

Tracking Tumors under Deformation from Partial Point Clouds using Occupancy Networks

Pit Henrich^{1†}, Jiawei Liu^{2†}, Jiawei Ge², Samuel Schmidgall², Lauren Shepard³,
Ahmed Ezzat Ghazi³, Franziska Mathis-Ullrich¹, Axel Krieger²

Abstract— To track tumors during surgery, information from preoperative CT scans is used to determine their position. However, as the surgeon operates, the tumor may be deformed which presents a major hurdle for accurately resecting the tumor, and can lead to surgical inaccuracy, increased operation time, and excessive margins. This issue is particularly pronounced in robot-assisted partial nephrectomy (RAPN), where the kidney undergoes significant deformations during operation. Toward addressing this, we introduce an occupancy network-based method for the localization of tumors within kidney phantoms undergoing deformations at interactive speeds. We validate our method by introducing a 3D hydrogel kidney phantom embedded with exophytic and endophytic renal tumors. It closely mimics real tissue mechanics to simulate kidney deformation during in vivo surgery, providing excellent contrast and clear delineation of tumor margins to enable automatic threshold-based segmentation. Our findings indicate that the proposed method can localize tumors in moderately deforming kidneys with a margin of 6mm to 10mm, while providing essential volumetric 3D information at over 60Hz. This capability directly enables downstream tasks such as robotic resection.

I. INTRODUCTION

Kidney cancer is one of the most common forms of cancer in the US, with over 65,000 new patients being diagnosed every year, leading to over 15,000 deaths [1]. The standard treatment for localized small renal masses has shifted from radical nephrectomy (complete kidney removal) toward the more minimally invasive approach of partial nephrectomy (removal of the tumor, retaining partial kidney function). One of the main challenges during tumor removal is ensuring the resection of adequate tumor margins. The margin needs to ensure that no cancer cells remain in the kidney, while also avoiding excessive removal of healthy tissue to preserve organ function.

For operations involving tumor resection, surgeons commonly interpret 2D preoperative scans and mentally construct

*Research reported in this paper was supported by the Advanced Research Projects Agency for Health (ARPA-H) under grant number AY1AX000023, and National Science Foundation (NSF/FRR CAREER) under grant number 2144348. The content is solely the responsibility of the authors and does not necessarily represent the official views of the Advanced Research Projects Agency for Health or the National Science Foundation.

[†]These authors contributed equally.

¹P. Henrich and F. Mathis-Ullrich are with the Department Artificial Intelligence in Biomedical Engineering (AIBE), Friedrich-Alexander University Erlangen-Nürnberg (FAU), 91052 Erlangen, Germany. {pit.henrich, franziska.mathis-ullrich}@fau.de

²J. Liu, J. Ge, S. Schmidgall and A. Krieger are with the Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218, USA {jliu298, jge9, sschmi46, axel}@jhu.edu

³L. Shepard and A.E. Ghazi are with the Department of Urology, Johns Hopkins University, Baltimore, MD 21218, USA {lshepar9, aghazil}@jh.edu

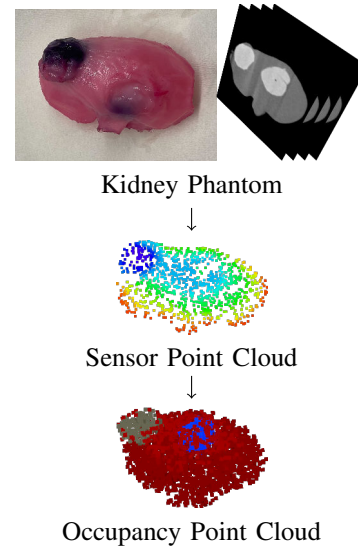


Fig. 1: Given a known kidney phantom with two tumors and a pre-operative CT scan, we estimate, from a single sensor point cloud derived from a depth image, a dense occupancy point cloud that encodes the locations of each tumor.

a 3D anatomical model during the surgery. This process of creating a mental model can extend the duration of the procedure, increasing the risk of damaging adjacent tissues [2]. The difficulty is exacerbated by tissue deformation during surgery, causing the internalized mental model to become inaccurate over time.

Registration methods that overlay pre-operative data onto the intra-operative scene can provide surgeons with valuable location information for regions of interest. Altamar et al. [3] develop a deformable registration method based on biomechanical elastic models to simulate in-vivo deformation during partial nephrectomy. While their method accurately analyzes kidney deformation, it doesn't address the tumor and margin deformation.

