



Toward Patient-Specific Partial Point Cloud to Surface Completion for Pre to Intra-operative Registration in Image-Guided Liver Interventions

Nakul Poudel¹(✉), Zixin Yang¹, Kelly Merrell¹, Richard Simon², and Cristian A. Linte^{1,2}

¹ Center for Imaging Science, Rochester Institute of Technology, Rochester, NY 14623, USA
np1140@rit.edu

² Biomedical Engineering, Rochester Institute of Technology, Rochester, NY 14623, USA

Abstract. Intra-operative data captured during image-guided surgery lacks sub-surface information, where key regions of interest, such as vessels and tumors, reside. Image-to-physical registration enables the fusion of pre-operative information and intra-operative data, typically represented as a point cloud. However, this registration process struggles due to partial visibility of the intra-operative point cloud. In this research, we propose a patient-specific point cloud completion approach to assist with the registration process. Specifically, we leverage VN-OccNet to generate a complete liver surface from a partial intra-operative point cloud. The network is trained in a patient-specific manner, where simulated deformations from the pre-operative model are used to train the model. First, we conduct an in-depth analysis of VN-OccNet’s rotation-equivariant property and its effectiveness in recovering complete surfaces from partial intra-operative surfaces. Next, we integrate the completed intra-operative surface into the Go-ICP registration algorithm to demonstrate its utility in improving initial rigid registration outcomes. Our results highlight the promise of this patient-specific completion approach in mitigating the challenges posed by partial intra-operative visibility. The rotation equivariant and surface generation capabilities of VN-OccNet hold strong promise for developing robust registration frameworks for variations of the intra-operative point cloud.

Keywords: Image-guided liver surgery · point cloud completion · pre-to intra-operative registration

1 Introduction

Registration methods play a vital role during image-guided interventions in assisting surgeons to target the key regions of interest, such as tumors and ves-

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
S. Ali et al. (Eds.): MIUA 2025, LNCS 15917, pp. 302–316, 2026.
https://doi.org/10.1007/978-3-031-98691-8_22

sels, that lie beneath the organ surface. Image-to-physical registration aligns pre-operative Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) data, represented as a point cloud or mesh, to intra-operative data acquired using surgical tracking devices or cameras, often represented as a point cloud [4, 17]. However, the registration [3, 17] faces challenges caused by the partial intra-operative visibility that arises due to constrained camera viewpoints and occlusions [3, 5, 7]. Therefore, it is necessary to address the issue of partial visibility, which prevents registration methods from performing with sufficient accuracy.

The reconstruction of an intra-operative 3D liver surface from a partially visible surface has the potential to mitigate the partial visibility issue faced by registration methods. Toward this effort, Jia *et al.* [5] proposed a non-rigid registration framework that integrates a learning-based point completion network to generate a complete surface from sparse intra-operative data to guide non-rigid registration. However, the proposed method still requires a rigid registration as the initialization. Foti *et al.* [2] adopted a Variational Autoencoder on pre-operative models to generate a full liver surface from a partial point cloud, which undergoes an iterative optimization procedure to generate an intra-operative surface. However, this approach also requires manually identified correspondences between the generated mesh and the intra-operative point cloud, which is non-trivial.

In this work, we propose a patient-specific pipeline to improve image-to-physical registration (and hence pre- to intra-operative registration) by completing partial intra-operative liver surfaces. The pipeline leverages a vector neuron-based occupancy network (VN-OccNet [1]) to recover a complete liver mesh from a partial intra-operative point cloud. VN-OccNet offers two key advantages for the registration: (1) rotation equivariance, which addresses the failure of conventional models under varying orientations of intra-operative data [12], and (2) the ability to generate watertight meshes rather than point clouds, enabling uniform surface sampling—a crucial feature for registration methods requiring consistent point density [15, 17]. We adopt a patient-specific training strategy that synthesizes intra-operative liver surfaces by deforming a pre-operative patient-specific liver model, hence allowing the model to focus on patient-specific geometry and deformation patterns. We first evaluate VN-OccNet’s ability to generate rotation-equivariant surfaces across diverse rotation settings. We then compare registration outcomes using completed surfaces against those using only partial point clouds. Results show that the reconstructed surfaces significantly reduce registration error, highlighting the benefit of surface completion towards enhancing registration accuracy.

