POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Development and Evaluation of an Automated Intraoperative CT–US Registration Pipeline for Endovascular Navigation

Laurea Magistrale in Biomedical Engineering - Ingegneria Biomedica

Author: Tommaso Vantini

Advisor: Prof. Emiliano Votta

Co-advisor: Veronica Ruozzi

Academic year: 2024-2025

1. Introduction

Endovascular procedures increasingly rely on patient-specific anatomical models derived from preoperative CT imaging to support planning and simulation. However, during the intervention, vascular geometry may differ from the preoperative configuration due to patient positioning, limb manipulation or soft tissue deformation. As a result, a direct use of the preoperative model without intraoperative update may introduce spatial inaccuracies.

Ultrasound (US) represents a non-invasive and widely available intraoperative imaging modality, but it provides partial and modality-specific representations of the anatomy. Establishing a reliable spatial correspondence between preoperative CT models and intraoperative US data is therefore essential to enable accurate navigation and model updating within clinically acceptable time constraints. The intrinsic differences between CT and US imaging physics make this task particularly challenging, and recent works have highlighted the need for dedicated strategies to bridge the multimodal gap between the two domains [1].

1.1. Limitations of Current CT–US Registration Strategies

CT–US registration remains challenging due to the intrinsic multimodal nature of the problem. Intensity-based approaches often struggle with appearance differences between modalities, while feature-based geometric methods such as Iterative Closest Point (ICP) are sensitive to initialization and may converge to incorrect local minima in the presence of partial overlap or rotational ambiguities. Hybrid and learning-based frameworks have attempted to address these limitations by integrating segmentation and data-driven matching within unified pipelines [2]. Although such approaches improve automation and anatomical awareness, they often require large annotated datasets and remain sensitive to acquisition variability. In intraoperative contexts, registration strategies must therefore be automatic, initialization-robust and computationally efficient, without relying on manually selected anatomical landmarks.

1.2. Aim of the Thesis

The aim of this thesis is to develop an automatic CT–US registration pipeline capable of establishing a reliable spatial correspondence be-

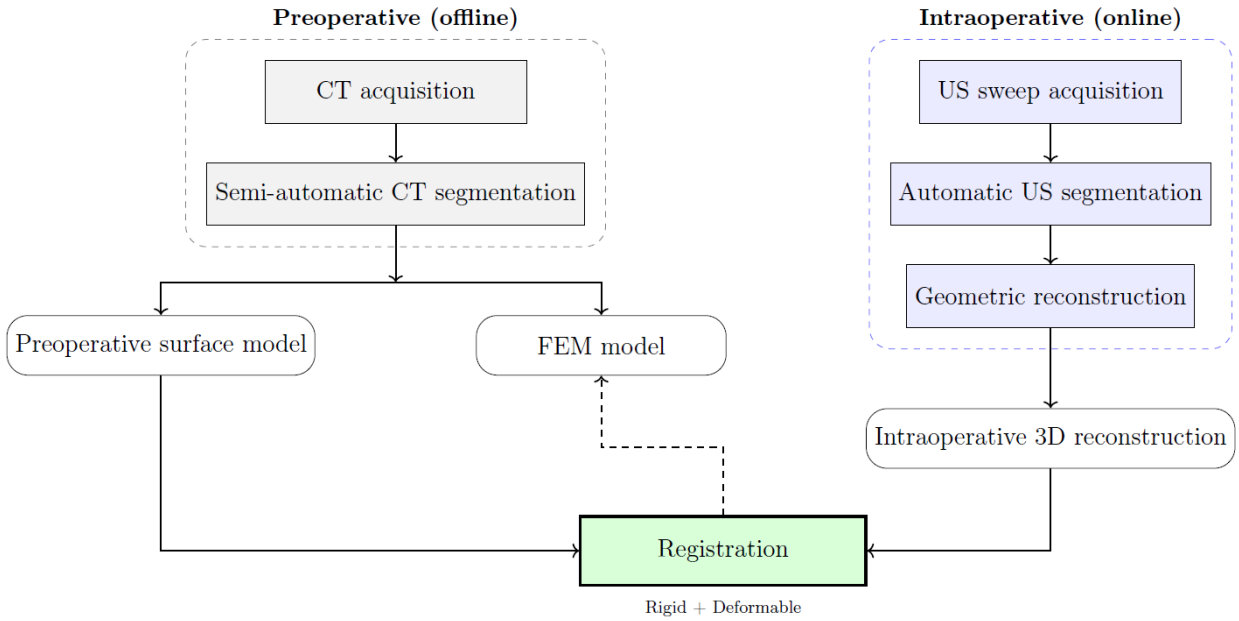

Figure 1: Overall CT–US registration pipeline.

tween a preoperative patient-specific vascular model and intraoperative US data. The proposed framework integrates a robust multi-start rigid registration strategy with a subsequent deformable refinement step applied to the reconstructed vascular geometry, based on a fast deep learning-based US segmentation method for automatic extraction of vascular structures (Figure 1).

