

Finding an Initial Probe Pose in Teleoperated Robotic Echocardiography via 2D LiDAR-Based 3D Reconstruction

Mariadas Capsran Roshan^{1*}, Edgar M Hidalgo¹, Mats Isaksson¹, Michelle Dunn¹, and Jagannatha Charjee Pyaraka¹

Abstract—Echocardiography is a key imaging modality for cardiac assessment but remains highly operator-dependent, and access to trained sonographers is limited in underserved settings. Teleoperated robotic echocardiography has been proposed as a solution. However, clinical studies report longer examination times than manual procedures, increasing diagnostic delays and operator workload. Automating non-expert tasks, such as automatically moving the probe to an ideal starting pose, offers a pathway to reduce this burden. Prior vision- and depth-based approaches to estimate an initial probe pose are sensitive to lighting, texture, and anatomical variability. We propose a robot-mounted 2D LiDAR-based approach that reconstructs the chest surface in 3D and estimates the initial probe pose automatically. To the best of our knowledge, this is the first reconstruction of robot-mounted 2D LiDAR used for 3D reconstruction of a human body surface. Through plane-based extrinsic calibration, the transformation between the LiDAR and robot base frames was estimated with an overall root mean square (RMS) residual of 1.82 mm and rotational uncertainty below 0.2° . The chest front surface, reconstructed from two linear LiDAR sweeps, was aligned with scale-augmented rigid registration to identify an initial probe pose. Mannequin-based study assessing reconstruction accuracy showed mean surface errors of 2.78 ± 0.21 mm. Human trials (N=5) evaluating the proposed approach found probe initial points typically 20–30 mm from the clinically defined initial point, while the variation across repeated trials on the same subject was less than 4 mm.

I. INTRODUCTION

Echocardiography is the application of ultrasound to visualize the heart, providing real-time assessment of cardiac anatomy and function. It is a primary diagnostic tool for detecting heart failure, evaluating valvular disorders, identifying congenital abnormalities, and guiding decision-making in emergency and critical-care contexts [1], [2]. The accuracy of echocardiography is highly operator-dependent, and access to trained sonographers is limited in regional and underserved areas [3]. Delays in cardiac assessment worsen outcomes, with one study [4] reporting that in acute heart attack patients each additional 30-minute emergency medical transport delay increased the 30-day mortality risk by 20%.

*Corresponding author

¹School of Engineering, Swinburne University of Technology, Victoria, Australia. Emails: mroshan@swin.edu.au, ehidalgoflorez@swin.edu.au, misaksson@swin.edu.au, jdunn@swin.edu.au, jagannathacharjee@swin.edu.au

This research was supported in part by the ARC Industrial Transformation Training Centre for Collaborative Robotics in Advanced Manufacturing (Project ID: IC200100001) and the Australia's Economic Accelerator (AEA) Ignite program (Grant No. IG240100320).

Teleoperated robotic echocardiography [5] has been proposed as a promising solution to overcome this limitation. Such a solution includes a robot equipped with an ultrasound probe is manipulated remotely by an echocardiographer, with real-time feedback enabling accurate image acquisition across distance. A clinical study [6] demonstrated that this approach can reduce diagnostic delays by nearly three months compared to conventional referral pathways that often require patients to travel to specialized facilities. However, as summarized in Table I, multiple clinical studies report longer durations for teleoperated robot echocardiography compared to manual examinations.

Previous research [10], [11] has proposed that not all stages of echocardiography require expert input and that repetitive preparatory steps can be automated to reduce examination time. One such step is the initial placement of the probe on the chest above the heart, which currently requires manual maneuvering from a neutral position. This process begins with localizing the cardiac region of interest (ROI) on the chest surface, followed by identifying a suitable position and orientation of the probe at that ROI, and finally moving the probe into this configuration. Automating this sequence can reduce procedure time and operator workload.

Several studies have investigated vision-based approaches to automate initial probe pose estimation within teleoperated and other robot-assisted ultrasound imaging systems, spanning applications in both echocardiography and other anatomical targets. Mustafa et al. [12] targeted the epigastric region by processing a frontal RGB photograph of the torso, automatically detecting anatomical landmarks such as the umbilicus and nipples, and then applying anthropometric distance ratios to triangulate an approximate target position for probe placement. In contrast, Ning et al. [13] avoided landmarks through a reinforcement learning framework. Their approach used a state-representation network to embed ultrasound and contact-force context into the RGB frame, while a pre-trained policy consumed these cues to guide

TABLE I: Comparison of examination durations between teleoperated and manual echocardiography.

Study	Teleoperated Duration	Manual Duration	Relative Increase
Liu et al. [7]	24 min 36 s	16 min 15 s	$\sim 1.5\times$
Wright et al. [8]	32.7 ± 8.9 min	5.9 ± 0.92 min	$\sim 5.5\times$
Solvin et al. [9]	26.4 min	Standard bedside	$\sim 1.9\times$

the probe toward the ROI. Both approaches highlight how RGB-only cues can initialize probe placement, though they remain sensitive to lighting and appearance variations [14]. Additionally, these approaches have not been demonstrated for cardiac imaging, where anatomical variability and motion add further challenges [15].

To address the limitations of RGB-based approaches, particularly their dependence on visual texture, some works employ depth sensors to capture geometric surface cues. Graumann et al. [16] and Kaminski et al. [17] registered pre-segmented anatomical models (e.g., cardiac or thyroid regions) onto live patient point clouds acquired with RGB-D cameras, using geometry-based alignment techniques to estimate the ROI and probe pose. In contrast, Tan et al. [18] fused multi-view structured-light point clouds, used key-points to segment the ROI, and estimated the probe’s initial position from surface points with its orientation aligned to local surface normals. While these approaches improve robustness over RGB-only approaches, they remain constrained by depth sensing, which is vulnerable to ambient light [19], limited range [20], and the need for multi-angle views or wide baselines [21] that are impractical in often space-limited teleoperated settings.

