

LETTER

Neural fields for 3D tracking of anatomy and surgical instruments in monocular laparoscopic video clips

Beerend G. A. Gerats^{1,2}  | Jelmer M. Wolterink^{3,4} | Seb P. Mol^{1,4} | Ivo A. M. J. Broeders^{1,2}

¹AI & Data Science Center, Meander Medical Center, Amersfoort, The Netherlands

²Robotics and Mechatronics, University of Twente, Enschede, The Netherlands

³Department of Applied Mathematics, University of Twente, Enschede, The Netherlands

⁴Technical Medical Center, University of Twente, Enschede, The Netherlands

Correspondence

Beerend G.A. Gerats, AI & Data Science Center, Meander Medical Center, Maatweg 3, 3813 TZ, Amersfoort, The Netherlands.
Email: bga.gerats@meandermc.nl

Funding information

Johnson & Johnson MedTech; NWO, Grant/Award Number: 18192

Abstract

Laparoscopic video tracking primarily focuses on two target types: surgical instruments and anatomy. The former could be used for skill assessment, while the latter is necessary for the projection of virtual overlays. Where instrument and anatomy tracking have often been considered two separate problems, in this article, a method is proposed for joint tracking of all structures simultaneously. Based on a single 2D monocular video clip, a neural field is trained to represent a continuous spatiotemporal scene, used to create 3D tracks of all surfaces visible in at least one frame. Due to the small size of instruments, they generally cover a small part of the image only, resulting in decreased tracking accuracy. Therefore, enhanced class weighting is proposed to improve the instrument tracks. The authors evaluate tracking on video clips from laparoscopic cholecystectomies, where they find mean tracking accuracies of 92.4% for anatomical structures and 87.4% for instruments. Additionally, the quality of depth maps obtained from the method's scene reconstructions is assessed. It is shown that these pseudo-depths have comparable quality to a state-of-the-art pre-trained depth estimator. On laparoscopic videos in the SCARED dataset, the method predicts depth with an MAE of 2.9 mm and a relative error of 9.2%. These results show the feasibility of using neural fields for monocular 3D reconstruction of laparoscopic scenes. Code is available via GitHub: <https://github.com/Beerend/Surgical-OmniMotion>.

1 | INTRODUCTION

Optical tracking of elements in laparoscopic videos provides valuable information for computer-assisted surgery. Tracking efforts primarily focus on two target types: surgical instruments and anatomy. For example, the former is used for skill assessment [1, 2], while the latter is necessary for the projection of virtual overlays [3].

Several 2D instrument trackers have been proposed [4–6], but these trackers generally lack depth perception and fail to track points adequately in 3D. This might be overcome by the inclusion of markers [7] or electromagnetic tracking [8]. However, such approaches require adaptations to the instruments or surgical workflow and are not widely available. 3D pose trackers designed for correcting position errors, for example, in robotic systems, require a known pose at the start of tracking [9]. Other methods use camera parameters to calculate 3D

instrument locations [10], which are typically unknown for regular surgery. In anatomical tracking, some approaches use an additional depth sensor [11]. However, these are difficult to implement in-vivo. Alternatively, depth could be estimated from stereoscopic [12] or monocular video [13, 14]. Although the depth estimates are highly informative, these methods do not model tissue deformation. Methods for 3D scene reconstruction based on simultaneous localization and mapping, do model for anatomical deformations [12, 15, 16], but are limited in handling complex deformations and color changes.

While instrument and anatomy tracking have often been considered two separate problems, here, we propose a method for joint tracking of *all* elements in a laparoscopic video. Based on a 2D video clip, we track the 3D location of all surfaces that are visible in at least one frame of the video. Key to our approach is the use of neural fields, which represent continuous spatio-(temporal) scenes in the weights of neural networks.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial License](https://creativecommons.org/licenses/by-nc/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2024 The Author(s). *Healthcare Technology Letters* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology.




