, a predefined reference line was set via blue physiotherapy tape on the human arm (Fig. 4). This reference, represented the ideal scanning path that the adapted trajectory should follow after accounting for surface movement. The initial demonstrated trajectory was obtained by asking an operator to manually follow this blue line on a static arm, ensuring that the intended scan path aligned with the reference. The position of the blue line was extracted in the robot's base frame, enabling a direct comparison with the transported trajectory.

B. Evaluation Metrics

To comprehensively assess the performance of the proposed Robotic Ultrasound System (RUSS), several quantitative metrics were employed, focusing on adaptation success, accuracy, robot motion quality, computational efficiency, and the quality of intermediate steps like point cloud registration.

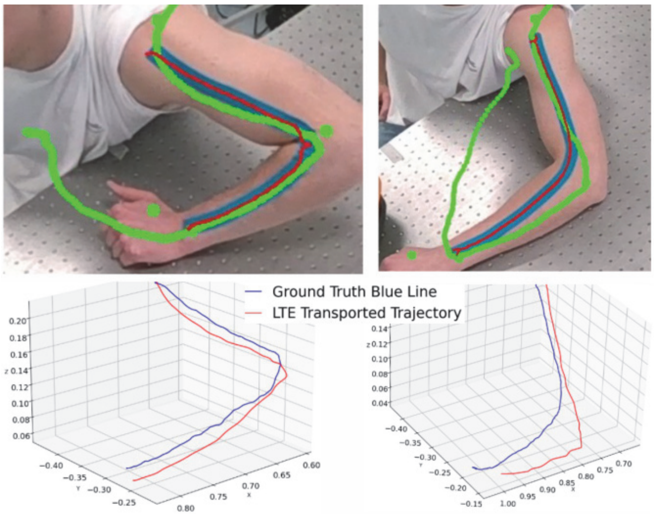

Fig. 4: Illustration of trajectory adaptation accuracy evaluation. The blue line represents the ground truth reference drawn on the arm, while the red line corresponds to the adapted trajectory transported using Laplacian Trajectory Editing (LTE). The deviation between the two lines quantifies adaptation accuracy.

1) Adaptation Success Rate: For each experimental condition, the outcome of every trial was categorized as a success or failure. A failure was defined as any instance where the robotic manipulator could not complete the intended adapted scan trajectory, for example as result of violating robot safety limits.

2) Trajectory Adaptation Accuracy (Non-Rigid Arm Motion): To evaluate adaptation accuracy, in the non-static experiment the Root Mean Squared Error (RMSE) is reported, which quantifies the average deviation between corresponding points on the adapted trajectory and the (blue) reference line.

3) Robot Motion Quality: Two key aspects of the robot's motion were evaluated:

Positional Tracking Error: This metric measures the deviation between the robot's actual executed end-effector position and the reference position from the adapted trajectory at every timestep. It is reported as the mean and maximum Euclidean distance error throughout a trial.

Motion Smoothness (Jerk): Jerk, was used to quantify the smoothness of the robot's motion. Higher jerk values indicate less smooth, potentially more abrupt movements. Both mean and maximum jerk are reported.

4) Computational Performance: To assess the computational feasibility, the execution times for the main components of the adaptation pipeline (Fig. 2) were measured. These include: (b) image processing, (c) displacement estimation, and (d) LTE trajectory adaptation. Average computation times per trial for each component are reported.

5) Point Cloud Registration Accuracy (for Phantom Experiments): Since surface displacement estimation in phantom-based scenarios relies on point cloud registration, its accuracy was assessed. For successful registration instances, two metrics were used:

Mean Point-to-Surface Distance: The average distance from points in the aligned source point cloud to the nearest surface point in the target point cloud. *Overlap Ratio:* The percentage of points in the source cloud that have a corresponding point in the target cloud within a certain proximity after registration.

IV. RESULTS

A. Baseline Performance

When replaying a recorded demonstration from the VSI protocol five times (Fig. 5), the trajectory tracking mean and maximum errors were 0.018 ± 0.001 m and 0.079 ± 0.002 m, respectively. The total mean jerk recorded in the motion was 4.03 ± 0.1 m/s^3, and the maximum jerk was 19.8 ± 1.4 m/s^3.