To overcome the limitations of vision-based approaches, we propose a two-dimensional light detection and ranging (2D LiDAR)-based initial probe pose estimation approach. A robot-mounted 2D LiDAR performs two linear sweeps across the patient’s chest front surface, which, following plane-based calibration, yield a reconstructed 3D point cloud. This point cloud is preprocessed to remove noise and non-relevant structures, after which scale-augmented rigid registration identifies the initial probe position. Finally, local plane fitting at that point estimates the surface normal to derive the probe orientation for placement.

The LiDAR-based approach directly addresses several limitations of vision-based approaches. LiDAR measurements are based on time-of-flight ranging, making them invariant to ambient lighting [22]. This avoids degradation from shadows, clothing textures, or background clutter that commonly affect camera-based methods [23]. Moreover, LiDAR sensing encodes spatial geometry through range measurements without capturing photometric or textural information, thereby inherently supporting privacy-preserving patient monitoring [24]. Additionally, robot-mounted LiDAR enables passive multi-view acquisition through platform motion, eliminating the need for wide camera baselines or fixed external sensor placements. To the best of our knowledge, this work presents the first demonstration of reconstructing a human body surface in 3D using a robot-mounted 2D LiDAR scanner.

Although 3D LiDAR systems can capture volumetric data in a single scan, even compact models like the Ouster OS0 Gen2 weigh between 445 g and 930 g [25], with high-resolution units exceeding 8 kg. In contrast, 2D LiDARs such as the RPLIDAR C1 weigh just 110 g [26], minimizing payload and inertia on robotic arms in clinical settings. Angular resolution is also finer in 2D LiDARs (0.25° - 1.0°) [27]

compared to 3D LiDARs’ sparser vertical resolution (0.4° - 2.0°) [28].

In this work, we present the first use of robot-mounted 2D LiDAR for probe pose initialization in tele-echocardiography. Our contributions are: (1) the first demonstration of full 3D human body surface reconstruction using robot-mounted 2D LiDAR; (2) a plane-based extrinsic calibration method, achieving an overall root mean square (RMS) residual of 1.8 mm and rotational uncertainty below 0.2° for robust robot–LiDAR alignment; (3) a scale-augmented rigid registration framework to estimate the initial probe position and orientation; and (4) two quantitative evaluations: one assessing the accuracy of the 3D chest front surface reconstruction against high-resolution scans of male and female mannequins, and another on human participants to evaluate the accuracy and repeatability of the complete approach. An open-source implementation of the complete approach, including LiDAR calibration, reconstruction, and scale-augmented rigid registration, is available at <https://mcr100.github.io/research/lippse/>.

II. METHODS

A. System Overview

While the proposed LiDAR-based initial probe pose estimation approach is system-agnostic, this study uses the configuration in Fig. 1 for implementation and evaluation. We focus on the robot-side components, as the work targets automation of the initial probe positioning step prior to teleoperation. The core platform is a UR5e collaborative robot (cobot), chosen for its compliance and built-in safety features. A Clarius C3 HD3 wireless ultrasound probe was mounted to emulate a realistic clinical setup and to verify that the probe geometry did not interfere with the proposed approach. The probe holder from [29], originally designed for the ultrasound probe, was modified to also accommodate a 2D LiDAR sensor on the robot mount. For this study, we




Fig. 1: System setup with UR5e cobot, probe holder with Clarius ultrasound probe, integrated 2D RPLIDAR C1 and emergency stop button. The bed includes a blue shoulder marker defining the safety plane for restricted robot motion.

used a Slamtec RPLIDAR C1, a lightweight 360° LiDAR with 0.72° angular resolution, chosen to demonstrate that the approach can be realized with low-cost hardware rather than requiring a high-end sensor. The control architecture and data flow were implemented in Robot Operating System (ROS) 1. Safety measures include constraining the robot’s tool center point (TCP) speed and joint limits in software, defining a virtual safety plane to prevent motion toward the participant’s head, and providing an emergency stop button for immediate interruption in case of discomfort.

B. Plane-Based LiDAR–TCP Extrinsic Calibration and Chest Surface 3D Reconstruction

For 3D reconstruction of the chest front surface and subsequent estimation of the initial probe pose, LiDAR points expressed in the LiDAR sensor frame $\{L\}$ must be mapped into the fixed robot base frame $\{B\}$. This requires the extrinsic transformation ${}^C\mathbf{T}_L$ between $\{L\}$ and the TCP frame $\{C\}$, since the robot kinematics provide ${}^B\mathbf{T}_C$. The overall mapping is thus

$$\tilde{\mathbf{x}}_B = {}^B\mathbf{T}_C {}^C\mathbf{T}_L \tilde{\mathbf{p}}_L, \quad (1)$$

where $\tilde{\mathbf{p}}_L$ is a homogeneous LiDAR point in $\{L\}$ and $\tilde{\mathbf{x}}_B$ its corresponding homogeneous coordinate in $\{B\}$.

We adopt a planar constraint strategy for this calibration in which a 600 mm × 900 mm flat board is placed within reach and the robot sweeps the LiDAR across it at multiple poses. The extrinsics are estimated by minimizing point to plane residuals, a common method for range sensor calibration that avoids external metrology and fiducials. Prior works [30], [31], [32] employing this strategy were primarily developed for high-precision laser range finders, which exhibit low noise and stable beam geometry. Unlike multi-plane approaches or joint robot–sensor identification methods such as the SCALAR variants [30], [31], we perform a separate extrinsics-only session with a single flat board. In this approach, plane fitting is applied with pose diversity and degeneracy checks to avoid near-parallel or collinear configurations. Finally, the transform is estimated directly to $\{C\}$ rather than to the flange. This reduces accumulated errors and ensures that the calibration is expressed directly in $\{C\}$ used for probe control, while retaining the simplicity and low setup cost of prior single plane methods [32]. In contrast to laser-based work, our method is adapted for low-cost 2D LiDARs, introducing angular filtering to mitigate scan non-uniformity, robustness measures for noisier point returns, and a lightweight single-board setup that simplifies calibration.