2 Methodology

2.1 Problem Definition

We define the intra-operative partial liver point cloud as the target point cloud $\mathcal{T} = \{t_i\}_{i=1}^N \in \mathbb{R}^3$ and the pre-operative liver point cloud as the source point cloud $\mathcal{S} = \{s_i\}_{i=1}^M \in \mathbb{R}^3$. Let $f_\theta(x_i, z)$ denote the trained VN-OccNet, where θ

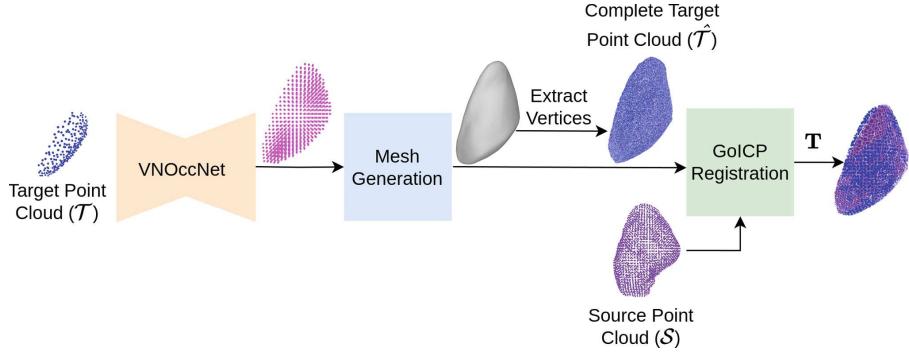


Fig. 1. Overview diagram of proposed pipeline: The target point cloud is input to the trained VN-OccNet to output occupancy points (magenta). An integrated mesh generation reconstructs a mesh from the predicted occupancy. The mesh vertices are extracted to represent the complete target point cloud and registered with a source point cloud. (Color figure online)

represents the learned parameters, x_i represents the i^{th} query point, and z is the latent representation of \mathcal{T} . The solution is to utilize $f_\theta(x_i, z)$ to reconstruct a complete estimate of the target point cloud, denoted as $\hat{\mathcal{T}}$, and subsequently perform registration with \mathcal{S} .

2.2 Proposed Pipeline

Our proposed pipeline illustrated in Fig. 1 consists of three key stages: (1) Generation of occupancy points from a target point cloud using VN-OccNet, (2) Mesh construction from the predicted occupancy points, and (3) Registration, where the complete target point cloud (obtained by extracting vertices from the generated mesh) is aligned with the source point cloud using the Go-ICP registration method.

Vector Neuron Occupancy Network (VN-OccNet): To accomplish the rotation-equivariant point cloud completion task, VN-OccNet [1] leverages the original OccNet architecture [10], with a vector-based equivariant encoder. The decoder is invariant, where all the operations are non-vector. The target point cloud $\mathcal{T} = \{t_i\}_{i=1}^N \in \mathbb{R}^3$ is input to the encoder network, and outputs equivariant latent vector-list features z . The query set $\mathcal{X} = \{x_i\}_{i=1}^K \in \mathbb{R}^3$, set of points that encloses the liver, and latent feature z are input to the decoder network to output the occupancy probabilities between 0 and 1. Points with probability greater than or equal to a threshold c correspond to occupancy points that implicitly represent the liver surface. Since VN-OccNet aims to predict whether or not each point in the query point set corresponds to a liver surface point, binary cross-entropy classification loss is utilized to learn the parameters θ .

Mesh Generation: Meshes are constructed by first identifying the voxels that intersect the liver surface using the Multiresolution Isosurface Extraction (MISE) algorithm [10]. The algorithm starts with a fixed query set with an initial resolution of 32^3 , consisting of 31^3 voxels, each containing 8 points. For all 32^3 points in the query set, the occupancy probability is determined using the trained VN-OccNet. The points inside the surface are identified by applying a threshold value of $c = 0.4$. Voxels that are entirely inside or outside the surface are separated from the surface-intersecting voxels by checking if at least two adjacent points of a voxel differ in occupancy (i.e., one point is inside and the other is outside). Each such ambiguous voxel that intersects the surface is further subdivided into 8 subvoxels, generating 19 additional points. The VN-OccNet is queried again to find the occupancy probability for the newly generated query points. This process is repeated until all ambiguous regions have been subdivided down to the finest voxel grid resolution of 128^3 . Finally, the surface-intersecting voxels obtained are passed to the Marching Cubes algorithm [8] to generate the surface mesh.

Rigid Registration Using Go-ICP: To evaluate the effectiveness of incorporating the complete intra-operative target surface generated from using the proposed protocol into the registration, we utilize the Go-ICP registration algorithm [15] to compare the registration outcomes between using the initial, incomplete target point clouds and the complete intra-operative surface meshes generated using the proposed method.

For the remainder of the manuscript, we refer to these registrations as registration w/ surface completion (when the complete surface mesh is utilized) and registration w/o surface completion (when only the original partial point cloud is utilized). The registration algorithm yields the transformation (\mathbf{T}), which aligns the source point cloud to the target point cloud.

3 Experiments

3.1 Datasets

The experiments utilize the following datasets: an *in silico* phantom dataset is used to train and test VN-OccNet, and an *in vitro* phantom dataset is used to assess the registration performance.