B. Adaptation Success and Accuracy

For all experimental conditions, the success and failure outcomes are presented in Table I. The system successfully adapted in most trials, with occasional failures occurring in the rigid and non-rigid adaptation scenarios. The initial error between the operator's demonstration and the blue line was averaged over 30 frames, resulting in a mean deviation of 0.021 ± 0.002 m. Over 54 LTE adaptations, the RMSE between the blue line and the adapted trajectory was computed (Fig. 6). The mean RMSE across all adaptations was 0.026 ± 0.012 m.

TABLE I: Success (S) and Failure (F) per Trial

Trial #	**Rigid Movements**			**Generalization**	**Non-Rigid**
	T 20 cm	**R 45°**	**Tilt 30°**		
1	S	S	F	S	S
2	S	F	S	S	S
3	S	F	S	S	S
4	S	S	F	S	F
5	S	S	S	S	S

C. Trajectory Tracking and Smoothness Analysis

Across scenarios, both mean and maximum positional tracking errors remained low compared to the static demonstration replay (Fig. 8). However, errors increased in the translational rigid scenario, generalization, and non-rigid adaptation, with non-rigid adaptations showing the highest spread. Positional tracking data for one trial is shown as an example in Fig. 7. Notably, the maximum positional tracking errors were highest in the non-rigid case. Mean jerk remained low and consistent in the static, generalization, and tilt scenarios (Fig. 8). However, translational and rotational rigid adaptations showed higher variability, with occasional spikes. Non-rigid adaptations showed the highest mean jerk values, exceeding 20 m/s^3 in some trials.

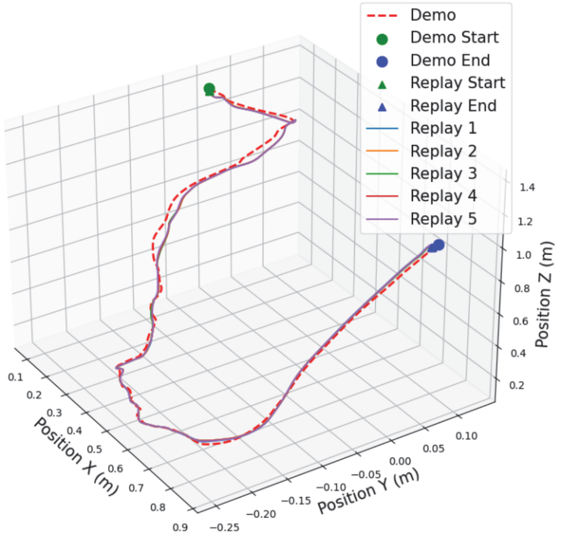

Fig. 5: The red-dotted line represents a VSI protocol demonstration, while the colored lines show recorded poses during replay. Due to close alignment, individual trajectories are difficult to distinguish.

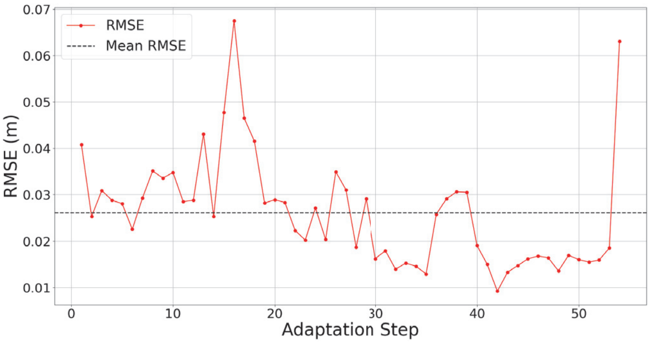

Fig. 6: RMSE between the adapted trajectory and the moving blue line reference over 54 adaptation steps. The mean RMSE across all adaptations is 0.026 ± 0.012 m.

D. Computational Performance

The reported execution times can be seen in Table II. Image processing times remained similar in the rigid and generalization scenarios, at 942 ms and 1090 ms, respectively, while the non-rigid scenario had a lower processing time of 55.2 ms. The displacement estimation time showed significant variation across scenarios, with the generalization case taking considerably longer than rigid and non-rigid adaptations. Lastly, LTE-based trajectory adaptation required 240 ms in the rigid scenario, 463 ms in the generalization scenario, and 415 ms in the non-rigid scenario.