The calibration setup is illustrated in Fig. 2, with the LiDAR mounted rigidly to the probe holder and a flat board, which serves as the calibration target, positioned such that the LiDAR scanning points lie on the same plane. Robot kinematics provide the known transform ${}^B\mathbf{T}_C \in SE(3)$, and the goal of calibration is to estimate the extrinsic ${}^C\mathbf{T}_L \in SE(3)$. Following the convention in Fig. 2, the LiDAR polar angle θ is defined such that 0° lies along the positive x_L




Fig. 2: Calibration setup showing the robot base $\{B\}$, TCP $\{C\}$, and LiDAR $\{L\}$ frames with the flat board as calibration plane. LiDAR polar coordinates are defined with 0° along $+x_L$ and 90° along $+y_L$; only the sector $[135^\circ, 225^\circ]$ is retained to capture board points for calibration.

axis and 90° along the positive y_L axis of frame $\{L\}$. A raw range return (r, θ) is converted as

$$\mathbf{p}^L(r, \theta) = \begin{bmatrix} r \cos \theta \\ r \sin \theta \\ 0 \end{bmatrix}, \quad \tilde{\mathbf{p}}^L = \begin{bmatrix} \mathbf{p}^L \\ 1 \end{bmatrix}, \quad (2)$$

where $\mathbf{p}^L \in \mathbb{R}^3$ is the LiDAR point in frame $\{L\}$ and $\tilde{\mathbf{p}}^L \in \mathbb{R}^4$ its homogeneous form. During calibration, only the sector $\theta \in [135^\circ, 225^\circ]$ is retained to restrict measurements to the calibration board.

The robot is moved to $K = 20$ distinct poses $\{{}^B\mathbf{T}_C^{(k)}\}_{k=1}^K$, providing sufficient positional and orientational diversity to fully determine the extrinsics while avoiding unnecessary data collection overhead. At each robot pose, multiple LiDAR scans are collected within the selected sector. The raw polar measurements are converted to Cartesian points \mathbf{p}^L using Equation 2, after which a RANSAC-based 2D line fitting procedure [33] is applied. This extracts the dominant line corresponding to the intersection of the LiDAR scan plane ($z_L=0$, since the sensor acquires points only in its local x_L - y_L plane) with the calibration board, while rejecting noisier point returns. The inlier set of LiDAR points at pose k after line extraction is denoted $\{\mathbf{p}_i^L\}_{i \in \mathcal{I}_k}$, where i indexes individual LiDAR points within the scan at pose k and \mathcal{I}_k is the set of indices corresponding to LiDAR points classified as inliers after angular filtering and line fitting.

Let the calibration plane in $\{B\}$ be

$$\mathbf{n}_B^\top \mathbf{x}_B + d_B = 0, \quad \|\mathbf{n}_B\| = 1, \quad (3)$$

where $\mathbf{n}_B \in \mathbb{R}^3$ is the unit normal and $d_B \in \mathbb{R}$ the offset. Each LiDAR point is mapped into $\{B\}$ as

$$\tilde{\mathbf{x}}_{B,i}^{(k)} = {}^B\mathbf{T}_C^{(k)} {}^C\mathbf{T}_L \tilde{\mathbf{p}}_i^L, \quad \mathbf{x}_{B,i}^{(k)} = [\tilde{\mathbf{x}}_{B,i}^{(k)}]_{1:3}, \quad (4)$$

where $\tilde{\mathbf{x}}_{B,i}^{(k)} \in \mathbb{R}^4$ is the homogeneous coordinate and $\mathbf{x}_{B,i}^{(k)} \in \mathbb{R}^3$ the Cartesian point in $\{B\}$. The point-to-plane residual is

$$r_i^{(k)} = \mathbf{n}_B^\top \mathbf{x}_{B,i}^{(k)} + d_B, \quad (5)$$

which ideally equals zero for all $i \in \mathcal{I}_k$.

We parameterize ${}^C\mathbf{T}_L$ by a rotation vector $\boldsymbol{\omega} \in \mathbb{R}^3$ and translation $\mathbf{t} \in \mathbb{R}^3$, and represent the plane normal as $\mathbf{n}_B = \mathbf{v}/\|\mathbf{v}\|$ for unconstrained $\mathbf{v} \in \mathbb{R}^3$. The unknown parameter vector is

$$\boldsymbol{\xi} = [\boldsymbol{\omega}^\top \quad \mathbf{t}^\top \quad \mathbf{v}^\top \quad d_B]^\top.$$

The calibration is then formulated as a robust nonlinear least-squares problem

$$\min_{\boldsymbol{\xi}} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \rho(\mathbf{n}_B(\mathbf{v})^\top \mathbf{x}_{B,i}^{(k)}(\boldsymbol{\omega}, \mathbf{t}) + d_B), \quad (6)$$

where $\rho(\cdot)$ is a robust loss (Cauchy). Although the LiDAR provides only planar measurements in its local frame ($z_L = 0$), each point can be embedded into 3D by expressing it in homogeneous form $\tilde{\mathbf{p}}^L$ and transforming it sequentially through ${}^C\mathbf{T}_L$ and the known ${}^B\mathbf{T}_C^{(k)}$. This operation, shown in Equation 4, maps the 2D scan into the robot base frame $\{B\}$ for pose k . When repeated over multiple diverse robot poses, the LiDAR’s scanning plane is effectively swept through 3D space, generating a family of lines that all lie on the calibration board. These intersections provide sufficient geometric constraints to estimate both the extrinsic transform ${}^C\mathbf{T}_L$ and the board parameters (\mathbf{n}_B, d_B) despite the LiDAR being only 2D.