Point-based registration enables the calculation of precise transformations between markers in pre-operative images and those in the physical space, as shown on intra-operative images. In the absence of distinct anatomical landmarks on organs like the kidney, applying fiducials to the renal surface is advocated for point-based registration [4], [5]. Nimmagadda et al. [6] develop a touch-based registration approach using ink fiducials tattooed on the kidney surface. They employ the Iterative Closest Point (ICP) algorithm for

initial alignment with a pre-operative point cloud, followed by point-based re-registration to account for kidney deformation. While precise, tattooing intra-operative tissue has been associated with increased patient morbidity [2]. Zhang et al. [7] propose a different approach. They gather surface data by capturing multiple views of the kidney, stitching them together to construct a 3D surface. They then apply Coherent Point Drift (CPD) for deformable registration. This technique achieves markerless deformable registration but requires multiple views of the organ, wherein the organ’s deformation during the image collection process could introduce errors.

In previous work [8], we demonstrate a method for reconstructing deformed objects composed of multiple parts from single-viewpoint point clouds using a neural multi-class occupancy function. This approach utilizes a combination of occupancy networks [9] and PointNet++ [10] to infer a 3D object that aligns with the observation of the deformable real-world object. It relies on deforming a pre-operative 3D model digitally to generate ground truth occupancy samples and a sensor-based depth image. Through supervised learning, the network is trained to label occupancy samples correctly, learning the adaptation of the labelling to observed deformations directly from the data. A distinct advantage of this method is its ability to estimate the position of deformable parts without necessitating initial registration.

To validate the accuracy of deformable registration and localization methods, organ “phantoms” are widely used, which aim to replicate the mechanical properties and medical imaging results of real organs. Birnbaum et al. [11] customized renal inserts with different Hounsfield Unit (HU) values in anthropomorphic phantoms to produce varying brightness between kidney components under CT imaging. However, the analysis of mechanical tissue properties was not undertaken. In our previous work, we develop a hydrogel kidney phantom capable of reproducing both the mechanical and functional properties of living tissue [12]. This phantom is specifically designed for Robot-Assisted Partial Nephrectomy (RAPN) training, with a focus on mechanical fidelity over CT imaging characteristics.

A method that can register or reconstruct deformable objects from an observation, can be applied to perform resection tasks autonomously. We recently presented the Autonomous System for Tumor Resection (ASTR) [13] method. It is a vision-guided robotic system that demonstrates success in tongue tumor resection (i.e. glossectomy).

In this work, we present an occupancy network-based tumor localization approach designed to use pre-operative CT images with intra-operative RGBD sensor data. Our method only requires a single depth image from a single viewpoint to estimate the deformation of the kidney and locate embedded tumors at over 60Hz, see Figure 1. Therefore, it can track tumors, enabling the delineation of margins in the presence of deformations occurring during resection. Furthermore, we evaluate our method through the use of a novel 3D hydrogel kidney phantom, which is embedded with renal tumors. This phantom possesses realistic mechanical properties and allows easy automatic segmentation by providing varying

brightnesses under CT imaging. The evaluation demonstrates the utility of our method in guiding a robotic resection, accounting for potential deformations during partial nephrectomy procedures.

II. METHODS

An overview of our method is shown in Figure 2. Here, we use an occupancy network which estimates the location of tumors based on a single-view point cloud of the kidney surface. The aim of our technique is to enable the localization to work in the presence of deformations. We make use of occupancy networks because they can directly encode deformable smaller structures inside of larger structures, such as a tumor inside of the kidney. This is in contrast to surface mesh based representations, which only provides a clear description of the surface of objects. Occupancy networks, therefore, directly provide 3D location and shape information of structures needed for autonomous resection. This innate suitability for robotic resection tasks is verified by integrating the occupancy-generated point cloud into an autonomous robot system, and virtually executing a resection plan using the robot.

For this, a pre-operative model of the kidney with tumors is deformed virtually and used to train an occupancy network. In training, the occupancy network learns to estimate the 3D shape of the kidney based on a single depth image, re-projected into 3D space as a point cloud. We will refer to the point cloud obtained from the depth image as a Sensor Point Cloud (SPC). Conditioned on a real-world SPC, the occupancy network produces registered volumetric information about the intra-operative kidney. The intra-operative localization of the tumors inside of a deformable kidney can be performed at over 60Hz on consumer hardware. This enables an interactive tracking of tumors during surgical applications.

