

# Towards Patient-Specific Deformable Registration in Laparoscopic Surgery

Alberto Neri<sup>1,2,3</sup>, Nazim Haouchine<sup>3</sup>, Veronica Penza<sup>1</sup>, and Leonardo S. Mattos<sup>1</sup>

<sup>1</sup> Biomedical Robotics Lab, Istituto Italiano di Tecnologia, Genoa, Italy

<sup>2</sup> Department of Computer Science, Bioengineering, Robotics and Systems Engineering (DIBRIS), University of Genoa, Italy

<sup>3</sup> Harvard Medical School, Brigham and Women's Hospital, Boston, USA  
alberto.neri@iit.it

**Abstract.** Unsafe surgical care is a critical health concern, often linked to limitations in surgeon experience, skills, and situational awareness. Integrating patient-specific 3D models into the surgical field can enhance visualization, provide real-time anatomical guidance, and reduce intraoperative complications. However, reliably registering these models in general surgery remains challenging due to mismatches between pre-operative and intraoperative organ surfaces—such as deformations and noise. To overcome these challenges, we introduce the first deep learning-based non-rigid point cloud registration method that is genuinely patient-specific, being both trained and tested on the same individual's anatomy. Our approach combines a Transformer encoder-decoder architecture with overlap estimation and a dedicated matching module to predict dense correspondences, followed by a physics-based algorithm for registration. Experimental results on both synthetic and real data demonstrate that our patient-specific method significantly outperforms traditional agnostic approaches, achieving 45% Matching Score with 92% Inlier Ratio on synthetic data, highlighting its potential to improve surgical care.

**Keywords:** Point cloud registration · Image-guided Surgery · Laparoscopy

## 1 Introduction

Unsafe surgical care remains a significant global health concern, contributing to patient complications, prolonged hospital stays, increased healthcare costs, and potential disability or death [16]. Adverse events are often related to the experience, skills, and situational awareness of surgeons, particularly in recognizing anatomical variations and integrating preoperative data intraoperatively [15]. Enhancing surgical awareness by overlaying patient-specific 3D models - such as organs, tumors, and vessels - onto the surgical field, can improve visualization and real-time anatomical guidance and reduce intraoperative complications [1]. The application in general surgery is particularly challenging due to the dynamic nature of organ tissues [3]. Factors like patient positioning, pneumoperitoneum

insufflation, and physiological movements cause intraoperative organ surfaces to differ from preoperative imaging, making model-to-image registration a critical yet unresolved challenge in laparoscopy.

Point cloud registration methods have been proposed to align the preoperative organ surface point cloud, extracted from imaging modalities like CT scan, with an intraoperative point cloud, obtained from e.g. stereo vision of the surgical field, by estimating the transformation needed to match corresponding anatomical structures. Early non-deep-learning methods rely on finite-element models of liver mechanics: [7] cast registration as an optimal control problem on a hyper-elastic FEM, solving for the surface forces that reproduce sparse displacement, while [4] build a linear modal basis from relaxed control-point perturbations and use a Levenberg–Marquardt solver to fit mode weights to surface and vascular landmarks. Moving to deep learning, a state-of-the-art algorithm in non-rigid partial point cloud registration applied to generic objects is Lepard [5]. Their model combines a fully convolutional feature extraction, Transformer-based self cross-attention and differentiable matching to predict point cloud correspondences. Despite the excellent performances on synthetic data of deep learning-based point cloud registration, challenges remain in real-world medical applications due to the disparity between dense, noise-free preoperative point clouds and intraoperative point clouds, which are often partial ( $< 20\%$ ), noisy [20], and deformable, making accurate alignment difficult. To this end, recent methods have been developed to improve registration in challenging surgical environments. V2S-Net [9] learns a data-driven deformation basis from simulated tissue deformation to align preoperative volumes with full intraoperative surfaces. [2] proposes a rigid registration algorithm for liver surfaces in a low overlap scenario. It combines local and global features to estimate the overlap mask used to filter the non-overlapping region of the point cloud. LiverMatch [17] employs a transformer-based network for matching complete and partial liver point clouds in deformed conditions. Furthermore, the authors propose their synthetic dataset starting from 16 livers in the 3D-IRCADb-01 dataset [11] and applying deformations and crops to them. Zhang et al. [21] propose KCR-Net, integrating a Neighborhood Feature Fusion Module (NFFM) for robust keypoint registration, even in low-overlap scenarios. Despite these advancements, challenges remain in handling intraoperative deformations, noise, and complete-to-partial registrations.