TABLE II: Computation Times (ms)

Component	**Rigid**	**Generalization**	**Non-Rigid**
Image processing	942	1090	55.2
Displacement estimation	1090	40500	0.380
LTE trajectory adaptation	240	463	415

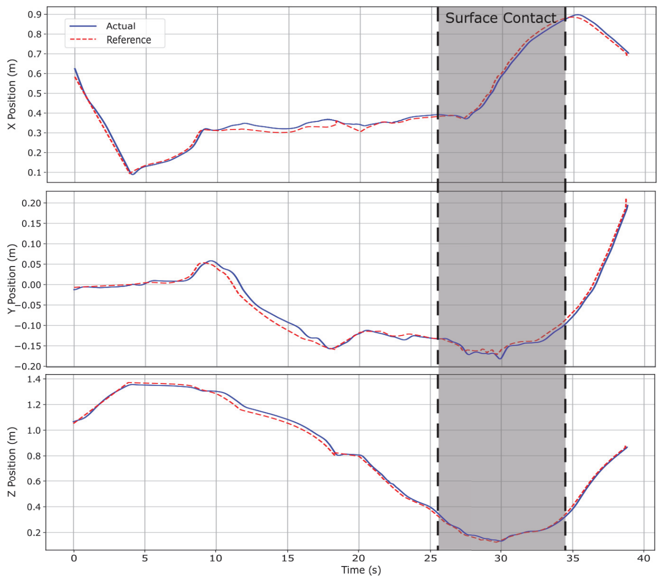

Fig. 7: Tracking errors across the X, Y, and Z dimensions. Each plot shows the deviation between the executed trajectory and the adapted reference trajectory over time. These results correspond to a 'Translation 20 cm' trial.

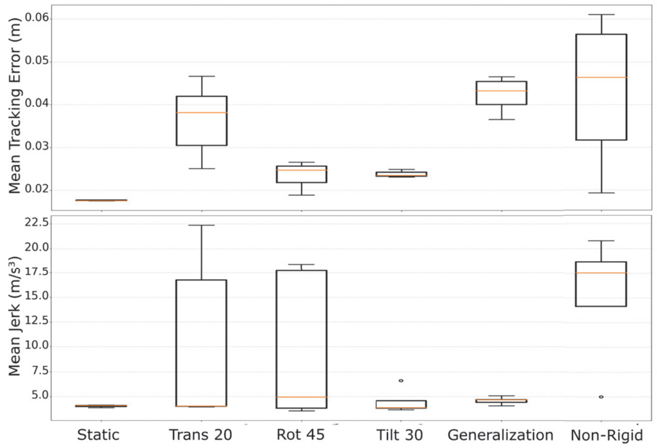

Fig. 8: Mean positional tracking errors (top) and mean jerk (bottom) across experimental conditions, with five trials per scenario.

E. *Point Cloud Registration Accuracy*

Across all tested scenarios, the mean point-to-surface distance ranged from 10.0 ± 0.5 mm to 14.7 ± 0.3 mm. Maximum distances varied more significantly, with the highest deviations occurring in the Tilt 30°scenario (97.3 ± 27.2 mm). Overlap ratios remained consistently high, exceeding 90% across all conditions, with the generalization scenario achieving full overlap (100%).

V. DISCUSSION

In this study we developed and validated a proof-of-concept RUSS capable of real-time trajectory adaptation to both rigid and non-rigid surface movements using our proposed framework leveraging LTE. Overall, this system proved successful in adapting demonstrated trajectories based on RGB-D camera streams, a step towards more autonomous and clinically adaptable RUSS. With regard to trajectory smoothness, maximum jerk remained low with larger jerks occurring only in response to key instances of the robot reaching workspace limits or singularities. Within the robot's feasible workspace, mean jerk closely aligns with the baseline value of 4.03 ± 0.1 m/s^3, demonstrating smooth trajectories across trials.