A. Localization

To locate the kidney tumors we use our deformable object reconstruction method as previously proposed [8], with all losses and hyper-parameters remaining the same. To elaborate, PointNet++ [10] is used as the encoder to distill a 1024 dimensional latent representation. The latent vector is appended with a query point and passed to a fully connected Multilayer Perceptron (MLP) with 8 hidden layers of 512 neurons. A skip connection is added to fifth hidden layer. ReLU is used as an activation function and batch normalization is used for the skip and the final layer. The network predicts the label for each query point using an argmax operation on the output logits, and the entire system is optimized end-to-end with a cross-entropy loss function.

Pre-Operative Training: Our method requires training data to be created to perform the localization. For this, a 3D surface mesh of the kidney and the tumors is required. For example, these models can be obtained using pre-operative Computed Tomography (CT) images. An example for this is given in Figure 2.

Pre-Operative CT Imaging and Training

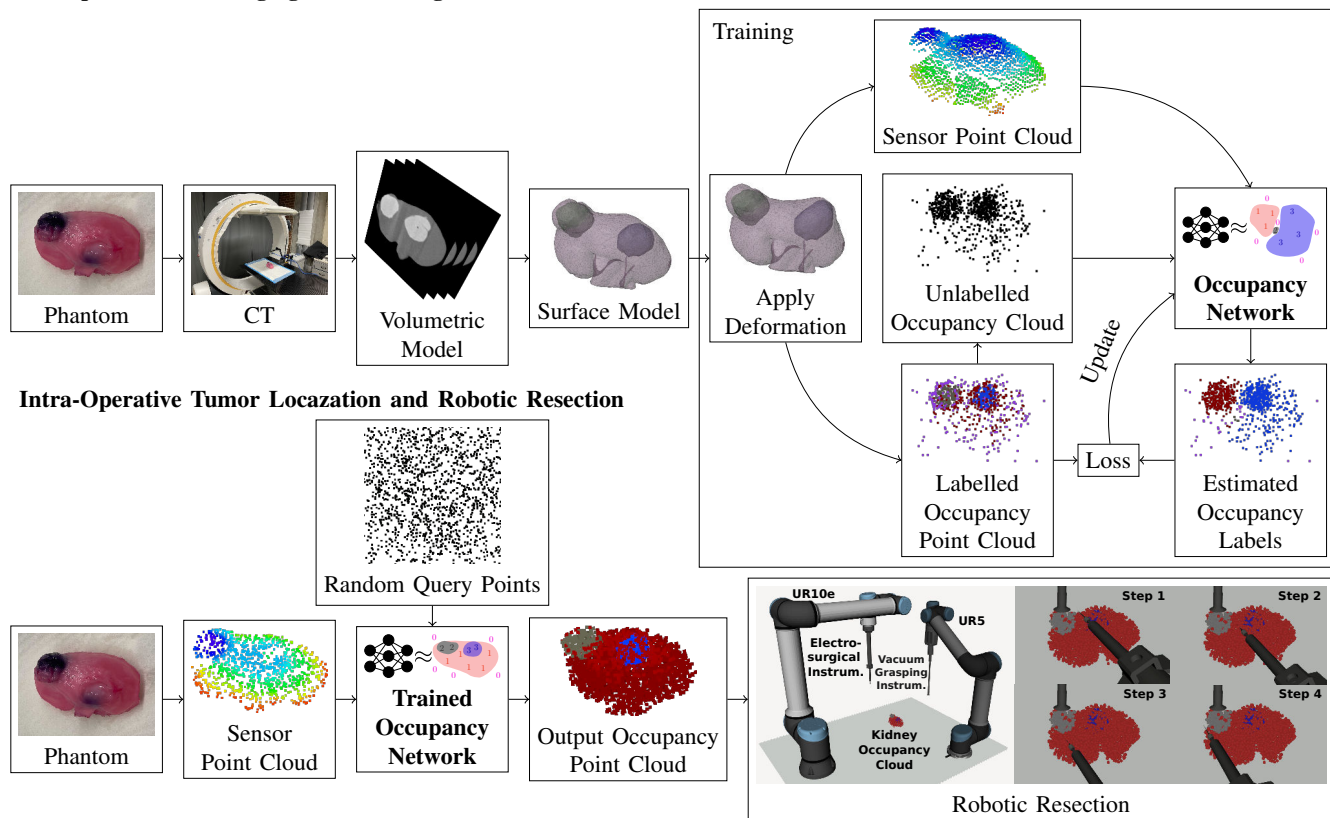


Fig. 2: Our processing pipeline used to perform an interactive tumor localization and robotic resection in simulation. There are two processing blocks. The **Pre-Operative** part uses a model obtained from an imaging modality, such as a CT image, to produce training data for our occupancy network. During training, the occupancy network learns to estimate the locations and shapes of the tumors under deformations. The **Intra-Operative** part uses a sensor point cloud obtained from an RGBD camera to output an occupancy point cloud that is innately compatible with our surgical resection method. The visualizations for the occupancy networks are 2D simplifications.

The occupancy network must learn how to estimate the tumor locations under deformation. Therefore, the training data needs to contain a deforming kidney. We propose a deformation system that ensures deformations are always partially visible to the camera. This ensures that reconstruction losses of parts not visible to the camera do not degrade the learning process.

To produce a deformed kidney, we follow these steps: A set of all vertices on the kidney surface is created. This set is filtered, such that only vertices visible to the camera are contained. From this set, random vertices v are moved by a random maximum distance, affecting each neighboring vertex n according to a Gaussian falloff $N(|v - n|, \sigma^2)$. This includes vertices on the tumor meshes. The number of vertices and the distance that they should be moved is evaluated in Section IV-B.

A SPC is taken from the deformed model and SortSample [8] is used to obtain an occupancy point cloud. The occupancy network is then trained using the occupancy point cloud and the SPC pairs. A total 31000 data pairs is used for training and 1000 are used for validation.