All these state-of-the-art methods adopt an agnostic strategy, training on large datasets to generalize to new cases. However, we argue this approach may not be ideal in surgical settings, where error margins are extremely low and optimal outcomes are critical for each patient. To address this limitation, we introduce a novel patient-specific non-rigid point cloud registration method that leverages preoperative information to tailor the registration process for each individual. Uniquely, our model is both trained and tested on the same patient’s anatomy. It employs a Transformer encoder-decoder architecture combined with overlap estimation and a dedicated matching module to predict dense correspondences. Additionally, we introduce a novel on-the-fly data generation strategy

during training, which includes dynamic patient-specific deformation generation and an innovative visible crop logic. Experimental results on both synthetic and real data confirm that our patient-specific approach significantly outperforms traditional agnostic methods.

## 2 Methods

The architecture of the proposed registration method is presented in Fig. 1. The pipeline consists of five main modules: *Keypoint Features Extraction*, *Keypoint Features Conditioning*, *Complete-to-Partial Overlap Prediction*, *Point-to-Node Decoder* and *Matching Non-Rigid Registration*. First, KPConv is utilized to extract keypoints and their corresponding features from the input point clouds. Specifically, let  $\mathbf{X} \in \mathbb{R}^{M \times 3}$  represent the complete point cloud of the organ of in-




Fig. 1: The network uses KPConv to detect keypoints and extract their features, refines these features through self- and cross-attention, and then feeds them into the decoder and matching module to predict dense point correspondences.

terest (e.g., derived from a 3D model obtained via a CT scan), and let  $\mathbf{Y} \in \mathbb{R}^{N \times 3}$  denote a partial point cloud of the same organ (e.g., captured using an endoscopic camera), where  $N \ll M$ . The extracted features are then processed by the encoder, which employs self-attention and cross-attention mechanisms to enhance feature representation and alignment. The *Complete-to-Partial Overlap Prediction* module predicts the regions of  $\mathbf{X}$  and  $\mathbf{Y}$  that correspond to the same anatomical area, helping to refine keypoint matching. Keypoint features and overlap scores are extended to the dense representation using a *Point-to-Node Decoder*. Subsequently, the *Matching* module determines correspondences by leveraging the similarity matrix computed between these features and the original 3D coordinates.

**Keypoint Features Extraction Module** Kernel Point Convolution (KPConv) [14] is used for feature extraction. The KPConv backbone converts the input point cloud into a reduced set of keypoints  $\mathbf{K}_X \in \mathbb{R}^{M' \times 3}$  and  $\mathbf{K}_Y \in \mathbb{R}^{N' \times 3}$  through a series of ResNet-like blocks and strided convolutions, effectively down-sampling the two clouds to achieve similar densities. In addition to the keypoints, the associated features  $\mathbf{F}_{K_X} \in \mathbb{R}^{M' \times D}$  and  $\mathbf{F}_{K_Y} \in \mathbb{R}^{N' \times D}$  are also produced.

**Keypoint Features Conditioning (Encoder) Module** We reduce the dimension of the keypoints features (i.e.  $\mathbf{F}_{K_X}$  and  $\mathbf{F}_{K_Y}$ ) to  $d = 256$  through a linear projection. Additionally, we apply sinusoidal positional encoding to the keypoint coordinates (i.e.  $\mathbf{K}_X$  and  $\mathbf{K}_Y$ ) to provide the transformer with an understanding of their spatial relationships. Both the high-level features and positional encodings are fed into the cross-encoder. High-level features enable the transformer to interpret and compare semantic content within the point clouds during alignment, while positional encoding supplies spatial context by emphasizing the structure and relative arrangement of points.

Following [18], each transformer cross-encoder layer consists of three sub-layers: (i) a multi-head self-attention layer that processes each point cloud independently; (ii) a multi-head cross-attention layer that updates the features of one point cloud by incorporating information from the other; and (iii) a position-wise feed-forward network.  $\mathbf{C}_{K_X} \in \mathbb{R}^{M' \times d}$  and  $\mathbf{C}_{K_Y} \in \mathbb{R}^{N' \times d}$  are the conditioned keypoint features produced by the encoder.