FIGURE 1 Overview of our method. A pixel p_i is selected in frame i , from where a ray is cast through the virtual scene. A sampled location x_i is mapped to location u in canonical volume G . This location is used to predict color c and material density σ , through a multi-layer perceptron (MLP). To predict the translation of this point to frame j , the location is mapped via an inverse transform and reprojected to the predicted pixel location \hat{p}_j .

Neural fields have previously been used in laparoscopic scene representation [17–19]. A limitation of these prior works is the need for stereoscopic vision, which limits the applicability of these methods in standard minimally invasive procedures. Instead, we propose to leverage the recently proposed OmniMotion [20] approach, in which monocular videos are used to reconstruct a 3D + time scene representation, without the need for camera calibration, stereoscopic vision, or depth maps. We evaluate the tracking performance on video clips from laparoscopic cholecystectomies for various tissue and instrument types. We show that instrument tracks can be obtained using enhanced class weighting. Last, we assess the quality of the depth maps obtained from the method’s scene reconstructions. We show that these pseudo-depths have comparable quality to a state-of-the-art pre-trained depth estimator.

2 | METHODS

We propose to use neural fields as representations of surgical scenes visualized in monocular laparoscopic video clips. The neural fields are optimized using OmniMotion [20], with additional class weighting to improve surgical instrument tracking. An overview of the method is given in Figure 1.

2.1 | Neural fields for pixel tracking

We use OmniMotion [20] for the reconstruction of 3D + time scene representations from 2D videos. The method captures camera movement and scene deformations in monocular video by mapping through a canonical volume G . Camera rays are cast through the scene at frame i , along which 3D locations x_i are sampled. Using bijective mapping $u = T_i(x_i)$, the 3D locations are translated to location u in the canonical volume. The mappings are invertible such that the new position of this point in frame j can be calculated with $x_j = T_j^{-1}(T_i(x_i))$. A canonical location u is used to query a multi-layer perceptron (MLP)

that returns material density σ and color c for this point. When all sampled points along a camera ray are mapped to canonical space and queried to the MLP to obtain their densities and colors, the quadrature rule is used to find the predicted color \hat{C}_i for the corresponding pixel p_i [21]. The predicted flow of a pixel between two frames $\hat{f}_{i \rightarrow j} = \hat{p}_j - p_i$ can be calculated similarly, using the material density values of the sampled points along the ray. In this way, the model can account for occlusions. The predicted flows are supervised with flows from a pre-trained optical flow model (RAFT) [22] (Equation 1), while the predicted pixel colors are supervised with the training images (Equation 2).

$$L_{\text{flow}} = \sum_{f_{i \rightarrow j} \in \Omega_p} \|\hat{f}_{i \rightarrow j} - f_{i \rightarrow j}\|_1, \quad (1)$$

$$L_{\text{color}} = \sum_{(i,p) \in \Omega_p} \|\hat{C}_i(p) - C_i(p)\|_2^2. \quad (2)$$

Although the method is supervised with 2D flows, the method’s representation of scene deformation is in 3D. This makes it possible to track locations on anatomical structures and to follow surgical instruments over the depth axis. Because of the flow regularization, the method does not require camera parameters, stereoscopic video, or pre-computed depth images. Additionally, as camera movement is compensated for implicitly by the method, camera locations are not needed as well. This makes the method applicable to video clips from a regular laparoscope. Note that for each scene, we optimize an individual model.

2.2 | Segmentation masks

Since OmniMotion simultaneously tracks all pixels in the video, it is not possible to distinguish instruments from anatomy pixels. Therefore, we propose to combine the method with segmentation masks that cover the elements of interest. The segmentation masks can be manually annotated or found by

pre-trained segmentation models [23]. Note that the masks are not necessary for training and can be applied afterward unless class weighting is used (see next section). Only a single mask at the start of the video is required to track these pixels throughout the rest of the video. An example mask can be seen in Figure 1.