Intra-Operative Inference: For the robotic resection, see Section II-B, a dense point cloud is needed which encodes volume and shape of the kidney and tumors. This dense point cloud is the same as an occupancy point cloud, without the outside points. To produce the needed dense point cloud, we query the trained occupancy network using 40000 random query points. We discard all points that are not inside the kidney or the tumors. The occupancy point cloud is already registered to the camera and can be used without any modification as the input to our robotic resection method.

B. Integration with Robotic Resection

The integration of the registered occupancy point cloud into the robotic system is achieved using the Robot Operating System (ROS). A robotic resection plan is implemented and executed virtually, visualized with the Robot Visualization Tool (RViz). This process aims to verify the feasibility of using the dense occupancy point cloud to guide robotic resection, paving the way for future real-world applications.

We adapt ASTR [13] for autonomous nephrectomy. As depicted in Figure 2, ASTR is equipped with an electro-surgical instrument on an UR10e manipulator (Universal Robots,

Odense, Denmark), and a vacuum grasping tool on an UR5 manipulator. The kidney’s occupancy point cloud is streamed through ROS, and presented near ASTR as the surgical target. A specialized planner in the ROS system directs the surgery step-by-step. Initially, it localizes the exophytic tumor using labels and spatial data from the occupancy point cloud. Subsequently, the tumor’s center of mass is calculated and mapped to the superior surface of its bounding box, delineating the target for the vacuum grasping tool, which then reaches in vertically and immobilizes the exophytic tumor for easier resection. In the following step, the tumor point cloud is projected onto the transverse plane to identify a concave hull, delineating the tumor boundary. This boundary is extended by 5mm for the resection margin, and the area of intersection with the kidney’s point cloud projected on the transverse plane is designated as the zone requiring cutting for tumor removal. Points along the intersected contour are evenly spaced to guide the cutting process. Currently, the planner employs a strategy of two sweeping cuts, mapping the intersected contour onto the central transverse plane of the tumor point cloud, and the posterior surface of its bounding box. The electro-surgical instrument, positioned horizontally, follows these two paths to execute the cutting. After completing the cuts, both robotic arms concurrently return to their initial poses.

III. EXPERIMENTAL EVALUATION

A. Experimental Setup

The experimental setup is used to validate the tumor localization accuracy. A computed tomography (CT) machine (Loop-X, BrainLab, Germany) was used to scan the kidney phantom. An Intel RealSense D405 RGBD camera (Intel Corp., Santa Clara, California) captured the SPC of the kidney phantom during the experiment, as shown in Fig. 3.