**Overlap prediction Module** The overlap module leverages conditioned features to compute keypoint overlap scores, denoted as  $\hat{\mathbf{s}} = [\hat{\mathbf{s}}_{K_X}, \hat{\mathbf{s}}_{K_Y}]$  using a linear fully connected layer with sigmoid activation, following:

$$\hat{\mathbf{s}} = \begin{cases} \hat{\mathbf{s}}_{K_X} = 1/(1 + e^{-(\mathbf{C}_{K_X} \mathbf{W}_3 + \mathbf{b}_3)}) \\ \hat{\mathbf{s}}_{K_Y} = 1/(1 + e^{-(\mathbf{C}_{K_Y} \mathbf{W}_3 + \mathbf{b}_3)}) \end{cases} \quad (1)$$

where  $\mathbf{W}_3$  and  $\mathbf{b}_3$  are learnable weights and biases parameters. These scores quantify the probability that a keypoint lies within the overlapping region. In our low-overlap scenario—where less than 20% of  $\mathbf{X}$  is covered by the target—accurate overlap score prediction is vital [22], as it directs the model to focus its correspondence search on the relevant region.

**Point-to-Node Decoder** To establish dense correspondences between the two clouds, we employ a decoder that propagates features and overlap scores across every dense point. Instead of using the conventional KPConv decoder, which relies on k-nearest neighbour search, we adopted a point-to-node grouping strategy similar to [19]. This approach offers two main advantages: (1) each point is uniquely assigned to one node (i.e. keypoint), ensuring no point is left unassigned, and (2) it inherently adapts to various scales. After grouping, we obtain dense point features ( $\mathbf{C}_X \in \mathbb{R}^{M \times d}$ ,  $\mathbf{C}_Y \in \mathbb{R}^{N \times d}$ ) and overlap scores ( $\hat{\mathbf{s}}_X, \hat{\mathbf{s}}_Y$ ). The features are concatenated with the point’s 3D coordinates ( $\mathbf{C}_X^* \in \mathbb{R}^{M \times d+3}$ ,  $\mathbf{C}_Y^* \in \mathbb{R}^{N \times d+3}$ ), and this combined vector is fed into a MLP to estimate the final coordinates for  $\mathbf{X}$ . Our MLP first projects the concatenated input from a  $d + 3$  dimensional space into a  $d$ -dimensional representation, applies non-linear activations, and ultimately outputs a 3-dimensional vector corresponding to the deformed coordinates ( $\hat{\mathbf{X}}, \hat{\mathbf{Y}}$ ). Specifically:  $\hat{\mathbf{X}} = \text{ReLU}(\mathbf{C}_X^* \mathbf{W}_1 + \mathbf{b}_1) \mathbf{W}_2 + \mathbf{b}_2$ , where  $\mathbf{W}_1, \mathbf{W}_2$  and  $\mathbf{b}_1, \mathbf{b}_2$  are learnable weights and biases, respectively. At this stage, the model attempts an initial non-rigid registration. Although the



deformed coordinates it produces are rough, they provide a valuable signal for the chamfer loss to effectively guide the training process. Separately, the overlap scores are refined through a linear layer.

**Matching and Physics-based Non-rigid Registration** Within the Matching module, we begin by refining the  $\mathbf{C}_X$  and  $\mathbf{C}_Y$  using a linear layer. Then, we compute a similarity matrix and convert it into a confidence matrix  $\mathbf{M}$  using a dual-softmax operation. Finally, following [17], matches  $\mathbf{m}$  are selected from  $\mathbf{M}$  based on a thresholded Mutual Nearest Neighbor criterion.

Once the matching is obtained, we can formulate the non-rigid registration problem as an energy minimization:  $\min_{\delta\mathbf{X}} \frac{1}{2}\delta\mathbf{X}^T\mathbf{S}\delta\mathbf{X} + k\|\mathbf{Y}_m - \mathbf{X}_m\|^2$ , where  $\delta\mathbf{X}$  is the displacement field to be estimated,  $\mathbf{S}$  is the stiffness matrix, which encodes the biomechanical properties of the organ,  $k$  is a scalar stiffness parameter that converts displacement to external forces.