To assess calibration quality, we report both per-pose and overall residual statistics. For each robot pose k , the RMS of the point-to-plane residuals $r_i^{(k)}$ in Equation 5 quantifies local fit accuracy. Aggregating across all poses yields the overall RMS residual, which reflects the global consistency of ${}^C\mathbf{T}_L$. Calibration stability is further characterized by summary statistics across poses, including the mean, extrema, and variability of the per-pose RMS values. In addition, parameter covariance is approximated via the Gauss–Newton Jacobian at the solution of Equation 6, yielding one-sigma uncertainties for the estimated extrinsic translation and rotation parameters. Together, these metrics quantify the accuracy and robustness of the LiDAR–TCP calibration.

As shown in Fig. 3a, the calibration achieved a mean per-pose RMS of 1.77 mm (range 1.17 mm – 3.03 mm), with an overall RMS residual of 1.82 mm. The residual distribution in Fig. 3b was tightly centered near zero, with most inliers lying within ± 3 mm, indicating accuracy without systematic bias. Uncertainty analysis further confirmed that the estimated extrinsic transform had one-sigma translation uncertainties within 1.1 mm and rotation uncertainties within 0.2° , demonstrating sufficient accuracy and stability for downstream probe initialization.

With ${}^C\mathbf{T}_L$ estimated, we reconstruct the chest’s front surface by executing two non-contact linear sweeps with the robot. In our setup, the bed’s longitudinal axis is approximately parallel to the robot-base y -axis (y_B). Accordingly,




Fig. 3: Calibration accuracy results. a) Per-pose RMS residuals across 20 robot poses, showing consistent < 3 mm errors. b) Histogram of all point-to-plane residuals, tightly centered around zero with most inliers within ± 3 mm.

the end-effector traverses two paths approximately parallel to y_B over the left chest’s front surface, each with a slightly different orientation to better capture the left-heart region, while maintaining a safe clearance d_{safe} from the body. During motion, the LiDAR continuously reports polar measurements (r, θ) within the sector $[135^\circ, 225^\circ]$, which are converted to Cartesian points $\mathbf{p}^L(r, \theta)$ through Equation 2 and mapped into the base frame using the time-stamped kinematic transform ${}^B\mathbf{T}_C(t)$ and the calibrated extrinsics ${}^C\mathbf{T}_L$ as in Equation 4. Accumulating samples over the two passes yields an unfiltered point cloud

$$\mathcal{P}_{\text{raw}} = \bigcup_k \bigcup_{i \in \mathcal{I}_k} \mathbf{x}_{B,i}^{(k)} \subset \mathbb{R}^3,$$

which forms a sparse 3D reconstruction that includes the chest’s front surface as well as incidental structures such as the arms and the bed, as shown in Fig. 4a.

C. Scale-augmented Rigid Registration Framework for Initial Probe Pose Estimation

Following 3D reconstruction, we apply the scale-augmented rigid registration framework illustrated in Fig. 5 to automatically estimate the probe’s initial pose. scale-augmented rigid registration has been applied in both general object localization [34] and robotic ultrasound systems with depth camera-based sensing [35] and we adapt this in the proposed approach. As shown in Fig. 4b, the reconstructed 3D point cloud is first preprocessed using voxel downsampling to ensure uniform point density, followed by statistical and radius outlier removal, density-based clustering, and Poisson surface reconstruction with signed-distance trimming. These steps remove floating artifacts and non-body structures while preserving chest geometry. Predefined chest surface templates in the form of point clouds (male and female variants), each annotated with an ideal initial probe point $\mathbf{p}_T \in \mathbb{R}^3$, were employed as references. To accommodate inter-subject variability, each template was scaled about its centroid to generate three variants: the original size, +10%, and -10%. For each variant, surface normals and fast point feature histograms (FPFH) were computed, and fast global registration (FGR) followed by iterative closest point (ICP) refinement was applied to obtain correspondences. If any variant achieved an alignment fitness $\geq 90\%$, that result was selected, resulting in the aligned template and preprocessed point cloud as shown in Fig. 4c. Otherwise, the variant




Fig. 4: Point clouds from mannequin at key stages: (a) raw 3D point cloud; (b) after preprocessing; (c) alignment of template (green) with preprocessed (blue) cloud; (d) Cropped chest surface with the template point (green sphere) and its corresponding nearest-neighbor point on the reconstructed surface (orange sphere); (e) estimated initial probe point with local PCA plane and surface normal vector.

with the highest fitness was used to guide further scaling: either enlarged or shrunk in additional 10% increments. This process was repeated iteratively until the target fitness threshold of 90% was reached.

After alignment, the template point \mathbf{p}_T was mapped to its corresponding point \mathbf{p}_S on the preprocessed point cloud, identified as the nearest neighbor to \mathbf{p}_T , yielding the initial probe position, as shown in Fig. 4d. Maintaining the probe normal to the chest surface offers an anatomically neutral and clinically ideal [36] starting orientation for subsequent teleoperated adjustments. To derive that orientation, principal component analysis (PCA) was applied to the $k = 30$ nearest neighbors of \mathbf{p}_S , yielding the eigenvector associated with the smallest eigenvalue as the local surface normal \mathbf{n}_S , as shown in Fig. 4e. The initialization pose for the probe was thus defined for automated positioning. The template point represents a typical clinical starting location but can be easily modified if alternative probe initialization poses are desired, for example by annotating a different point on the template or selecting another location on the reconstructed chest surface to accommodate patient-specific anatomical variations.