In Silico Phantoms: We used the undeformed No. 1 synthetic phantom constructed based on a patient-specific CT image according to the methods proposed by Yang *et al.* [18] briefly described in further detail in the next sub-section. To simulate deformed models, the deformation simulation pipeline described in [11] is followed, which models the liver as a neo-Hookean hyperelastic material with a Young's modulus of 2–5 kPa and a Poisson's ratio of 0.35. Three forces of magnitude up to 3N and zero boundary conditions, along with the material properties, were input to the finite element solver to output deformed liver models. The complete target point clouds are obtained by extracting the mesh vertices. In total, 4,969 deformed models were generated.

Training an occupancy network requires transforming liver models into a representation comprising a liver point cloud, a set of query points enclosing the liver, and corresponding occupancy values that indicate whether each query point lies inside or outside the liver surface. We followed the pipeline of Stutz *et al.* [12] to generate this data format. The generated data was divided into training, validation, and testing sets with a ratio of 8:1:1.

In Vitro Phantoms: We used the *in vitro* phantom dataset described and released in [18], which includes four pairs of underformed and deformed models. Phantoms were manufactured using synthetic gelatin (Humimic Gelatin # 0, Humimic Medical, Fort Smith, AK, USA) poured into a 3D-printed mold generated [9] from a patient-specific liver model obtained from OpenHELP [6]. They were deformed by placing wedges of different gradients underneath part of the liver phantom. The surfaces and fiducial marker locations are segmented from CT scans. As Go-ICP requires a long time to estimate rigid registration, we utilized No. 1 and No. 3 phantoms within the dataset in this work, as shown in Fig. 2. No. 1 and No. 3 phantoms feature 53 and 176 fiducials, respectively.

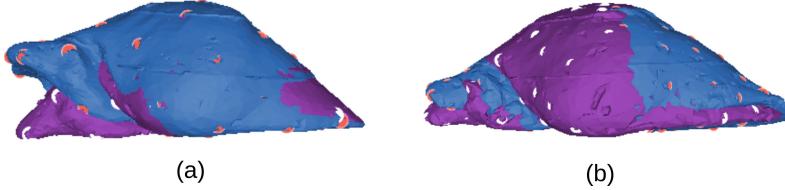


Fig. 2. The purple phantoms in (a) and (b) represent the undeformed No. 1 and phantom No. 3, respectively, while the blue phantoms represent their deformed counterparts. (Color figure online)

3.2 VN-OccNet Training and Testing Setup

Target Point Cloud Generation: In training, target point clouds are dynamically generated from the complete point clouds. First, we crop the posterior surfaces of the deformed liver models. The anterior surface points are then down-sampled to D points. A viewpoint is randomly selected, and the N nearest surface points relative to this viewpoint are extracted to create the target point cloud. For testing, we follow the approach proposed in [19] to simplify the qualitative analysis, where five distinct viewpoints are selected for each liver model, as illustrated in Fig. 3.

Rotation Setup: VN-OccNet is trained in three different modes based on the rotation imposed on the target point clouds: without any imposed rotations, with rotations restricted to the Z-axis, and with SO(3) rotations. We imposed random rotations within the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$ for Z-axis and SO(3) rotation modes.

In the non-rotation mode, all target point clouds are aligned and share the same pose. The evaluation involves four train/test rotation settings: I/I, where both training and testing are conducted on non-rotated target point clouds; Z/Z, where both training and testing involve rotations of target point clouds along the Z-axis; Z/SO(3), where training is performed with Z-axis rotations, while testing involves SO(3) rotations; and SO(3)/SO(3), where both training and testing involve SO(3) rotations.

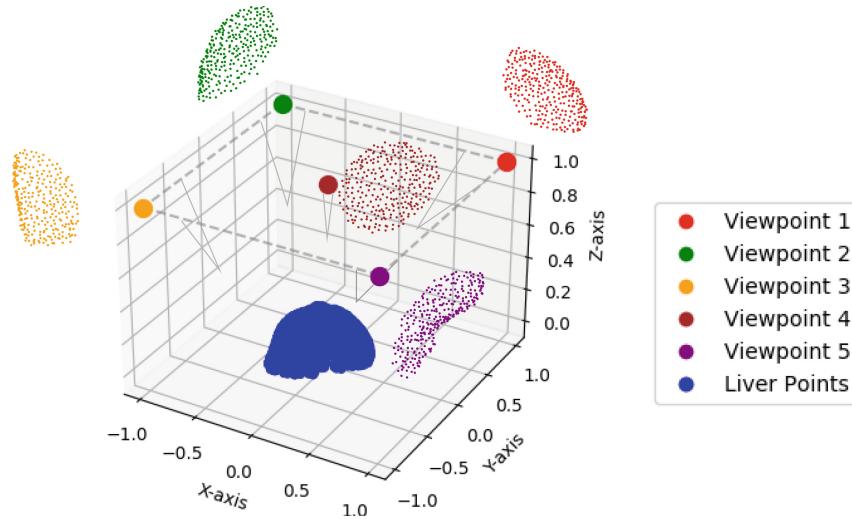


Fig. 3. Visualization of five target point clouds generated for each intra-operative liver model during testing. Each target is captured from a unique viewpoint, with different regions of the liver obscured to mimic varying intra-operative visibility.