The first term  $\frac{1}{2}\delta\mathbf{X}^T\mathbf{S}\delta\mathbf{X}$  represents the internal elastic energy, ensuring that deformations follow biomechanical constraints, in particular in non-visible parts of the organ. The second term,  $k\|\mathbf{Y}_m - \mathbf{X}_m\|^2$ , enforces alignment with the observed correspondences. Since this formulation is static (no acceleration), the system does not include mass or damping terms. This system is solved using the Conjugate Gradient method.

**Optimization Losses** Our method employs a weighted sum of three losses: (i) a matching loss (ML) [6] to supervise the confidence matrix  $\mathbf{M}$  with the ground truth matches; (ii) a chamfer loss (CL) weighted on the overlap score, defined as:  $CL(\hat{\mathbf{X}}, \mathbf{Y}, \hat{\mathbf{s}}_X) = \frac{1}{|\hat{\mathbf{X}}|} \sum_{\hat{x} \in \hat{X}} \hat{s}_x \min_{y \in Y} \|\hat{x} - y\|^p$ . This loss is not focused on finding exact 1-1 correspondences but rather a global geometrical alignment. Finally, (iii) an overlap loss (OL) similar to [18] designed to optimize the overlap scores.

**Patient-specific Training Dataset Generation** Generating a patient-specific training dataset involves creating pairs of complete and partial 3D point clouds from pre-operative CT scans, incorporating both rigid and non-rigid transformations. We employ the As-Rigid-As-Possible (ARAP) algorithm [12] to generate two types of deformation: (i) compression deformation, which simulates the effect of  $CO_2$  insufflation by displacing control points along their normal vectors with randomly assigned magnitudes in the range  $[0, 0.1]$ ; and (ii) lobe deformations, introduced by applying random displacements to control points within the lobe region, with magnitudes ranging from  $[0, 0.25]$ .

To realistically generate partial point clouds similar to an endoscopic scene, we generate a dummy camera position that points toward the visible surface of the organ. We sample random spherical coordinates while constraining the polar and azimuthal angles to a realistic range for intraoperative settings. Using the camera direction and the surface normals of the organ, we compute the dot

product to assess visibility. Only 5% of points closest to the camera, with a dot product below  $80^\circ$ , are retained.

Following the criteria above, we apply deformation, and cropping is applied randomly to constitute  $\mathbf{Y}$ , while a random rigid transformation is applied to  $\mathbf{X}$ . In addition, both  $\mathbf{X}$  and  $\mathbf{Y}$  are shuffled. This pipeline is executed on the fly for each intraoperative cloud, with a processing time of 0.5 seconds per sample. Since random parameters govern deformations and cropping, the model is exposed to a continuously varying dataset during training, eliminating the need for a pre-computed offline dataset. Conversely, validation is performed using deterministic seeds to ensure consistency. The test set follows the same generation process as validation but with a different seed. Since our approach is patient-specific, a new dataset must be generated and the model retrained for every patient. By applying early stopping based on the validation loss, training completes in approximately 12 hours.

### 3 Results

**Matching Performances** To evaluate the matching performance of our model, we employed “liver 1” from the 3D-IRCADb-01 dataset and generated a test set comprising 50 examples, following the procedure described in Section 2. We com-

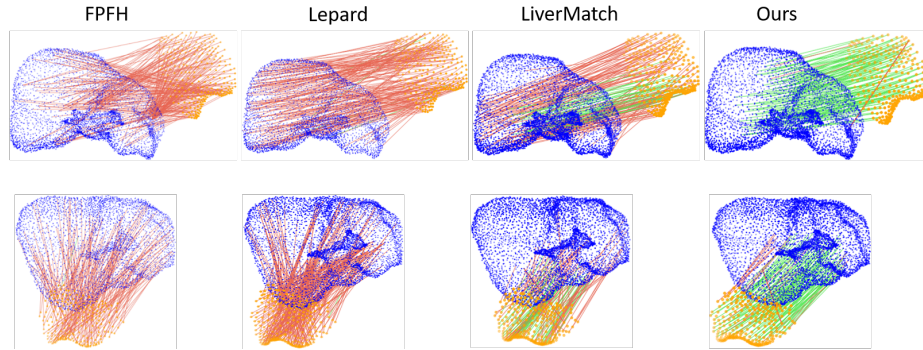


Fig. 2: Qualitative results of predicted matches. Blue points denote the preoperative point cloud, while orange points represent the partial and deformed intraoperative cloud. Red lines indicate incorrect matches, and green lines highlight exact matches relative to the ground truth.