III. EVALUATIONS

A. Mannequin-Based Geometric Validation of LiDAR-based 3D Reconstruction

The purpose of this experiment was to quantify the geometric accuracy of the LiDAR-based reconstruction against high-resolution ground truth (GT) point clouds. The front

chest surfaces of male and female mannequins were manually scanned using a high-resolution Revopoint 3D scanner, yielding the GT point cloud \mathcal{P}_{GT} for each mannequin, as shown in Fig. 6a-b. For each mannequin, 3D reconstructions were obtained using five distinct LiDAR configurations that varied both the sensor’s distance from the chest surface and its orientation to assess robustness, resulting in ten LiDAR-based point clouds, each denoted as \mathcal{P}_L . These point clouds were manually cropped in MeshLab to remove background and non-chest structures.

The cleaned \mathcal{P}_L was then aligned to \mathcal{P}_{GT} using FGR followed by local refinement with ICP. To verify alignment quality prior to geometric error analysis, we computed the final ICP fitness score f_{ICP} and the inlier root-mean-square error e_{ICP} . In addition, we calculated the percentage of LiDAR points within an 8 mm tolerance of the GT point cloud, denoted as c_L . This threshold corresponds to the native spatial resolution of the RPLIDAR C1 (0.72° angular step, projecting to 6-10 mm lateral spacing at the tested stand-off distances) and represents the smallest reliable discrimination of the LiDAR. For reconstruction accuracy, we calculated the RMSE of nearest-neighbor distances from \mathcal{P}_L to \mathcal{P}_{GT} , denoted e_{RMSE} , along with the 95th percentile error e_{95} . Results were aggregated across repeats to compute the mean and standard deviation per condition, and overall statistics across all ten trials were reported.

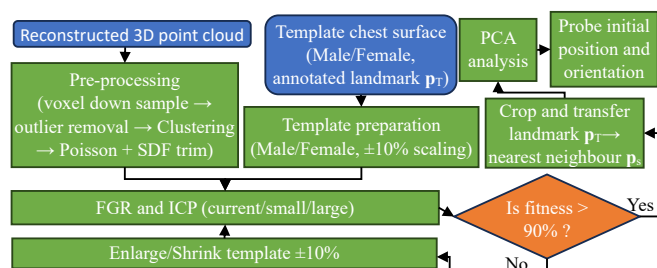


Fig. 5: Framework for initial probe pose estimation via scale-augmented rigid registration with enlarge/shrink alignment loop and PCA-based orientation estimation.




Fig. 6: Cropped chest surfaces of the male (a) and female (b) mannequins. In (c) and (d), the LiDAR point clouds for the male and female mannequins are color-coded by distance to the GT surface (GT shown in yellow): blue < 2 mm, green $[2, 8]$ mm, red > 8 mm.

B. Initial Probe Position Validation on Human Subjects

The objective of this evaluation was to quantify how accurately the complete end-to-end LiDAR-based reconstruction and scale-augmented rigid registration approach identifies an initial probe position.

We recruited $N=5$ healthy participants under Swinburne University of Technology institutional ethics approval (Approval No. 20258952-22425). As shown in Fig. 7a, each participant was positioned supine. No additional breathing restraint beyond standard echocardiography practice was required. Participants remained in a relaxed supine position while breathing normally during the short LiDAR scanning process. A 3D-printed spherical marker of 30 mm diameter was used as a physical ground-truth target. The size was chosen to ensure robust detectability in the LiDAR point cloud while remaining small enough to be comfortably placed on the chest surface without obstructing the scanning region. The marker was positioned at the point of maximal impulse (PMI), corresponding to the cardiac apex beat typically located on the left chest between the fifth and sixth ribs [37], approximately in line with the nipple, as shown in Fig. 7b. This site represents the strongest outward cardiac motion and provides a clinically practical starting point for echocardiography.

For each participant, two linear LiDAR sweeps were performed for 3D reconstruction, and the procedure was repeated across three independent trials. In each trial, before pre-processing and applying the scale-augmented rigid registration framework, a point on the top surface of the spherical marker was manually selected on the raw LiDAR point cloud via a single-point pick, thereby recording the marker location $\mathbf{p}_M^{\text{sphere}} \in \mathbb{R}^3$ in the robot base frame $\{B\}$. This selection was made on the raw point cloud to ensure the marker was preserved, since subsequent preprocessing could otherwise remove it. The complete approach described in Section II was then applied to estimate the initial probe position and orientation for comparison against the ground-truth marker.

To validate the accuracy of the estimated initial point \mathbf{p}_S produced by the proposed approach, the comparison was made against the physical ground-truth marker location.




Fig. 7: Human evaluation setup and outcomes. (a) Participant positioned supine during LiDAR-based chest surface reconstruction. (b) 3D-printed spherical marker placed at the point of maximal impulse (PMI) as a ground-truth reference. (c) Aligned raw (blue) and preprocessed (green) point clouds showing the manually selected marker point $\mathbf{p}_M^{\text{sphere}}$ (red), the estimated initial probe position \mathbf{p}_S (orange), and the projected skin-level ground-truth point \mathbf{p}_M^{\parallel} on the local chest plane (purple).

However, since the spherical marker protrudes above the chest surface, direct comparison could introduce bias. As shown in Fig. 7c, to address this bias, $\mathbf{p}_M^{\text{sphere}}$ was orthogonally projected onto the local chest surface at \mathbf{p}_S as

$$\mathbf{p}_M^{\parallel} = \mathbf{p}_S + (\mathbf{I} - \hat{\mathbf{n}}_S \hat{\mathbf{n}}_S^T) (\mathbf{p}_M^{\text{sphere}} - \mathbf{p}_S), \quad (7)$$

where $\hat{\mathbf{n}}_S$ is the local surface normal at \mathbf{p}_S . This yielded a corrected, skin-level ground-truth location $\mathbf{p}_M^{\parallel} \in \mathbb{R}^3$. Thereafter, the tangential positional error was calculated as

$$e_{\parallel} = \left\| \mathbf{p}_M^{\parallel} - \mathbf{p}_S \right\|_2. \quad (8)$$

Repeatability was assessed by computing the within-subject standard deviation of e_{\parallel} across the three repeated trials for each participant. Alignment quality was further characterized using the final ICP fitness f_{ICP} together with the inlier root-mean-square error e_{ICP} .