3.3 VN-OccNet Implementation

We employed the PyTorch implementation of the VN-OccNet by Deng *et al.* [1]. The anterior liver points were downsampled to $D = 1000$, and $N = 300$ nearest points from the viewpoint were chosen to create the target liver point cloud, yielding a surface visibility of approximately 30%. The network was trained with a batch size of $B = 8$ and a learning rate of 0.0001, using the Adam optimizer. The optimal threshold value c was determined to be 0.4 based on the validation dataset. Training was performed for 604 epochs on a NVIDIA A100 GPU, with all other hyperparameters kept unchanged from the original implementation. To verify the rotation equivariance property, we further compared VN-OccNet with OccNet [10], which does not include rotation equivariance design.

3.4 Registration Evaluation Setup

In a manner similar to the five target point clouds generated for each liver model during testing in the previous Subsect. 3.2, five target point clouds are generated for each *in vitro* phantom (No. 1 and No. 3) to perform registration. A set of 100 random $SO(3)$ rotations within the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$ are applied to the target point clouds, resulting in 500 target point clouds at different orientations for each phantom. Registration w/ surface completion is performed by utilizing the complete target point clouds reconstructed from target point clouds featuring a minimum of 30% visibility. To test the influence of the increment of surface coverage, registration w/o surface completion was performed at three different visibilities of the target point cloud at around 30%, 40%, and 50% by varying the value of D , where $D = 1000, 750$, and 600, respectively.

3.5 Evaluation Metrics

Chamfer Distance (CD-L2), F-Score, and Intersection over Union (IoU) were used to evaluate the performance of complete target surface generation. Moreover, the Target Registration Error (TRE) computed across a set of fiducials was used to assess registration performance.

CD-L2 measures the bidirectional similarity between the complete generated target and ground truth target points:

$$d_{CD}(\hat{\mathcal{T}}, \mathcal{G}) = \frac{1}{|\hat{\mathcal{T}}|} \sum_{\hat{\mathbf{t}} \in \hat{\mathcal{T}}} \min_{\mathbf{g} \in \mathcal{G}} \|\hat{\mathbf{t}} - \mathbf{g}\|_2 + \frac{1}{|\mathcal{G}|} \sum_{\mathbf{g} \in \mathcal{G}} \min_{\hat{\mathbf{t}} \in \hat{\mathcal{T}}} \|\mathbf{g} - \hat{\mathbf{t}}\|_2, \quad (1)$$

where $|\hat{\mathcal{T}}|$ and $|\mathcal{G}|$ represent the total number of points in the complete generated target and ground truth target point cloud, respectively, and $d_{CD}(\hat{\mathcal{T}}, \mathcal{G})$ denotes the Chamfer Distance between the point sets $\hat{\mathcal{T}}$ and \mathcal{G} .

The F-Score measures the quality of the reconstructed surface by balancing precision $P(d)$ and recall $R(d)$. Precision measures the fraction of complete target points that lie within a threshold $d = 1$ mm of the nearest ground truth target points, and recall measures the fraction of ground truth target points that lie within d of the nearest complete target points:

$$\text{F-Score}(d) = \frac{2P(d)R(d)}{P(d) + R(d)}. \quad (2)$$

IoU measures the ratio of the intersection to the union of the occupancy values of the generated and ground truth targets, as defined in Eq. (3):

$$\text{IoU} = \frac{|\text{occ}_{\text{gt}} \cap \text{occ}_{\text{pred}}|}{|\text{occ}_{\text{gt}} \cup \text{occ}_{\text{pred}}|}, \quad (3)$$

where occ_{gt} and occ_{pred} represent the ground truth and predicted occupancy values, respectively, and $|\cdot|$ denotes the number of true occupancy labels.

TRE measures the Euclidean distance between the estimated target fiducial markers $\mathbf{F}_{tgt,i}$ and the transformed source fiducial markers $\mathbf{T}(\mathbf{F}_{src,i})$:

$$TRE = \frac{1}{F} \sum_{i=1}^F \|\mathbf{F}_{tgt,i} - \mathbf{T}(\mathbf{F}_{src,i})\|_2, \quad (4)$$

where F is the total number of fiducial markers.

4 Results

4.1 Mesh Construction Results

Table 1 presents the performance of the intra-operative surface completion for OccNet and VN-OccNet across four different train/test configurations: I/I, Z/Z, Z/SO(3), and SO(3)/SO(3). OccNet achieves the best performance across all metrics in the I/I setting. However, its performance significantly deteriorates when tested with target liver point clouds with rotations along the Z-axis and SO(3).