pared our approach with three alternative algorithms: (i) FPFH – where FPFH features [10] are extracted from both complete and partial point clouds, and correspondences are established by matching each target point to its nearest source point in the feature space using Euclidean distance; (ii) Lepard [5] – by leveraging their pre-trained model on the 4DMatch dataset, a collection of dynamic,

Table 1: Matching Metrics

Methods	MS (%)	IR (%)	Avg. #MP
<i>FPFH</i>	$0.03 \pm 0.11$	$0.03 \pm 0.11$	$< 1$
<i>Lepard</i>	$1.6 \pm 1.5$	$2.4 \pm 2.2$	$4 \pm 3.75$
<i>LiverMatch</i>	$13 \pm 9$	$23 \pm 13$	$39 \pm 27$
<i>Ours - ML + CL</i>	$36 \pm 8$	$90 \pm 4$	$90 \pm 20$
<b><i>Ours - Full</i></b>	$45 \pm 7$	$92 \pm 3$	$112.5 \pm 17.5$

real-world, partially overlapping point cloud pairs; and (iii) LiverMatch [17] – utilizing the pre-trained model derived from liver data within the 3D-IRCADb-01 dataset. Notably, “liver 1” is absent from the LiverMatch training set, as it was also used for test set generation in their work, we call IRCAD-Liver1. We evaluated performance using the Matching Score (MS)—defined as the number of correctly predicted matches relative to the total available matches—and the Inlier Ratio (IR), which represents the proportion of correct matches among all predictions. Additionally, we report the absolute number of exact matches predicted (#MP) by each algorithm. Table 1 demonstrates that our algorithm significantly outperforms its competitors in identifying exact matches. Furthermore, we evaluated our algorithm under various loss configurations. Our results indicate that optimal performance is achieved when all the three losses are active (**Full**), confirming the synergistic relationship between overlap score estimation and correspondence estimation. Figure 2 visually compares the quality of the matches predicted by each method.

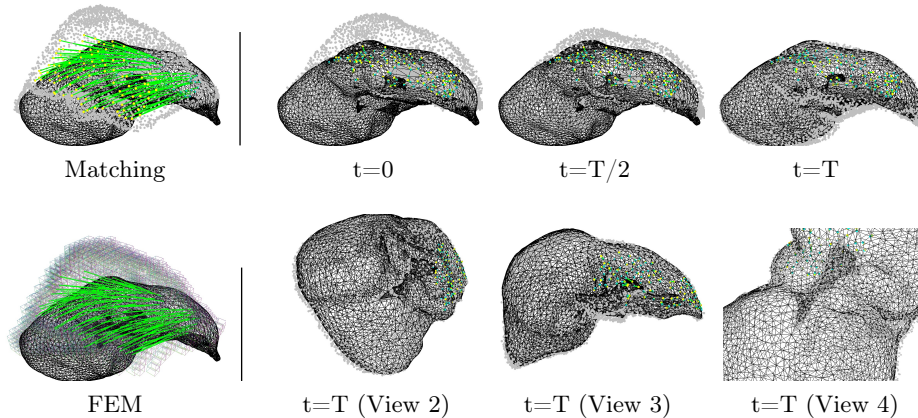


Fig. 3: Non-rigid registration results on IRCAD-Liver1: Source point cloud is represented with gray points while target point clouds are represented in black wireframe. Green lines represent the matching correspondences and FEM models are represented with grid cells. Registered point clouds are presented at  $t=0$ ,  $t=T/2$  and  $t=T$  are shown in 4 different views.