IV. RESULTS AND DISCUSSION

In the mannequin evaluation, quantitative results from ten trials are summarized in Table II, separated by male and female mannequins with overall statistics reported for the combined dataset. Alignment quality was consistently high, with f_{ICP} exceeding 0.95 in all cases and e_{ICP} remaining below 3.1 mm. The coverage c_L confirmed that over 93% of LiDAR points were within the 8 mm tolerance of the ground truth, ensuring robust alignment prior to error analysis. In terms of reconstruction accuracy, e_{RMSE} averaged 2.97 mm for male and 2.59 mm for female mannequins, with e_{95} of 5.18 mm and 4.53 mm, respectively. The slightly lower errors for female scans likely reflect smoother chest geometry, which yields fewer local discontinuities during registration. In contrast, the male mannequin exhibits more pronounced anatomical contouring (sternum and rib cage), where minor mismatches in sampling and alignment accumulate as larger local deviations. Across all conditions, the overall reconstruction error remained low, with $e_{\text{RMSE}} = 2.78 \pm 0.21$ mm and $e_{95} = 4.86 \pm 0.36$ mm. These trends are also visible in the color-coded error maps for one pair of alignment in Fig. 6c-d, where LiDAR point clouds for the male and female mannequins are shown relative to the GT surface: blue < 2 mm, green $[2, 8]$ mm, and red > 8 mm. All these results confirm that the proposed method provides reliable and repeatable surface 3D reconstructions suitable for subsequent template-based initial probe pose estimation.

In the human subject evaluation, the e_{\parallel} across all participants and trials ranged from 21–30 mm, with mean values

TABLE II: Quantitative results of mannequin evaluation (mean \pm SD across five trials per mannequin).

Condition	f_{ICP}	e_{ICP} [mm]	c_L [%]	e_{RMSE} [mm]	e_{95} [mm]
Male (N=5)	0.974 ± 0.012	3.01 ± 0.10	95.1 ± 1.3	2.97 ± 0.08	5.18 ± 0.12
Female (N=5)	0.978 ± 0.013	2.62 ± 0.14	98.6 ± 1.9	2.59 ± 0.02	4.53 ± 0.04
Overall (N=10)	0.976 ± 0.012	2.82 ± 0.24	96.8 ± 2.4	2.78 ± 0.21	4.86 ± 0.36

clustering in the mid–20 mm range. Coarse placement errors are operationally acceptable, since routine echocardiography relies on sliding and minor probe adjustments to refine the acoustic window and achieve usable images [36]. Within-subject variability was small, with standard deviations between 0.3–3.9 mm as summarized in Table III, and trial-level errors were tightly clustered around subject-specific means with minimal outliers as shown in Fig. 8. Alignment quality metrics further supported this stability, with f_{ICP} converging to 1.0 in all cases and e_{ICP} between 4–6 mm, confirming that the template registration was geometrically well constrained. Taken together, these results demonstrate that the system reliably converges to the same initial probe position on repeated runs, and that the absolute offsets of 21–30 mm are most likely attributable to inter-subject anatomical variability rather than systematic errors in the proposed algorithm.

Overall, the findings indicate that the LiDAR-based method provides a consistent and robust initial probe pose estimation, with initialization errors well within the anatomical variability encountered in clinical echocardiography.

V. CONCLUSION AND FUTURE WORK

This work presented the first demonstration of using a robot-mounted 2D LiDAR to automate probe pose initialization in teleoperated echocardiography. A plane-based calibration strategy enabled extrinsic accuracy of less than 2 mm in translation and under 0.2° rotation, after which two linear LiDAR sweeps produced consistent 3D models of the chest front surface. A scale-augmented rigid registration framework with sex-specific chest front surface priors and PCA-based orientation estimation then localized a clinically meaningful initial probe pose for placement. Quantitative validation on mannequins confirmed reconstruction accuracy within 3 mm RMSE and < 5 mm at the 95th percentile, while human trials demonstrated repeatable convergence to subject-specific initial probe position with lateral deviations of 20–30 mm, consistent with the anatomical variability of the cardiac acoustic window. Together, these findings demonstrate that LiDAR-based initial probe pose estimation offers a robust strategy for automating probe placement in tele-echocardiography.

Future work will extend this foundation in three directions. First, LiDAR scanning trajectories will be optimized adaptively to patient size, with the dual goals of reducing acquisition time and limiting irrelevant point captures or noise. Additionally, the total runtime of the system, including LiDAR scanning duration and registration computation time, will be quantified to better characterize the requirements for real-time performance. Second, we will define and evaluate

TABLE III: Tangential error e_{\parallel} (mm) for each participant, reported as mean \pm standard deviation over three independent trials.

	P1	P2	P3	P4	P5
Mean \pm SD [mm]	26.32 \pm 3.93	24.48 \pm 2.81	28.00 \pm 0.30	23.75 \pm 3.15	26.14 \pm 1.72

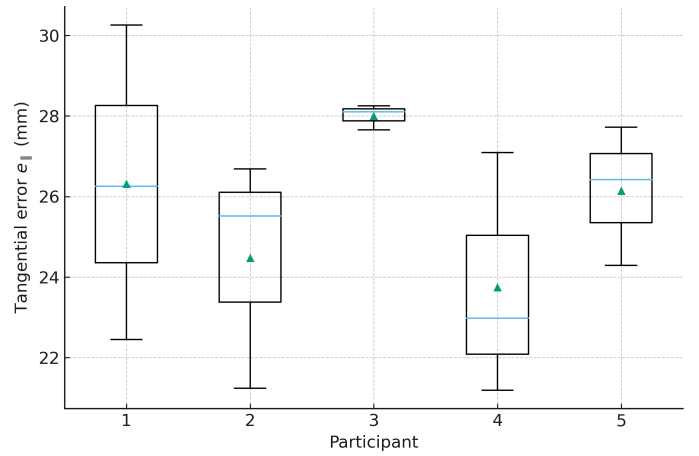


Fig. 8: Distribution of tangential error e_{\parallel} across three trials per participant (P1–P5). The central line indicates the median, the triangle marker denotes the mean, and the whiskers represent the range of observed values.

strategies for safe and efficient motion from a rest position to the identified probe initial pose, ensuring clearance of the human and bed while maintaining clinical safety margins. Finally, we plan a full clinical study to compare time-to-initialization between the proposed LiDAR-based automation and manual sonographer manipulation, as well as to capture structured feedback from sonographers regarding usability and integration into routine teleoperated workflows. These efforts will advance the system from proof-of-concept toward clinical translation, ultimately supporting more efficient and accessible tele-echocardiography.