Table 1. Performance of OccNet and VN-OccNet: I/I represents training and testing without any rotations of target point cloud, Z/Z denotes training and testing with rotations along the Z-axis, Z/SO(3) involves training with Z-axis rotations and testing with SO(3) rotations, and SO(3)/SO(3) refers to both training and testing with SO(3) rotations. CD-L2 and IoU represent the Chamfer Distance (in mm) and Intersection over Union.

	I/I	Z/Z	Z/SO(3)	SO(3)/SO(3)
OccNet				
F-Score	0.57 ± 0.20	0.47 ± 0.17	0.07 ± 0.07	0.33 ± 0.12
CD-L2	2.75 ± 1.52	3.51 ± 1.88	29.43 ± 13.41	4.87 ± 2.27
IoU	0.89 ± 0.06	0.86 ± 0.07	0.35 ± 0.18	0.18 ± 0.07
VN-OccNet				
F-Score	0.52 ± 0.17	0.51 ± 0.18	0.51 ± 0.18	0.50 ± 0.18
CD-L2	3.30 ± 1.91	3.29 ± 1.89	3.25 ± 1.83	3.41 ± 2.02
IoU	0.87 ± 0.07	0.87 ± 0.07	0.87 ± 0.06	0.86 ± 0.07

While data augmentation along SO(3) improves performance to some extent, it remains impractical to encode all possible SO(3) rotations during training, making the model susceptible to failure when encountering unseen liver poses during testing. Although VN-OccNet does not outperform OccNet in the I/I setting, it maintains consistent performance across Z/Z, Z/SO(3), and SO(3)/SO(3), demonstrating robustness to rotations.

Unlike OccNet, VN-OccNet does not show performance improvement due to data augmentation, indicating that its inherent rotation-equivariant properties reduce reliance on augmented training data. The statistical significance

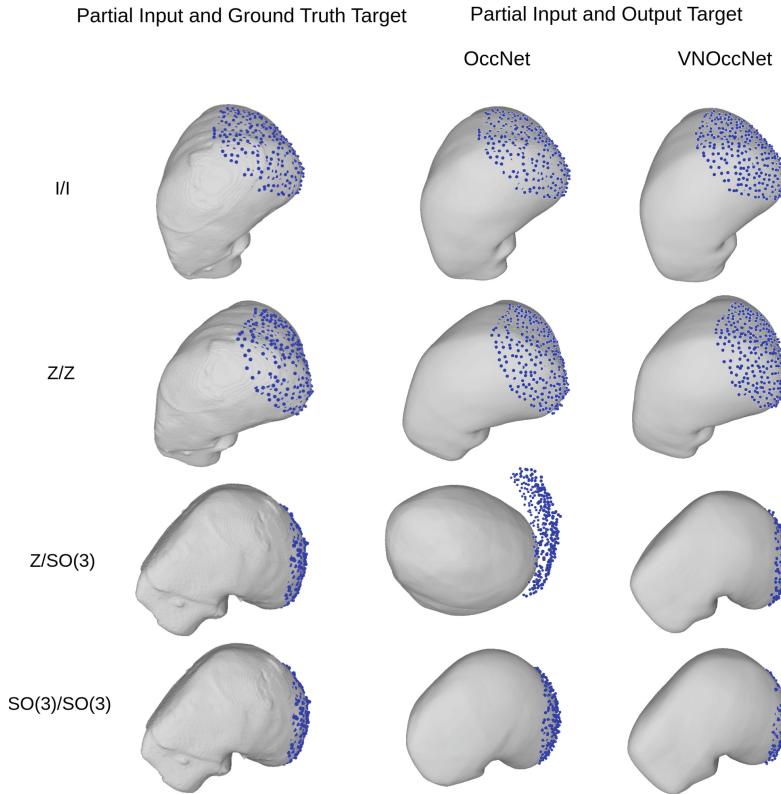


Fig. 4. Visualization of mesh reconstruction results: The first column displays the target input point clouds alongside the ground truth target mesh. The second column presents the complete target liver surfaces generated by OccNet and VN-OccNet. Qualitative results are shown for four train/test configurations: I/I (first row), where both training and testing involve non-rotated target point clouds; Z/Z (second row), where training and testing are restricted to Z-axis rotations; Z/SO(3) (third row), where training uses Z-axis rotations and testing employs SO(3) rotations; and SO(3)/SO(3) (fourth row), where both training and testing involve SO(3) rotations.

of OccNet and VN-OccNet performance was evaluated by comparing all three evaluation metrics using a Wilcoxon Rank Sum test at a significance level of $\alpha = 0.05$. The test yielded p-values less than α ($p < \alpha$), confirming that the performance differences between VN-OccNet and OccNet across various train/test setups are statistically significant.

Figure 4 presents the target point clouds, ground truth target mesh surfaces, and complete target mesh surfaces generated by OccNet and VN-OccNet across four different train/test configurations: I/I, Z/Z, Z/SO(3), and SO(3)/SO(3). The performance gap in surface generation is evident in the Z/SO(3) case, where OccNet struggles to generate a target mesh surface that accurately resembles the

ground truth target mesh. In the SO(3)/SO(3) setting, the mesh generated by OccNet lacks the correct geometry, indicating an incorrect reconstruction. In contrast, the mesh generated by VN-OccNet is able to capture correct geometry and better alignment of the target point cloud with the reconstructed mesh.