TABLE III: Point Cloud Registration Accuracy Metrics

Metric	**T20 cm**	**R 45°**	**Tilt 30°**	**Generalization**
Point-to-Surface Mean Distance (mm ± SD)	10.0 ± 0.5	12.9 ± 0.3	14.7 ± 0.3	11.2 ± 0.1
Point-to-Surface Max Distance (mm ± SD)	51.0 ± 8.1	52.7 ± 3.2	97.3 ± 27.2	69.0 ± 9.8
Overlap Ratio (% ± SD)	98.1 ± 2.5	90.7 ± 2.9	97.6 ± 0.1	100.0 ± 0.0

For non-rigid trials, the error between the blue reference line and the adapted trajectory remained within 0.035 and 0.015 m for most adaptations (Fig. 6). The mean RMSE after adaptation (0.026 ± 0.012 m) was slightly higher than the initial deviation in the operator-drawn line (0.021 ± 0.002 m), indicating effective adaptations that are able to maintain operator-defined trajectories on moving targets with minor deviations. However, some trials showed errors exceeding 0.06 m, while others had lower RMSE than the initial demonstration, highlighting variability in adaptation accuracy influenced by arm movement (Fig. 6). Much of this variability can be attributed to an inherent property of LTE optimization when using sparse keypoints as LTE seeks to minimize changes in local Laplacian coordinates and distribute these adjustments across all unconstrained points. This can result in excessively curved trajectories between keypoints as the arm bends close to 45°. The use of additional keypoints along the target trajectory would likely improve this result substantially.

An analysis of the trial failures reveals that they stemmed not from the adaptation algorithm itself, but from the robot's physical limitations. Since our framework adapts trajectories purely in Cartesian space, it does not currently account for kinematic constraints. In several instances, an adapted trajectory commanded the robot to move near a singularity, a configuration where small movements in Cartesian space require disproportionately large and fast joint movements. This led to sudden spikes in joint velocity and acceleration, triggering the robot's safety limits and causing the system to disconnect.

During execution of the adapted trajectories, the system demonstrated consistent and low tracking errors, with a mean and maximum tracking error of 0.018 ± 0.001 m and 0.079 ± 0.002 m during static replay (Fig. 5). Across adaptation scenarios, mean tracking errors remained low but were slightly higher in trials where the probe made contact with the surface, likely due to force interactions (Fig. 8). In the non-rigid case, mean tracking errors were further influenced by high maximum errors, resulting from robot workspace limitations. Future work will incorporate these workspace bounds to impose feasible limits on trajectories.

VI. LIMITATIONS AND FUTURE WORK

While this work successfully demonstrates a proof-of-concept framework for applying LTE to robotic ultrasound to achieve robust dynamic trajectory adaptation to both rigid and non-rigid patient motion, several limitations warrant consideration and guide future work.

The primary limitations concern the system's **readiness for clinical application**. The focus on trajectory adaptation, rather than US imaging itself, limits the system's immediate readiness for clinical validation. Since the system only adapts to external surface motion and there is no imaging feedback, there is no guarantee that the internal imaging plane remains correct, as tissue deforms differently from the skin surface. Similarly, while basic force and orientation control were implemented, they were not extensively optimized. These factors are important to achieve high-quality imaging. Finally, the non-rigid validation was conducted with a single human subject, which is insufficient to establish statistical validity and ensure the system generalizes across wider inter-subject variability. Future work should therefore prioritize the integration of an ultrasound probe, the implementation of more advanced force-aware controllers (like in [25], [26]), and validation across a larger and more diverse human cohort.

From a **technical point of view**, the performance reveals two areas for improvement. First, the system adapts trajectories without considering the robot's kinematic constraints. This occasionally led to unreachable commanded poses or approaching a kinematic singularity, causing trial failures and spikes in joint velocities and acceleration, highlighting the need to incorporate workspace-aware and singularity-avoidance constraints directly into the trajectory planning process. Second, the real-time capability of the overall pipeline is currently limited by its pre-processing stages. While LTE adaptation is consistently fast, the image processing and point cloud registration created a significant bottleneck, which future work can address for example by exploring GPU-accelerated methods.

REFERENCES

[1] C. Martinoli, "Musculoskeletal ultrasound: technical guidelines," *Insights into imaging*, vol. 1, no. 3, p. 99, 2010.

[2] B. Ihnatsenka and A. P. Boezaart, "Ultrasound: Basic understanding and learning the language," *International journal of shoulder surgery*, vol. 4, no. 3, p. 55, 2010.

[3] D. Won, J. Walker, R. Horowitz, S. Bharadwaj, E. Carlton, and H. Gabriel, "Sound the alarm: the sonographer shortage is echoing across healthcare," *Journal of Ultrasound in Medicine*, vol. 43, no. 7, pp. 1289–1301, 2024.