Fig. 3: The kidney phantom is placed in the supine position onto the CT machine gantry. The RGBD camera is set up 15cm above the phantom.

The CT scans were performed using a tube voltage of 90 kV and an anode current of 13.6mA, with a CT Dose Area Product (DAP) of $500\text{cGy} \cdot \text{cm}^2$ over 400 slices.

After each SPC acquisition, and before CT scanning, the RGBD camera was moved out of the CT scanner’s field of

view to avoid artifacts from its metallic components. Care is taken not to touch or disturb the phantom and underlying supports during camera repositioning. This is to ensure the SPCs correspond to the same deformed state captured in the CT scans.

B. Kidney Phantom

We create two phantoms to fine-tune and evaluate the accuracy of our tumor localization method. The first phantom is used for the evaluation. It is created by combining two real-world patient cases into one phantom with a 38mm endophytic and 32mm exophytic tumor in one right kidney. The second phantom is used to fine-tune the amount of deformation that should be used during training, see Section II-A. This phantom is also based on a combination of two real-world patient cases. We use this second different phantom to ensure that we do not optimize our training data generation against our evaluation data.

Two batches of 10% polyvinyl alcohol (PVA) hydrogel were prepared, one with 2.27% barium sulfate (BaSO_4 , Sigma-Aldrich, St. Louis, MO) for contrast enhancement and the other without contrast agent. The contrast-enhanced PVA was injected into the tumor casts, while the plain PVA was used for the parenchyma cast. The exophytic tumor phantom was embedded directly into the cast. For the endophytic case, the tumor was suspended using wires before parenchyma injection. 10% PVA was then injected into the parenchyma casts, embedding the endophytic tumor. The entire phantom underwent freeze-thaw cycles to solidify and mimic mechanical tissue properties, with the final cycle facilitating cohesion between structures to replicate connective tissues.

The experiment aimed to simulate the in vivo kidney deformation that occurs when the kidney is subjected to a parenchymal incision under controlled renal arterial pressure during partial nephrectomy. Altamar et al. developed a kidney elastic deformation model describing the deformation caused by incision and loss of turgor [3]. The model indicates a general “sinking in” of the kidney surface due to gravity, accompanied by a small expansion along the lateral edges. In this biomechanical model, the maximum deformation is the indentation at the midsection of the kidney, which was approximately 8mm. The average displacement of the lateral edges is around 4.4mm.

Based on this elastic model, we force the kidney phantoms to deform. To simulate the “sinking in” surface configuration with expansion along the edges while keeping the surface relatively smooth without sudden indentations, we use a soft cardboard adhered to the anterior and posterior surfaces of the kidney phantom, see Fig. 5(a) and (b). Nylon cable ties are punctured through the soft cardboard and phantoms, locking on the other side. By adjusting the tightness of the cable ties, varying degrees of compression can be applied to the cardboard, and subsequently to the phantoms.

C. Experiments

To evaluate the accuracy of our tumor localization method, we use the first kidney phantom described in section III-B.




Fig. 4: From real patient CT cases, we segmented the patient’s kidney and tumors, generated a CAD model, and 3D printed the phantom mold. By injecting contrast-enhanced PVA liquid, we created a hydrogel phantom that exhibits clear tumor and parenchyma visualization under CT imaging.

Through deformation and rotation, we create 6 kidney variants from our original phantom. Three variants contain increasing deformations and three contain different orientations.




Fig. 5: Figure (a),(b) and (c) depict the kidney phantom in a normal supine configuration with a midsection height of approximately 57.87mm. Figures (d), (e) and (f) present the measurements of the midsection under compression, which are 51.21mm, 55.00mm, 59.06mm

As described in Section III-B, we generate three different degrees of deformations with midsection compressed by 4mm, 8mm, and 12mm, see Figure 5. We denote them as $D1$, $D2$ and $D3$, respectively. The three orientations are generated by rotating the kidney phantom by 90, 180 and 270 degrees. We denote them as $Rot90$, $Rot180$ and $Rot270$. For each kidney variant, the experimental setup described in section III-A is used to capture a CT image and a SPC.