4.2 Go-ICP Registration Results

Table 2 summarizes the Target Registration Error (in mm) measured between target fiducial markers and source fiducial markers. The measurements are reported for registration w/ and w/o surface completion. The registration w/o surface completion is performed using the target point clouds at three visibility levels (30%, 40%, and 50%), and registration w/ surface is performed using the completed surface meshes generated from the 30% visible target point clouds.

Table 2. Target Registration Error (TRE) (in mm) measured between target fiducial markers and transformed source fiducial markers for phantoms No. 1 and No. 3. For registration w/o surface completion, TRE is evaluated at intra-operative visibility levels of 30%, 40%, and 50%. For registration w/ surface completion, the full surfaces are reconstructed from the target point clouds with 30% visibility.

Phantom	W/O surface completion			W/ surface completion
	30% Visibility	40% Visibility	50% Visibility	
No. 1	33.13 ± 13.89	30.39 ± 12.60	26.04 ± 7.81	5.19 ± 1.34
No. 3	39.87 ± 15.07	32.79 ± 14.97	25.87 ± 10.36	3.35 ± 0.61

The registration for phantom No. 1 yields the TRE of 33.13 ± 13.89 mm for the 30% visibility of target point clouds. However, when using the complete target point clouds generated from the same 30% visible target, the TRE significantly decreased to 5.19 ± 1.34 mm. This notable improvement demonstrates the effectiveness of VN-OccNet’s surface completion in enhancing registration performance. Furthermore, as the visibility of the target point clouds increased, TRE progressively decreased, confirming that a lower visibility ratio leads to higher registration errors.

A Wilcoxon Rank Sum test, conducted at a significance level of $\alpha = 0.05$, showed statistically significant differences ($p < \alpha$) between the three cases of registration w/o surface completion and registration w/ surface completion.

Figure 5 displays the results for registration w/ surface completion and w/o surface completion for *in vitro* phantom No. 1 and No. 3. The registration w/o surface completion struggles to find the correct alignment between the source and target point clouds. The incorrect registration is evident when the complete ground truth target is overlaid onto the registered results. However, registration w/ surface completion is able to correctly align the source and target point clouds.

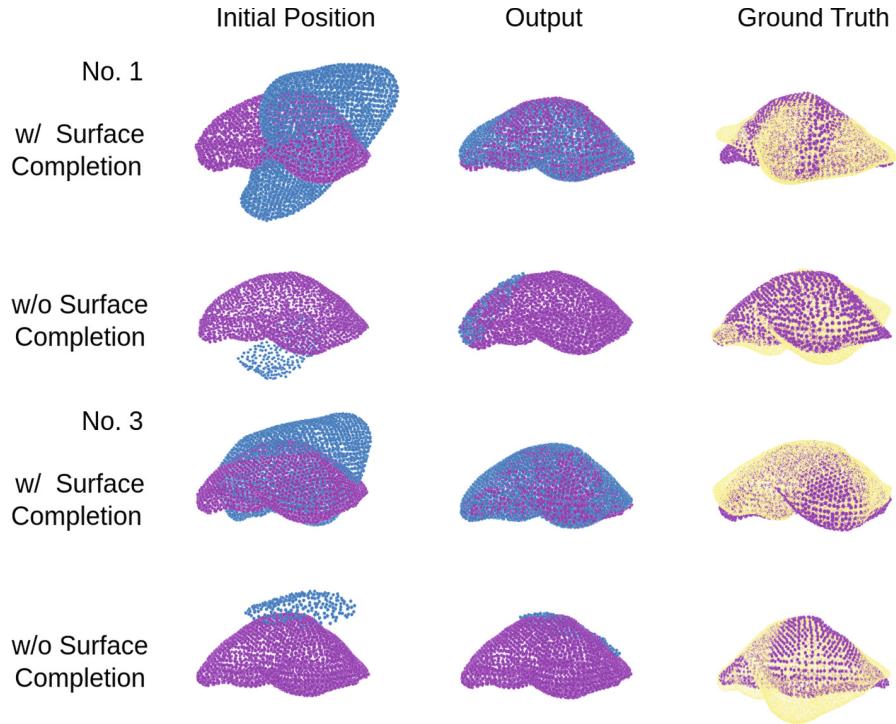


Fig. 5. The registration results are displayed for the source point cloud (purple) and the target point cloud (blue) for phantoms No. 1 and No. 3. For easy visualization, the source point cloud is fixed and the target point clouds are moving. The first column represents the initial position of the source and target point clouds. The second column displays the registered results. To see the correctness of the registration, the third column overlays the ground truth target point cloud (yellow). The registrations are performed using partial surface (30% visibility) and full surface generated from the same partial surface. (Color figure online)

Similarly, Fig. 6 shows the registration between source and target fiducial markers for phantom No. 1 and phantom No. 3. The transformations were determined using a target point cloud with 30% visibility for registration w/o surface completion and a reconstructed completed surface mesh from the same target point cloud for registration w/ surface completion. For the registration w/o surface completion, the registered source fiducial markers are misaligned with respect to the target fiducial markers. In contrast, in registration w/ complete surface, the source and target fiducial markers are closely aligned, demonstrating improved registration.