[4] O. Pallotta and A. Roberts, "Musculoskeletal pain and injury in sonographers, causes and solutions," *Sonography*, vol. 4, no. 1, pp. 5–12, 2017.

[5] G. Ning, H. Liang, X. Zhang, and H. Liao, "Autonomous robotic ultrasound vascular imaging system with decoupled control strategy for external-vision-free environments," *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 11, pp. 3166–3177, 2023.

[6] Y. Bi, C. Qian, Z. Zhang, N. Navab, and Z. Jiang, "Autonomous path planning for intercostal robotic ultrasound imaging using reinforcement learning," *arXiv preprint arXiv:2404.09927*, 2024.

[7] M. Chen, Y. Huang, J. Chen, T. Zhou, J. Chen, and H. Liu, "Fully robotized 3d ultrasound image acquisition for artery," in *2023 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2023, pp. 2690–2696.

[8] Z. Jiang, Y. Zhou, D. Cao, and N. Navab, "Defcor-net: Physics-aware ultrasound deformation correction," *Medical Image Analysis*, vol. 90, p. 102923, 2023.

[9] Z. Jiang, Y. Gao, L. Xie, and N. Navab, "Towards autonomous atlas-based ultrasound acquisitions in presence of articulated motion," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7423–7430, 2022.

[10] 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 Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–11, 2022.

[11] 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, 2025.

[12] S. Schaal, "Dynamic movement primitives-a framework for motor control in humans and humanoid robotics," in *Adaptive motion of animals and machines*. Springer, 2006, pp. 261–280.

[13] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," *Advances in neural information processing systems*, vol. 26, 2013.

[14] Y. Huang, L. Rozo, J. Silvério, and D. G. Caldwell, "Kernelized movement primitives," *The International Journal of Robotics Research*, vol. 38, no. 7, pp. 833–852, 2019.

[15] S. Calinon, "A tutorial on task-parameterized movement learning and retrieval," *Intelligent service robotics*, vol. 9, pp. 1–29, 2016.

[16] S. Arora and P. Doshi, "A survey of inverse reinforcement learning: Challenges, methods and progress," *Artificial Intelligence*, vol. 297, p. 103500, 2021.

[17] T. Nierhoff, S. Hirche, and Y. Nakamura, "Spatial adaption of robot trajectories based on laplacian trajectory editing," *Autonomous Robots*, vol. 40, pp. 159–173, 2016.

[18] T. J. Marini, D. C. Oppenheimer, T. M. Baran, D. J. Rubens, M. Toscano, K. Drennan, B. Garra, F. R. Miele, G. Garra, S. J. Noone, L. Tamayo, C. Carlotto, L. Trujillo, E. Waks, K. Garra, M. S. Egoavil, J. Berrospi, and B. Castaneda, "New ultrasound telediagnostic system for low-resource areas: Pilot results from peru," *Journal of Ultrasound in Medicine*, vol. 40, no. 3, pp. 583–595, 2021.

[19] C. Lugaresi, J. Tang, H. Nash, C. McClanahan, E. Uboweja, M. Hays, F. Zhang, C.-L. Chang, M. G. Yong, J. Lee *et al.*, "Mediapipe: A framework for building perception pipelines," *arXiv preprint arXiv:1906.08172*, 2019.

[20] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 3212–3217.

[21] A. Myronenko and X. Song, "Point set registration: Coherent point drift," *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, no. 12, pp. 2262–2275, 2010.

[22] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image and vision computing*, vol. 10, no. 3, pp. 145–155, 1992.

[23] G. Franzese, R. Prakash, and J. Kober, "Generalization of task parameterized dynamical systems using gaussian process transportation," *arXiv preprint arXiv:2404.13458*, 2024.

[24] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[25] J. Fu, I. Burzo, E. Iovene, J. Zhao, G. Ferrigno, and E. De Momi, "Optimization-Based Variable Impedance Control of Robotic Manipulator for Medical Contact Tasks," *IEEE Transactions on Instrumentation and Measurement*, vol. 73, pp. 1–8, Mar. 2024.

[26] J. Jiang, J. Luo, H. Wang, X. Tang, F. Nian, and L. Qi, "Force Tracking Control Method for Robotic Ultrasound Scanning System under Soft Uncertain Environment," *Actuators*, vol. 13, no. 2, p. 62, Feb. 2024.