The occupancy network requires the SPC to only contain the kidney. Therefore, we discard points further than 15cm from the camera. We also discard points closer than 11cm to remove some floating artifacts. We scale up the SPC by a factor of 1000, to match the scaling of the CT. 1000 unique points are randomly sampled from the SPC, to be used as the input to the point cloud encoder [8].

All meshes required for training are automatically extracted from the CT data. For this, the kidney and tumors are segmented using an upper-lower threshold segmentation, with the value ranges being $(-300, 300)$ and $(300, 1000)$ HU, respectively. The segmentation masks are converted to surface meshes. To remove small floating artifacts, only large connected triangle surfaces are kept.

The occupancy network is trained using training data created from the original kidney model. The kidney’s relative position and rotation to the camera may be unknown. We randomly rotate the original kidney virtually. As the kidney is round with few features, estimating the rotation is difficult from a single point cloud. We therefore produce two datasets. The first rotates the kidney randomly by up to 360 degrees around all axes. The second rotates the kidney by up to 360 degrees around the Z axis, but the random rotations for X and Y are uniformly sampled from the interval $[-15, 15]$ (degrees). The Z axis points upwards in Figure 3. The model trained on the fully random rotations is denoted as M_{full} , the other is M_{lim} , both of which are evaluated.

The kidney variants are used to obtain the evaluation data. For this, the point cloud and the meshes of the variants are manually registered. A rigid object placed next to the CT is used as a registration aid.

The point clouds of the kidney variants are then passed to the trained occupancy network. The occupancy network is sampled on an equidistant grid with a resolution of $400 \times 400 \times 400$ to produce a volumetric model. Although our robotic application does not require such a dense volumetric representation, it enables a more accurate evaluation. The volumetric model is then compared to the reference kidney models to compute the Hausdorff Distance (HD). In the context of estimated and reference tumors, the HD quantifies the maximum distance from any point on the estimated tumor surface to the nearest point on the real tumor surface. Consequently, this measurement provides an estimate for safety margin that should be added into the surgical excision of the tumor to ensure that it is completely resected. Additionally, we compute the center of the axis aligned bounding boxes for both the volumetric model and the reference object. The distance between these centers are computed, to evaluate how well the positioning of the tumors was estimated.

We also evaluate how varying the degree of deformation during training affects the accuracy of the output. For this, the second phantom is used. This evaluation is performed using limited rotations, like for M_{lim} . We apply a varying number of successive vertex-pull operations, as described in Section II-A. We use Mean Intersection over Union (mIoU) as a single metric that describes how well the kidney with all integrated tumors is estimated.

IV. RESULTS

A. Training and Inference Time

All measurements are performed using a single Nvidia RTX 4090 GPU (Nvidia, Santa Clara, California). During training, each model achieves its maximum validation accuracy after approximately 2 hours or 200 epochs. At inference time, the time needed to label 40000 query points, with a SPC containing 1000 points, is approximately 14ms.

B. Deformation Strength

TABLE I: Effect of using different deformations configurations on mIoU. Each (*Distance*, σ^2 , *Times Repeated*) triple represents a vertex-drag operation. Where a random vertex v is moved according to Section II-A.

Conf.	Deformations (Distance, σ^2 , Times Repeated)	mIoU
1	(16, 33.3, 1)	0.70
2	(16, 33.3, 1), (16, 16.7, 1), (8, 11.7, 1)	0.65
3	(8, 11.7, 1), (6, 10, 1), (4, 6.7, 1), (4, 3.3, 5), (2, 3.3, 10)	0.71
4	(16, 13.3, 1), (16, 10, 1), (16, 6.7, 1)	0.71
5	(10, 33.3, 1), (16, 13.3, 1), (16, 6.7, 1)	0.68
6	(16, 33.3, 1), (16, 16.7, 1), (8, 11.7, 1), (6, 10, 1), (4, 6.7, 1), (4, 3.3, 5), (2, 3.3, 10)	0.66

Six different deformation configurations are tested, see Table I. Overall, configuration 3 and 4 performed best, each with an mIoU of 0.71. For all following evaluations, we use deformations as defined in configuration 3 for creating training data.