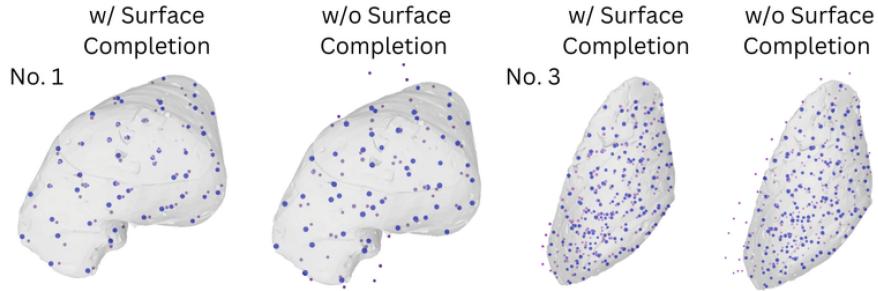


Fig. 6. Visualization of target fiducial markers (purple) and registered source fiducial markers (blue) for *in vitro* phantom No. 1 and No. 3. The target point clouds utilized for registration w/o surface completion have 30% visibility, and the full point clouds are reconstructed from the same targets for registration w/ surface completion. The ground truth target liver surface is overlaid for improved visualization. (Color figure online)

5 Discussion

The partial visibility of intra-operative point clouds poses challenges for both initial rigid and non-rigid registration methods in image-guided liver surgery. To address this issue, we propose a patient-specific, learning-based surface completion approach. In this work, we initially validate the effectiveness of this approach in facilitating the initial rigid registration.

Initial rigid registration has relied on global optimization techniques [15] and the establishment of reliable correspondences using handcrafted [13] or learning-based methods [17]. However, these approaches are often sensitive to partial liver surfaces and inconsistent point cloud densities. In contrast, our method tackles the problem at its root by employing surface completion through learning-based techniques. Unlike previous completion approaches that require rigid initialization [5] or manual interaction [2], our method operates without such dependencies. We utilize a vector-based occupancy network—VN-OccNet—due to its robustness to rotation and its ability to generate watertight meshes in just 0.25 seconds per mesh, making it a strong candidate for this task.

We first verify the rotation-equivariance property of VN-OccNet, which is essential for our application. Recent work [12] evaluated several widely used point completion methods and found them unsuitable for this task. We compared VN-OccNet with the original OccNet, which lacks a rotation-invariant design, under various training and testing conditions, using extensive quantitative and qualitative evaluations. Furthermore, we assessed registration performance by integrating completed surfaces with Go-ICP, a method known for its robustness to partial visibility [16]. Results show that our reconstructed surfaces significantly reduce registration errors, demonstrating the value of surface completion in improving registration accuracy. However, Go-ICP, which takes approximately 16 s per registration, is computationally intensive, suggesting that integrating

surface completion directly into learning-based registration methods may offer a more efficient solution.

An additional challenge lies in the variation of liver geometries across different patients, which can hinder surface completion. To overcome this, we adopt a patient-specific training strategy that simulates intra-operative surfaces by deforming a pre-operative patient-specific liver model. This approach enables the network to learn patient-specific geometries and deformation patterns. Our results show that this strategy allows the model to generalize effectively to realistic deformations observed in *in vitro* phantoms. Nonetheless, the requirement to train a new model for each patient may limit its practicality in resource-constrained environments.

It is also important to note that real intra-operative data may exhibit more complex patterns, such as noise, holes, and occlusions, than those seen in our simulations. This discrepancy poses a challenge for learning-based methods, which generally assume the test data distribution is similar to that of the training data. As such, improving generalization to diverse, realistic intra-operative point clouds remains an open challenge.

6 Conclusion and Future Work

This paper presents a pipeline that incorporates VN-OccNet to generate a complete liver surface from partial intra-operative point cloud data, using the reconstructed surface for subsequent registration. We demonstrate the effectiveness of VN-OccNet in producing accurate and complete surface reconstructions, which significantly enhance registration performance. While partial intra-operative data typically hampers the accuracy of most registration methods, our results indicate that the proposed pipeline provides a promising solution to this challenge.

In future work, we plan to investigate the integration of surface completion with non-rigid registration, explore patient-generic surface completion strategies, incorporate completion into end-to-end learning-based registration pipelines, and develop a more robust simulation pipeline capable of generating realistic intra-operative point clouds with noise, holes, and occlusions.