Beyond the integration of these components, the main contributions of this work are:

- the design of a symmetry-aware multi-start rigid registration strategy to improve robustness under rotational ambiguities and partial overlap;
- the structured integration of segmentation, 3D reconstruction and deformable registration into a fully automated intraoperative-oriented pipeline;
- an experimental analysis of the impact of acquisition heterogeneity on segmentation accuracy and its implications for geometric alignment.

The overall objective is to achieve a geometrically consistent, robust and computationally efficient solution compatible with intraoperative constraints.

2. Materials and Methods

The proposed framework is designed to establish geometric alignment between a preoperative CT-derived vascular model and intraoperative ultrasound (US) data. The workflow is structured into two distinct phases:

1. **Preoperative phase:** a patient-specific vascular surface model is generated from CT imaging. This step is performed offline and does not affect intraoperative timing.
2. **Intraoperative phase:** a short robot-tracked US sweep is acquired prior to the intervention. US images and corresponding probe poses are processed to obtain a 3D vascular reconstruction, which is subsequently aligned to the preoperative model through rigid and deformable registration.

This structure ensures minimal operator interaction and concentrates the computational effort in a preparatory step compatible with clinical workflow constraints.

2.1. System Architecture and Reference Frames

All spatial entities are expressed within a consistent chain of reference frames. The world frame W is attached to the robotic base, while additional frames describe the end effector E, the probe P, and the image frame I. Rigid transformations are represented in homogeneous form

as ${}^{A}\mathbf{T}_{B}$, mapping coordinates from frame B to frame A. The overall transformation chain linking image and world coordinates is defined as

$$ {}^{W}\mathbf{T}_{I} = {}^{W}\mathbf{T}_{E} \circ {}^{E}\mathbf{T}_{P} \circ {}^{P}\mathbf{T}_{I} $$

This formulation, together with image spacing, enables consistent mapping of segmented image contours into the global coordinate system for geometric reconstruction and registration.

2.2. Data Acquisition and Synchronization

During the intraoperative phase, US images and probe poses are acquired through a ROS2-based communication framework and synchronized to ensure spatial consistency between each frame and its associated probe configuration. A calibration procedure estimates the fixed transformation between probe and image frames.

2.3. Ultrasound Segmentation

Automatic vessel segmentation is performed using convolutional neural networks of the U-Net [3] family, selected for their robustness in biomedical image segmentation under limited data conditions.

Model training. Training is conducted offline using k-fold cross-validation to obtain stable performance estimates. Data augmentation and dropout are adopted to mitigate overfitting. Independently trained models are combined through parameter averaging into a single consolidated network. Probability calibration is applied to improve reliability of confidence estimates, followed by hysteresis thresholding to improve spatial coherence.

Inference within the pipeline. During the intraoperative procedure, US frames undergo pre-processing before being fed to the network. The resulting probability maps are converted into binary masks through hysteresis thresholding. Morphological post-processing is applied to obtain geometrically consistent vessel contours suitable for reconstruction.

2.4. 3D Reconstruction

From each binary mask, vessel contours are extracted in image coordinates (r, c) and mapped into the world frame through the calibrated transformation chain, producing spatially located 3D samples. Aggregating successive frames yields an intraoperative surface point cloud. Point cloud downsampling and filtering are applied to homogenize density and suppress outliers. A surface mesh is reconstructed using Poisson surface reconstruction, followed by smoothing to improve regularity. An intraoperative centerline is then extracted from the mesh and used to initialize rigid alignment.