REFERENCES

- [1] “Echocardiography in heart failure,” *Australian Journal of General Practice*, vol. 39, pp. 904–909, Nov. 2012.
- [2] P. Lancellotti, S. Price, T. Edvardsen, B. Cosyns, A. N. Neskovic, R. Dulgheru, F. A. Flachskampf, C. Hassager, A. Pasquet, L. Gargani, M. Galderisi, N. Cardim, K. H. Haugaa, A. Ancion, J.-L. Zamorano, E. Donal, H. Bueno, and G. Habib, “The use of echocardiography in acute cardiovascular care: Recommendations of the European Association of Cardiovascular Imaging and the Acute Cardiovascular Care Association,” *European Heart Journal. Acute Cardiovascular Care*, vol. 4, no. 1, pp. 3–5, Feb. 2015.
- [3] M. C. Roshan, A. Pranata, and M. Isaksson, “Robotic Ultrasonography for Autonomous Non-Invasive Diagnosis—A Systematic Literature Review,” *IEEE Transactions on Medical Robotics and Bionics*, vol. 4, no. 4, pp. 863–874, Nov. 2022.
- [4] A. Alrawashdeh, Z. Nehme, B. Williams, K. Smith, A. Brennan, D. T. Dinh, D. Liew, J. Lefkovits, and D. Stub, “Impact of emergency medical service delays on time to reperfusion and mortality in STEMI,” *Open Heart*, vol. 8, no. 1, p. e001654, May 2021.
- [5] E. M. Hidalgo, L. Wright, M. Isaksson, G. Lambert, and T. H. Marwick, “Current Applications of Robot-Assisted Ultrasound Examination,” *JACC: Cardiovascular Imaging*, vol. 16, no. 2, pp. 239–247, Feb. 2023.
- [6] K. Boman, M. Olofsson, P. Berggren, P. P. Sengupta, and J. Narula, “Robot-Assisted Remote Echocardiographic Examination and Teleconsultation: A Randomized Comparison of Time to Diagnosis With Standard of Care Referral Approach,” *JACC: Cardiovascular Imaging*, vol. 7, no. 8, pp. 799–803, Aug. 2014.
- [7] Y. Liu, F. Meng, N. Zhou, Y. Wu, C. Hu, X. Fang, H. Chen, C. Pan, X. Shu, and J. Ge, “Initial experience with a 5G cellular network and robotic arm-based remote echocardiographic system in an all-