Acknowledgments. We would like to acknowledge the generous support for this work by the National Institutes of Health – National Institute of General Medical Sciences under Award No. R35GM128877 and the National Science Foundation - Division of Chemical, Bioengineering and Transport Systems under Award No. 2245152. We would like to thank the Research Computing team at Rochester Institute of Technology [14] for providing computing resources for this research. We gratefully acknowledge Bidur Khanal for his insightful feedback and encouragement throughout this work.

References

1. Deng, C., Litany, O., Duan, Y., Poulenard, A., Tagliasacchi, A., Guibas, L.J.: Vector neurons: a general framework for so (3)-equivariant networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 12200–12209 (2021)
2. Foti, S., et al.: Intraoperative liver surface completion with graph convolutional VAE. In: Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis: Second International Workshop, UNSURE 2020, and Third International Workshop, GRAIL 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, 8 October 2020, Proceedings 2, pp. 198–207. Springer (2020)
3. Heiselman, J.S., et al.: Characterization and correction of intraoperative soft tissue deformation in image-guided laparoscopic liver surgery. *J. Med. Imaging* **5**(2), 021203 (2018)
4. Heiselman, J.S., Collins, J.A., Ringel, M.J., Peter Kingham, T., Jarnagin, W.R., Miga, M.I.: The image-to-physical liver registration sparse data challenge: comparison of state-of-the-art using a common dataset. *J. Med. Imaging* **11**(1), 015001 (2024)
5. Jia, M., Kyan, M.: Improving intraoperative liver registration in image-guided surgery with learning-based reconstruction. In: ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1230–1234. IEEE (2021)
6. Kenngott, H.G., et al.: OpenHELP (Heidelberg laparoscopy phantom): development of an open-source surgical evaluation and training tool. *Surg. Endosc.* **29**, 3338–3347 (2015)
7. Lin, B., Sun, Y., Qian, X., Goldgof, D., Gitlin, R., You, Y.: Video-based 3D reconstruction, laparoscope localization and deformation recovery for abdominal minimally invasive surgery: a survey. *Int. J. Med. Robot. Comput. Assist. Surg.* **12**(2), 158–178 (2016)
8. Lorensen, W.E., Cline, H.E.: Marching cubes: a high resolution 3D surface construction algorithm. In: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987, pp. 163–169. Association for Computing Machinery, New York (1987). <https://doi.org/10.1145/37401.37422>
9. Merrell, K., Jackson, P., Simon, R., Linte, C.: Developing and evaluating the fidelity of patient specific kidney emulating phantoms for image-guided intervention applications. In: Medical Imaging 2023: Image-Guided Procedures, Robotic Interventions, and Modeling, vol. 12466, pp. 318–323. SPIE (2023)
10. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: learning 3D reconstruction in function space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4460–4470 (2019)
11. Pfeiffer, M., et al.: Non-rigid volume to surface registration using a data-driven biomechanical model. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12264, pp. 724–734. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59719-1_70
12. Poudel, N., Yang, Z., Merrell, K., Simon, R., Linte, C.A.: Evaluation of intraoperative patient-specific methods for point cloud completion for minimally invasive liver interventions. In: Medical Imaging 2025: Image-Guided Procedures, Robotic Interventions, and Modeling, vol. 13408, pp. 376–384. SPIE (2025)

13. Robu, M.R., et al.: Global rigid registration of CT to video in laparoscopic liver surgery. *Int. J. Comput. Assist. Radiol. Surg.* **13**, 947–956 (2018)
14. Rochester Institute of Technology: Research computing services (2025). <https://doi.org/10.34788/0S3G-QD15>, <https://www.rit.edu/researchcomputing/>
15. Yang, J., Li, H., Campbell, D., Jia, Y.: Go-ICP: a globally optimal solution to 3D ICP point-set registration. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(11), 2241–2254 (2015)
16. Yang, Z., et al.: Resolving the ambiguity of complete-to-partial point cloud registration for image-guided liver surgery with patches-to-partial matching. arXiv preprint [arXiv:2412.19328](https://arxiv.org/abs/2412.19328) (2024)
17. Yang, Z., Simon, R., Linte, C.A.: Learning feature descriptors for pre-and intraoperative point cloud matching for laparoscopic liver registration. *Int. J. Comput. Assist. Radiol. Surg.* **18**(6), 1025–1032 (2023)
18. Yang, Z., Simon, R., Merrell, K., Linte, C.A.: Boundary constraint-free biomechanical model-based surface matching for intraoperative liver deformation correction. *IEEE Trans. Med. Imaging* **1** (2024). <https://doi.org/10.1109/TMI.2024.3515632>
19. Yu, X., Rao, Y., Wang, Z., Liu, Z., Lu, J., Zhou, J.: PoinTr: diverse point cloud completion with geometry-aware transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 12498–12507 (2021)