2.5. Point Cloud Registration

Geometric alignment between the reconstructed ultrasound geometry and the preoperative model is performed in two stages:

Rigid registration. A symmetry-aware rigid alignment strategy is adopted. Given a set of input points, Principal Component Analysis (PCA) is used to estimate the main axis of the geometry. Candidate orientations are generated by rotating around the principal axis and by testing possible axis inversions to address axial symmetries typical of tubular vascular structures. For each candidate initialization, ICP [4] is applied to minimize the root mean square error (RMSE) between corresponding points. The solution giving the lowest RMSE is selected. This rigid registration block is applied first to the extracted centerline to obtain a robust coarse alignment and subsequently to the full surface geometry, initialized with the transformation estimated from the centerline.

Deformable registration. After rigid alignment, a non-rigid refinement step is applied to account for local geometric discrepancies between intraoperative and preoperative surfaces. Deformable registration is performed using Coherent Point Drift (CPD) [5], a probabilistic framework that estimates a continuous and smooth displacement field. The deformable step improves local correspondence while preserving geometric plausibility and avoiding overfitting to local noise.

The sequential application of multi-start ICP rigid alignment and CPD deformable refinement provides a robust balance between global consistency and local accuracy.

3. Experimental part

All experiments were conducted on a laptop workstation equipped with an Intel Core i7-13620H CPU, 16 GB RAM and an NVIDIA GeForce RTX 4060 Laptop GPU (8 GB). The

same hardware configuration was used for reconstruction, registration and segmentation experiments.

3.1. In-vitro Phantom Experiments

Experimental evaluation of the registration pipeline was conducted on three vascular phantoms named **Artificial**, **Vein** and **Artery**. The phantoms were fabricated using gel-wax with additives (paraffin wax 7% w/w and talcum powder 1% w/w) to obtain suitable mechanical stability and ultrasound visibility. The Vein and Artery geometries were derived from the same preoperative CT model used as fixed point cloud in the pipeline, ensuring consistency between digital and physical representations. For each phantom, 8 tracked US sweeps were acquired using the robotic setup.

Reconstruction consistency was evaluated by comparing repeated reconstructions of the same phantom geometry. Surface discrepancies were quantified using mean surface distance and 95th percentile Hausdorff Distance (HD95) computed pairwise across acquisitions.

Rigid registration robustness was assessed by comparing the proposed symmetry-aware multi-start strategy with a conventional single-start ICP baseline under controlled perturbations of the initial alignment. Convergence behavior and alignment quality were quantified using RMSE and success rate.

Registration accuracy was evaluated by comparing alignment results after rigid registration and after deformable refinement. Distance-based metrics (RMSE and HD95) were computed between the intraoperative reconstruction and the corresponding reference model.

Computational performance was assessed by measuring the execution time of the main processing stages, from ultrasound acquisition to final registration, in order to estimate overall workflow duration.

3.2. Segmentation Experiments on Phantom and In-vivo Data

As segmentation constitutes the first intraoperative processing step and directly determines the geometric reliability of the reconstructed surface, a dedicated experimental analysis was conducted to assess its robustness across both controlled phantom data and heterogeneous in-vivo acquisitions. The segmentation component was evaluated on four datasets defined in the thesis:

- **Phantom**: ultrasound acquisitions of the Artificial, Vein and Artery phantoms;
- **Carotid**: publicly available in-vivo carotid ultrasound dataset;
- **San Raffaele**: in-house in-vivo femoral vein ultrasound dataset;
- **Mus-V**: heterogeneous in-vivo vascular ultrasound dataset, reformulated into binary tasks (*Mus-V Artery* and *Mus-V Vein*).

Performance was quantified using Dice Similarity Coefficient (DSC) and HD95.

Model comparison was performed across three encoder–decoder architectures within the U-Net family (U-Net, Attention U-Net and U-Net3+) on all datasets, in order to assess how architectural complexity influences segmentation performance under different levels of variability.

Dataset variability analysis was conducted on the San Raffaele dataset by comparing *Single-session*, *Full* and *Pruned* configurations. This experiment was designed to investigate the impact of intra-dataset variability and dataset size on segmentation performance.

Vessel type and **domain aggregation analysis** was carried out on the Mus-V dataset to evaluate segmentation performance with respect to artery versus vein discrimination and to analyze the effect of aggregating data from different anatomical districts.

4. Results and Discussion

4.1. In-vitro Registration Results

Reconstruction consistency. Repeated acquisitions of the same phantom geometry produced comparable reconstructed surfaces, with median HD95 values below 0.8 mm across geometries. Slightly higher variability was observed for the Artery phantom, consistent with its smaller diameter and increased curvature. Overall, reconstruction stability was sufficient to support subsequent registration.