**Non-rigid Registration** We present results on non-rigid registration using biomechanical modeling. We use the framework SOFA [13] to generate a Finite Element model from the liver mesh and generated around 12000 tetrahedra elements. We use a Young’s modulus of 1.5 KPa and a Poisson ratio of 0.45. We compute Target Registration Errors (TREs) as the point-to-point Euclidean distance between vertices of the non-visible part of the preoperative and intraoperative meshes. We also report Fiducial Registration Errors (FREs). Results shown in Table 1 using 10 random deformations demonstrate that we obtain very low TREs with an average of  $4.82 \pm 3.33$  mm. FREs are also very low,  $1.68 \pm 1.11$  mm, suggesting few matching outliers and adequate registration. Moreover, our method outperforms LiverMatch, which exhibits a high number of mismatches, resulting in a TRE of  $17.36 \pm 13.68$  mm and FRE of  $28.76 \pm 21.30$  mm. Visualization of the non-rigid registration and FEM model is provided in Figure 3. We then tested our method on the DePoLL porcine dataset [8] where 3D intraoperative point clouds reconstructed from the endoscopic images. We use a Young’s modulus of 3 KPa and a Poisson ratio of 0.45. We report in Table 2 registration results on 13 cases from DePoll dataset.

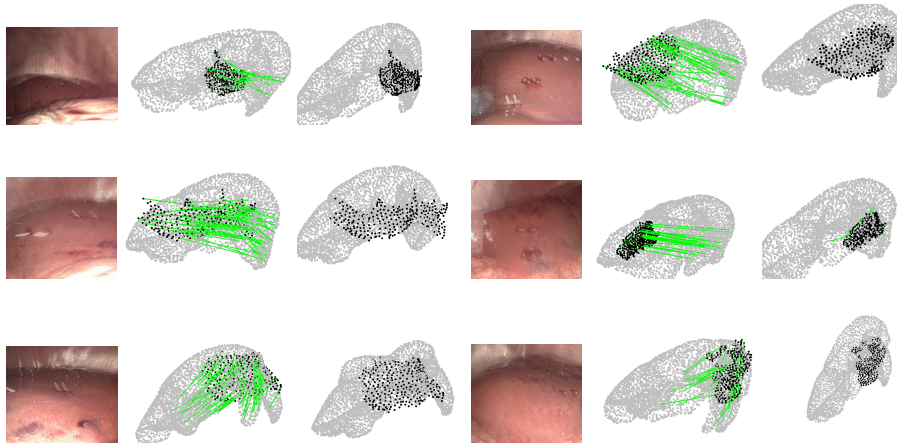


Fig. 4: Qualitative results of non-registration on six cases from the DePoLL dataset are presented, showing the endoscopic image, the matching, and the registration from left to right. Preoperative point clouds are represented in gray, intraoperative point clouds in black, and correspondences with green lines.

We obtain an average Hausdorff distance (HD) of  $8.45 \pm 3.60$  mm and an average FRE of  $15.90 \pm 6.41$  mm, suggesting a good surface-to-surface registration and good matching. Visualizations of six cases are provided in Figure 4. It is important to note that this test differs from those proposed in other works [21], as it involves registering the reconstructed point cloud with the preoperative one, rather than with the reconstruction from the intraoperative CT scan.

Table 2: Registration results on DePoLL dataset (results in mm).

Met.	Case0	Case1	Case2	Case3	Case4	Case8	Case10	Case5	Case6	Case7	Case9	Case11	Case12	Avg
HD	11.12	8.44	6.62	6.23	15.45	10.17	12.04	8.12	5.65	6.93	6.88	5.98	7.75	8.45 $\pm$ 3.60
FRE	27.68	18.68	12.53	12.23	32.98	26.68	23.02	10.43	9.46	15.04	8.84	10.54	9.83	15.90 $\pm$ 6.41

## 4 Conclusion

We presented a novel method for intraoperative point cloud registration. By leveraging a patient-specific partial and deformable data generation strategy and training a Transformer-based matching network, our approach outperforms traditional agnostic methods, demonstrating improved accuracy and robustness, on synthetic and real surgical dataset.

**Acknowledgments.** This research was partially funded by the European Union - NextGenerationEU, by the Ministry of University and Research (MUR), National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.5, project “RAISE - Robotics and AI for Socio-economic Empowerment” (ECS00000035), and NIH grant K25EB035166.