C. Localization Accuracy

We evaluate the accuracy of the occupancy networks M_{full} and M_{lim} . Trained on fully and limited rotations respectively, see Section III-C. The results are shown in Table II. For scenes with small deformation, such as *D1*, M_{full} and M_{lim} perform similarly. Both achieve a HD of approximately 4mm and 7mm for the exophytic and endophytic renal tumors, respectively. For *D2*, both models again perform similarly, with tumor HDs of approximately 6mm and 9mm. For the strongest deformation *D3* we evaluate, M_{lim} better estimates the tumors, with HDs of 11.3mm and 12.9mm versus 15.4mm and 18.5mm. For the evaluated variants *Rot90-Rot270*, M_{full} estimates both tumors with HDs of less than or equal to 13.4mm. For all kidney and model combinations, except *D3* with M_{full} , the tumor centers are estimated with an error of less than 10mm.

V. DISCUSSION AND CONCLUSION

For scenes with little deformations, our models trained on limited rotations M_{lim} , and trained on full rotations M_{full} perform similarly. However, as deformations get stronger (*D3*), M_{full} is unable to accurately reconstruct the tumors. This is likely due to orientations being difficult to distinguish in the presence of strong deformations. Conversely, M_{lim} demonstrates a relatively consistent ability to estimate the tumor positions across different deformations. As it has only seen a limited amount of rotations during training, it has

more inductive bias. It is able to make more assumptions about the input data when making predictions. During minimally invasive surgery, the view of the kidney is limited, we therefore suggest limiting the rotations accordingly.

Generally, the HD of the exophytic tumor, compared to the endophytic tumor, is lower. The exophytic tumor is often, at least partially, visible to the depth sensor, providing features for estimating the location.

The HD values for both tumors suggest a safety margin of approximately 16mm to 19mm around the exophytic and endophytic tumors respectively. For smaller kidney deformations, safety margins of 6mm and 10mm are sufficient.

We plan to transfer our simulated resection to the real-world to perform a tumor resection on our phantom. For this, we must take visual obstructions into account. Further, we aim to refine the dual-arm resection technique, advancing from the current basic two-sweep strategy. This will include active tumor retraction through vacuum grasping, and finer forward movements at diverse angles with the electrosurgical instrument during a single sweep.

Method Limitations: Our results indicate that our current approach struggles to capture local deformations on the surface of the kidneys. For this, we will further improve our deformation system, used to produce training data. Adding regularizers, such as volume-retention, may improve the tumor locations estimation.

Evaluation Limitations: When registering the SPC and the CT images, we noticed distortions and artifacts in the SPC. This prevents an optimal registration, directly affecting our results. The move away from our low cost depth camera to a more accurate model can resolve this. The cardboard and cable tie method for registration causes the SPC to have a systematic offset due to the thickness of the cardboard. Using sutures to deform the kidney could resolve this issue. Finally, our manual registration introduces an error. Even with registration guides, we notice that our SPC and CT images do not align correctly. Again, this may be resolved with a better depth sensor. Using a virtual camera, to produce the SPC directly from the CT images, could provide a quantitative measure for the error caused by our depth sensor.

Conclusion: We present and evaluate an occupancy network based method to determine the location of exophytic and endophytic kidney tumors in a deformable phantom. This method only requires a single depth image. The accuracy of the location estimate varies with the strength of the kidney deformation. As the deformation increases, a larger safety margin must be chosen. For smaller kidney deformations, a safety margins of 10mm and 12mm for the exophytic and endophytic tumors are needed. This estimation can be performed over 60 times a second. Further, we show that the use of an occupancy network directly enables downstream tasks, such as a robotic resection. Currently, multiple limitations affect our evaluation and method accuracy. We discuss their possible solutions.

TABLE II: Results of using real-world point clouds to estimate the deformation of the kidney and the positions of the tumors. **Error** shows **over-reconstruction** and **under-reconstruction** when compared with the reference. H and C are the Hausdorff and Center distance between the reference and the reconstruction (in mm). D1-D3 and Rot90-Rot270 are the Deformed and Rotated kidneys. The model used to produce the reconstructions is indicated on the right.