Rigid registration robustness. The proposed strategy showed improved robustness compared to a conventional single-start ICP approach. While the baseline method exhibited convergence failures under unfavorable rotational initialization, the proposed strategy maintained high success rates across perturbations.

In particular, the multi-start approach preserved stable RMSE values below 1 mm in configurations where the baseline frequently diverged, confirming the effectiveness of PCA-based axis estimation and symmetry handling.

Registration accuracy. Figure 2 compares registration accuracy after rigid alignment and deformable refinement. For all phantom geometries, deformable registration reduced both RMSE and HD95. For the Artery phantom, RMSE decreased from approximately 1.4 mm after rigid alignment to below 0.6 mm after refinement, with similar relative reductions (40–60%) observed for the Artificial and Vein geometries. Final residual errors remained within the millimetric range, indicating effective compensation of geometric discrepancies not resolved by rigid alignment alone.

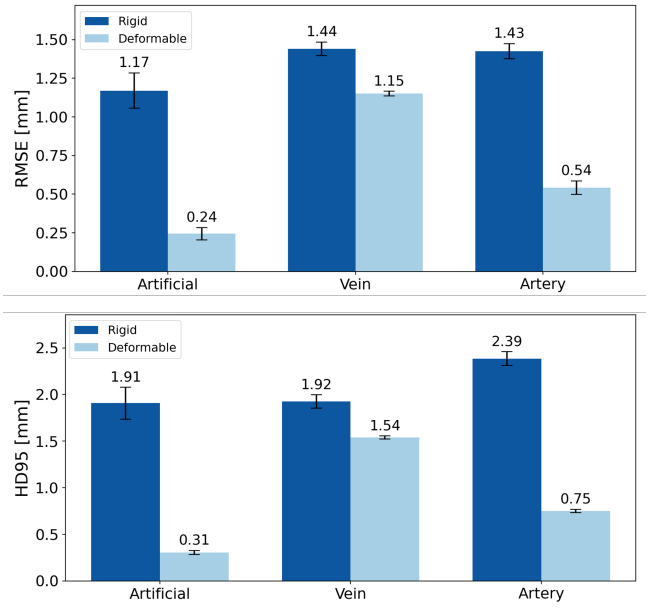

Figure 2: Registration accuracy comparison between rigid and deformable alignment across phantom geometries.

Computational performance. Figure 3 reports the runtime breakdown of the pipeline. The overall acquisition-to-registration workflow required approximately 40–70 s depending on phantom size. US acquisition and pre-processing represent the dominant contributions, while rigid and deformable registration accounted for a limited portion of the total runtime. These results confirm compatibility of the pipeline with intraoperative preparatory phases.

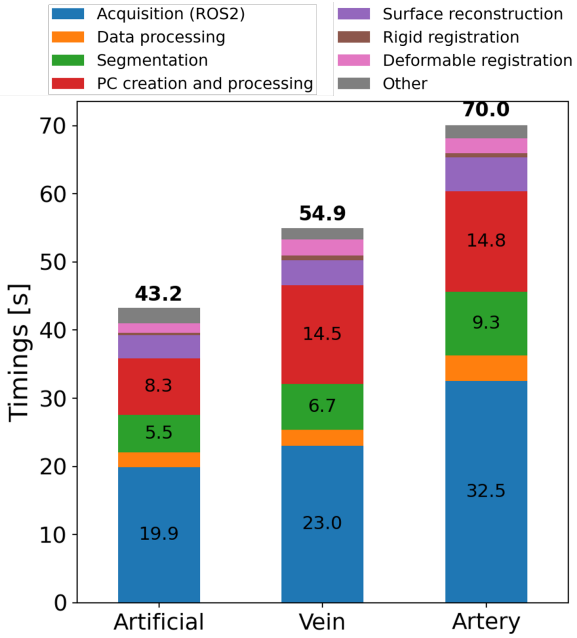

Figure 3: Runtime breakdown of the registration pipeline across phantom geometries.

4.2. In-vivo Segmentation Results

Model comparison. Across datasets, U-Net-based architectures achieved comparable performance. On the Phantom and Carotid datasets, Dice scores ranged between 89% and 93%, reflecting controlled acquisition conditions and homogeneous appearance. Increased architectural complexity did not consistently improve generalization.