**Disclosure of Interests.** The authors have no competing interests to declare.

## References

- Bernhardt, S., Nicolau, S.A., Soler, L., Doignon, C.: The status of augmented reality in laparoscopic surgery as of 2016. *Medical image analysis* **37**, 66–90 (2017)
- Guan, P., Luo, H., Guo, J., Zhang, Y., Jia, F.: Intraoperative laparoscopic liver surface registration with preoperative ct using mixing features and overlapping region masks. *International Journal of Computer Assisted Radiology and Surgery* **18**(8), 1521–1531 (2023)
- Haouchine, N., Dequidt, J., Peterlik, I., Kerrien, E., Berger, M.O., Cotin, S.: Image-guided simulation of heterogeneous tissue deformation for augmented reality during hepatic surgery. In: 2013 IEEE international symposium on mixed and augmented reality (ISMAR). pp. 199–208. IEEE (2013)
- Heiselman, J.S., Jarnagin, W.R., Miga, M.I.: Intraoperative correction of liver deformation using sparse surface and vascular features via linearized iterative boundary reconstruction. *IEEE transactions on medical imaging* **39**(6), 2223–2234 (2020)
- Li, Y., Harada, T.: Leopard: Learning partial point cloud matching in rigid and deformable scenes. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5554–5564 (2022)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
- Mestdagh, G., Cotin, S.: An optimal control problem for elastic registration and force estimation in augmented surgery. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 74–83. Springer (2022)
- Modrzejewski, R., Collins, T., Seeliger, B., Bartoli, A., Hostettler, A., Marescaux, J.: An in vivo porcine dataset and evaluation methodology to measure soft-body

- laparoscopic liver registration accuracy with an extended algorithm that handles collisions. *International journal of computer assisted radiology and surgery* **14**, 1237–1245 (2019)
9. Pfeiffer, M., Riediger, C., Leger, S., Kühn, J.P., Seppelt, D., Hoffmann, R.T., Weitz, J., Speidel, S.: Non-rigid volume to surface registration using a data-driven biomechanical model. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part IV* 23. pp. 724–734. Springer (2020)
  10. Rusu, R.B., Blodow, N., Beetz, M.: Fast point feature histograms (fpfh) for 3d registration. In: *2009 IEEE international conference on robotics and automation*. pp. 3212–3217. IEEE (2009)
  11. Soler, L., Hostettler, A., Agnus, V., Charnoz, A., Fasquel, J.B., Moreau, J., Osswald, A.B., Bouhadjar, M., Marescaux, J.: 3d image reconstruction for comparison of algorithm database (2010)
  12. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: *Symposium on Geometry processing*. vol. 4, pp. 109–116. Citeseer (2007)
  13. Talbot, H., Haouchine, N., Peterlik, I., Dequidt, J., Duriez, C., Delingette, H., Cotin, S.: Surgery Training, Planning and Guidance Using the SOFA Framework. In: Hege, H.C., Ropinski, T. (eds.) *Eurographics 2015 - Dirk Bartz Prize*. The Eurographics Association (2015). <https://doi.org/10.2312/egm.20151028>
  14. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6411–6420 (2019)
  15. Visser, A., Geboers, B., Gouma, D.J., Goslings, J.C., Ubbink, D.T.: Predictors of surgical complications: A systematic review. *Surgery* **158**(1), 58–65 (2015)
  16. World Health Organization: WHO Guidelines for Safe Surgery 2009: Safe Surgery Saves Lives. No. WHO/IER/PSP/2008.08-1E, World Health Organization (2009)
  17. Yang, Z., Simon, R., Linte, C.A.: Learning feature descriptors for pre-and intra-operative point cloud matching for laparoscopic liver registration. *International journal of computer assisted radiology and surgery* **18**(6), 1025–1032 (2023)
  18. Yew, Z.J., Lee, G.H.: Regtr: End-to-end point cloud correspondences with transformers. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6677–6686 (2022)
  19. Yu, H., Li, F., Saleh, M., Busam, B., Ilic, S.: Cofinet: Reliable coarse-to-fine correspondences for robust pointcloud registration. *Advances in Neural Information Processing Systems* **34**, 23872–23884 (2021)
  20. Zha, R., Cheng, X., Li, H., Harandi, M., Ge, Z.: Endosurf: Neural surface reconstruction of deformable tissues with stereo endoscope videos. In: *International conference on medical image computing and computer-assisted intervention*. pp. 13–23. Springer (2023)
  21. Zhang, Y., Zou, Y., Liu, P.X.: Point cloud registration in laparoscopic liver surgery using keypoint correspondence registration network. *IEEE Transactions on Medical Imaging* (2024)
  22. Zhao, Y., Zhang, J., Xu, S., Ma, J.: Deep learning-based low overlap point cloud registration for complex scenario: The review. *Information Fusion* **107**, 102305 (2024)