	Input Point Cloud	Kidney Reconstr.	Kidney Reference	Kidney Error	Tumor Reconstr.	Tumor Reference	Tumor Error	
Trained on Full Rotation (\mathcal{M}_{full})	D1			H: 10.4 C: 2.2			H: 6.1 C: 2	
	D2			H: 12.0 C: 4			H: 9.2 C: 7	
	D3			H: 16.4 C: 4.1			H: 18.5 C: 13	
	Rot90			H: 20.8 C: 5.3			H: 11.6 C: 5.3	
	Rot180			H: 11.4 C: 2.5			H: 4.3 C: 2.2	
	Rot270			H: 27.6 C: 4			H: 9.6 C: 8	
Trained on Limited Rotation (\mathcal{M}_{lim})	D1			H: 8.7 C: 2			H: 7.2 C: 3.8	
	D2			H: 9.1 C: 3.4			H: 4.2 C: 2.4	
	D3			H: 18.1 C: 8.3			H: 9.1 C: 3.7	
							H: 5.9 C: 3.6	
							H: 12.9 C: 7.3	
							H: 11.3 C: 4.8	

REFERENCES

- [1] R. L. Siegel, K. D. Miller, and A. Jemal, "Cancer statistics, 2019," *CA: a cancer journal for clinicians*, vol. 69, no. 1, p. 7–34, January 2019.
- [2] P. Cannon, S. A. Setia, S. Klein-Gardner, N. Kavoussi, R. J. Webster III, and D. Herrell, "Are 3d image guidance systems ready for use? a comparative analysis of 3d image guidance implementations in minimally invasive partial nephrectomy," *Journal of Endourology*, no. ja, 2024.
- [3] H. O. Altamar, R. E. Ong, C. L. Glisson, D. P. Viprakasit, M. I. Miga, S. D. Herrell, and R. L. Galloway, "Kidney deformation and intraoperative registration: A study of elements of image-guided kidney surgery," *Journal of Endourology*, vol. 25, no. 3, pp. 511–517, 2011, pMID: 21142942. [Online]. Available: <https://doi.org/10.1089/end.2010.0249>
- [4] J. Ge, H. Saeidi, J. D. Opfermann, A. S. Joshi, and A. Krieger, "Landmark-guided deformable image registration for supervised autonomous robotic tumor resection," in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I* 22. Springer, 2019, pp. 320–328.
- [5] T. Simpfendorfer, C. Gasch, G. Hatiboglu, M. Müller, L. Maier-Hein, M. Hohenfellner, and D. Teber, "Intraoperative computed tomography imaging for navigated laparoscopic renal surgery: first clinical experience," *Journal of endourology*, vol. 30, no. 10, pp. 1105–1111, 2016.
- [6] N. Nimmagadda, J. M. Ferguson, N. L. Kavoussi, B. Pitt, E. J. Barth, J. Granna, R. J. Webster III, and S. D. Herrell III, "Patient-specific, touch-based registration during robotic, image-guided partial nephrectomy," *World journal of urology*, vol. 40, no. 3, pp. 671–677, 2022.
- [7] X. Zhang, J. Wang, T. Wang, X. Ji, Y. Shen, Z. Sun, and X. Zhang, "A markerless automatic deformable registration framework for augmented reality navigation of laparoscopy partial nephrectomy," *International journal of computer assisted radiology and surgery*, vol. 14, no. 8, p. 1285–1294, August 2019. [Online]. Available: <https://link.springer.com/content/pdf/10.1007/s11548-019-01974-6.pdf>
- [8] P. Henrich, B. Gyenes, P. M. Scheikl, G. Neumann, and F. Mathis-Ullrich, "Registered and segmented deformable object reconstruction from a single view point cloud," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2024, pp. 3129–3138.
- [9] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.
- [10] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *Advances in neural information processing systems*, vol. 30, 2017.
- [11] B. A. Birnbaum, N. Hindman, J. Lee, and J. S. Babb, "Renal cyst pseudoenhancement: influence of multidetector ct reconstruction algorithm and scanner type in phantom model," *Radiology*, vol. 244, no. 3, pp. 767–775, 2007.
- [12] R. Melnyk, B. Ezzat, E. Belfast, P. Saba, S. Farooq, T. Campbell, S. Mcleavey, M. Buckley, and A. Ghazi, "Mechanical and functional validation of a perfused, robot-assisted partial nephrectomy simulation platform using a combination of 3d printing and hydrogel casting," *World Journal of Urology*, vol. 38, 07 2020.
- [13] J. Ge, M. Kam, J. D. Opfermann, H. Saeidi, S. Leonard, L. J. Mady, M. J. Schnermann, and A. Krieger, "Autonomous system for tumor resection (astr)-dual-arm robotic midline partial glossectomy," *IEEE robotics and automation letters*, vol. 9, no. 2, pp. 1166–1173, 2024.