- comers setting,” *European Heart Journal*, vol. 45, no. Supplement-1, p. ehae666.3526, Oct. 2024.
- [8] L. Wright, M. Hildalgo, M. Isaksson, T. Marwick, and G. Lambert, “Assessing the Viability and Patient Response to a Haptically-Enabled Robot for Remote Echocardiography,” *Heart, Lung and Circulation*, vol. 34, pp. S179–S180, Aug. 2025.
- [9] H. Solvin, S. Sajadi, M. Lippert, R. Massey, H. Holmström, O. Elle, K. Mathiassen, and H. Brun, “Feasibility of teleoperated robotic echocardiography – A pilot study,” *WFUMB Ultrasound Open*, vol. 1, no. 2, p. 100018, Dec. 2023.
- [10] E. M. Hidalgo, M. Isaksson, M. C. Roshan, T. H. Marwick, L. Wright, and G. Lambert, “Evaluating the impacts of network latency, haptics, and ergonomics in a Haptically-Enabled Robot for Teleoperated Echocardiography,” *Computers in Biology and Medicine*, vol. 195, p. 110450, Sept. 2025.
- [11] Z. Jiang, S. E. Salcudean, and N. Navab, “Robotic ultrasound imaging: State-of-the-art and future perspectives,” *Medical Image Analysis*, vol. 89, p. 102878, Oct. 2023.
- [12] A. S. Bin Mustafa, T. Ishii, Y. Matsunaga, R. Nakadate, H. Ishii, K. Ogawa, A. Saito, M. Sugawara, K. Niki, and A. Takanishi, “Human abdomen recognition using camera and force sensor in medical robot system for automatic ultrasound scan,” in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Osaka: IEEE, July 2013, pp. 4855–4858.
- [13] G. Ning, X. Zhang, and H. Liao, “Autonomic Robotic Ultrasound Imaging System Based on Reinforcement Learning,” *IEEE Transactions on Biomedical Engineering*, vol. 68, no. 9, pp. 2787–2797, Sept. 2021.
- [14] Y. Bi, Z. Jiang, F. Duellmer, D. Huang, and N. Navab, “Machine Learning in Robotic Ultrasound Imaging: Challenges and Perspectives,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, no. Volume 7, 2024, pp. 335–357, July 2024.
- [15] 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 in Robotics and AI*, vol. 11, May 2024.
- [16] C. Graumann, B. Fuerst, C. Hennersperger, F. Bork, and N. Navab, “Robotic ultrasound trajectory planning for volume of interest coverage,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 736–741.
- [17] J. T. Kaminski, K. Rafatizand, and H. K. Zhang, “Feasibility of Robot-Assisted Ultrasound Imaging with Force Feedback for Assessment of Thyroid Diseases,” *Proceedings of SPIE—the International Society for Optical Engineering*, vol. 11315, p. 113151D, Feb. 2020.
- [18] J. Tan, B. Li, Y. Li, B. Li, X. Chen, J. Wu, B. Luo, Y. Leng, Y. Rong, and C. Fu, “A Flexible and Fully Autonomous Breast Ultrasound Scanning System,” *IEEE Transactions on Automation Science and Engineering*, vol. 20, no. 3, pp. 1920–1933, July 2023.
- [19] L. A. Seewald, V. F. Rodrigues, M. Ollenschläger, R. S. Antunes, C. A. da Costa, R. da Rosa Righi, L. G. da Silveira, A. Maier, B. Eskofier, and R. Fahrig, “Toward analyzing mutual interference on infrared-enabled depth cameras,” *Computer Vision and Image Understanding*, vol. 178, pp. 1–15, Jan. 2019.
- [20] B.-S. Seo, B. Park, and H. Choi, “Sensing Range Extension for Short-Baseline Stereo Camera Using Monocular Depth Estimation,” *Sensors*, vol. 22, no. 12, p. 4605, Jan. 2022.
- [21] F. Imam, A. Setyanto, and K. ., “Optimising Baseline Distance in Stereo Cameras: An Experimental Approach to Enhance Object Distance Accuracy,” *International Journal of Research and Review*, vol. 12, no. 5, pp. 94–107, May 2025.
- [22] P. Padmanabhan, C. Zhang, and E. Charbon, “Modeling and Analysis of a Direct Time-of-Flight Sensor Architecture for LiDAR Applications,” *Sensors*, vol. 19, no. 24, p. 5464, Jan. 2019.
- [23] R. Whyte, L. Streeter, M. J. Cree, and A. A. Dorrington, “Application of lidar techniques to time-of-flight range imaging,” *Applied Optics*, vol. 54, no. 33, p. 9654, Nov. 2015.
- [24] X. Wu, H. Zhang, C. Kong, Y. Wang, Y. Ju, and C. Zhao, “A Lidar-based computer vision system for monitoring of patients in medical scenes,” Apr. 2023.
- [25] “Ouster OS2: High-Precision Long-Range Lidar for Autonomous Systems | Ouster,” <https://ouster.com/products/hardware/os2-lidar-sensor>, accessed: Sept. 15, 2025.
- [26] “RPLIDAR C1 Fusion Lidar | High-Res & High-Perf | SLAMTEC,” <https://www.slamtec.com/en/C1>, accessed: Sept. 15, 2025.
- [27] V. Mazzari, “How to select the right LiDAR?” <https://www.generationrobots.com/blog/en/how-to-select-the-right-lidar/>, accessed: Sept. 15, 2025.
- [28] B. Zhou, Y. He, K. Qian, X. Ma, and X. Li, “S4-SLAM: A real-time 3D LIDAR SLAM system for ground/watersurface multi-scene outdoor applications,” *Autonomous Robots*, vol. 45, no. 1, pp. 77–98, Jan. 2021.
- [29] M. C. Roshan, E. Hidalgo, M. Isaksson, L. Wright, and A. Pranata, “A low-cost probe changer for robotic ultrasonography,” in *2022 Australasian Conference on Robotics and Automation, (ACRA 2022), Queensland University of Technology, Garden Point Campus, Queensland, Australia, 6 - 8 December 2022*, vol. 2022-December. Australian Robotics and Automation Association, 2022.
- [30] T. S. Lembono, F. Suárez-Ruiz, and Q.-C. Pham, “SCALAR - Simultaneous Calibration of 2D Laser and Robot’s Kinematic Parameters Using Three Planar Constraints,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 5570–5575.
- [31] T. S. Lembono, F. Suarez-Ruiz, and Q.-C. Pham, “SCALAR: Simultaneous Calibration of 2-D Laser and Robot Kinematic Parameters Using Planarity and Distance Constraints,” *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 4, pp. 1971–1979, Oct. 2019.
- [32] S. Sharifzadeh, I. Biro, and P. Kinnell, “Robust hand-eye calibration of 2D laser sensors using a single-plane calibration artefact,” *Robotics and Computer-Integrated Manufacturing*, vol. 61, p. 101823, Feb. 2020.
- [33] M. A. Fischler and R. C. Bolles, “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” *Commun. ACM*, vol. 24, no. 6, pp. 381–395, June 1981.
- [34] D. Lopez, C. Haas, and S. Narasimhan, “Specific object finding in point clouds based on semantic segmentation and iterative closest point,” *Automation in Construction*, vol. 156, p. 105116, Dec. 2023.
- [35] M. C. Roshan, M. Isaksson, A. Pranata, and E. M. Hidalgo, “A sensor fusion approach to autonomous ultrasound imaging of the lumbar region,” *Biomedical Signal Processing and Control*, vol. 108, p. 106818, Oct. 2025.
- [36] C. Mitchell, P. S. Rahko, L. A. Blauwet, B. Canaday, J. A. Finstuen, M. C. Foster, K. Horton, K. O. Ogunyankin, R. A. Palma, and E. J. Velazquez, “Guidelines for Performing a Comprehensive Transthoracic Echocardiographic Examination in Adults: Recommendations from the American Society of Echocardiography,” *Journal of the American Society of Echocardiography: Official Publication of the American Society of Echocardiography*, vol. 32, no. 1, pp. 1–64, Jan. 2019.
- [37] J. B. Williams, D. Harmon, and J. Lindenfeld, “Physical Exam for Presence and Severity of Heart Failure,” in *Managing Heart Failure in Primary Care: A Case Study Approach*, K. M. S. Hayes and N. R. Dellise, Eds. Cham: Springer International Publishing, 2023, pp. 35–53.