Dataset variability. Figure 4 shows the impact of dataset heterogeneity in the San Raffaele dataset. Dice scores remained above 82% in the homogeneous *Single-session* configuration, while performance decreased markedly in the more heterogeneous *Full* setting, reaching values as low as 61.4% for U-Net3+. The *Pruned* configuration partially recovered performance, indicating that variability has a stronger impact than moderate dataset size reduction. Similar trends, but more pronounced, were observed in HD95, highlighting the relevance of boundary accuracy for downstream geometric reconstruction.

Vessel type and domain variability. On the Mus-V dataset, arterial segmentation consistently outperformed venous segmentation. Aggregated arterial training reached 89.8–91.2% Dice, whereas venous aggregation ranged between 58.4–72.1%. Domain aggregation im-

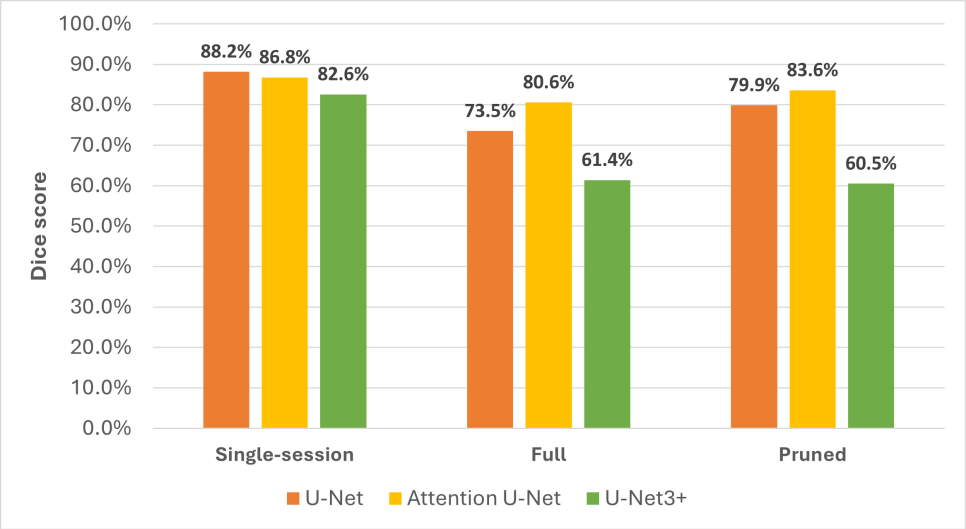

Figure 4: DSC comparison across dataset configurations (*Single-session*, *Full*, *Pruned*) and architectures, showing the effect of acquisition heterogeneity.

proved performance for arteries but degraded it for veins, indicating higher sensitivity of venous segmentation to cross-domain heterogeneity.

5. Conclusions

This work presented an automated CT-US registration framework integrating US segmentation, 3D reconstruction and geometric alignment in a unified pipeline. The proposed strategy achieved millimetric-level registration accuracy in controlled phantom experiments, with deformable refinement consistently improving rigid alignment while maintaining computational times compatible with a pre-interventional setup phase (40–70 s overall workflow duration). Segmentation experiments demonstrated robust performance under coherent training conditions, while acquisition heterogeneity emerged as the main factor limiting generalization, particularly at the boundary level relevant for geometric reconstruction. Since surface alignment accuracy directly depends on contour precision, these findings highlight the importance of managing dataset variability in clinically oriented pipelines. Overall, the proposed framework supports the feasibility of automated CT-US alignment with millimetric accuracy and clinically compatible runtime within a controlled setting. Further validation on in-vivo data remains necessary.

References

[1] Y. Velikova, W. Simson, M. F. Azampour, P. Paprottka, and N. Navab, "Cactuss: Common anatomical ct-us space for us examinations," *International Journal of Computer Assisted Radiology and Surgery*, vol. 19, pp. 861–869, 2024.

[2] M. F. Azampour, M. Tirindelli, J. Lameski, M. Gafencu, E. Tagliabue, E. Fatemizadeh, I. Hacihaliloglu, and N. Navab, "Anatomy-aware computed tomography-to-ultrasound spine registration," *Medical Physics*, vol. 51, no. 3, pp. 2044–2056, 2024.

[3] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *arXiv preprint arXiv:1505.04597*, 2015.

[4] P. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.

[5] 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.