2.3 | Class weighting

One of our goals is to track surgical instruments through the video. However, due to their size, instruments generally cover a small part of the image only. In the regular training approach of neural radiance fields, there is no incentive to represent a pixel more accurately than others. This means that reconstructions of irrelevant background pixels are reconstructed with similar importance as the surgical instruments. Therefore, we propose to use weighting in the loss function, where the weight depends on the class of the pixel, that is, “instrument” or “background”:

$$L = w_{\text{class}}(L_{\text{flow}} + \lambda L_{\text{color}}) + L_{\text{other}}, \quad (3)$$

where $w_{\text{class}} = 1$ if the pixel is in the background and a larger value otherwise, λ is a loss weighting factor, and L_{other} is a combination of regularizing losses. The pixel classes could be obtained by manual or automatic segmentation masks.

2.4 | Data

We evaluate our model on two tasks. First, we assess the method’s performance in 2D tracking of anatomical structures and surgical instruments. We train OmniMotion on all 80-frame video clips in the CholecSeg8k dataset [24]. The set includes 101 short videos of 25 fps that display laparoscopic cholecystectomies. Each video frame is accompanied by manually annotated segmentation masks of tissues and instruments. We pre-process the videos by cropping to a rectangle, to remove the circular border of the scope as much as possible, and resizing to 256×256 pixels. After pre-processing, we find 84 clips that contain a segmentation in the first frame for the “gallbladder” class. For “liver” and “gastrointestinal tract” we find 101 and 46 clips, respectively. For the instruments “grasper” and “l-hook electrocautery” there are 65 and 24 clips available, respectively. In total, we track 231 anatomical structures and 89 surgical instruments.

Second, we explore the accuracy of the third axis, by comparing predicted depths with ground-truth depth measurements. For the evaluation of the depth predictions, we use laparoscopic videos of the SCARED dataset [25]. The videos display a static surgical scene, recorded by various moving cameras. We select the first camera of each scene, as we evaluate the method for the monocular setting. The videos in the dataset are attributed with ground-truth depth values, measured via a structured light pattern projected on the field of view. We train OmniMotion on all nine videos individually, where the video length ranges between 88 and 728 frames.

2.5 | Evaluation metrics

For 2D tracking, we calculate the mean accuracy of pixels in a segmentation mask that are tracked throughout the video. For each pixel p_1 in the segmentation mask of the first frame, we get a predicted pixel location \hat{p}_j at consecutive frames $j \in [2, 3, \dots, 80]$. Predictions that are located within the segmentation mask of that frame are considered accurate, while those that fall outside the mask are inaccurate. Predicted locations that are outside the image boundaries are excluded from evaluation. In this way, we calculate the per-frame accuracy. When no segmentation mask is available, for example, when an instrument is temporarily out of view, the frame is discarded. Finally, we calculate the mean accuracy by averaging the per-frame accuracies over all frames first, and over all videos afterwards.

For the evaluation of the predicted depth maps, we follow the general approach for evaluating pseudo-depths [26], meaning that we shift and scale the depth predictions towards the ground-truth values before calculating the metrics.

3 | EXPERIMENTS AND RESULTS

We train the method in batches of 256 correspondences from 8 image pairs. Along each camera ray, 32 points are sampled. Learning rates 3×10^{-4} and 1×10^{-4} are used for color and flow loss, respectively. Using an NVIDIA A40 GPU, it takes ± 10 min for RAFT pre-processing an 80-frame video clip. Training OmniMotion for 10K iterations requires an additional ± 45 min.

3.1 | Evaluation of anatomy and instrument tracking

A selection of qualitative results is given in Figure 2. The top row displays two videos where gallbladder pixels are tracked. In the left example, the tracks accurately follow the gallbladder deformation under retraction with a grasper. In the right example, the camera zooms out while the gallbladder is released from the grasper grip. The pixel tracks accurately follow the anatomical structure. In the bottom row, the grasper is tracked. Left, the instrument leaves the field of view during the middle section of the video. It can be seen that the method accurately re-identifies the grasper when re-entering the field of view. In the right example, the grasper is heavily displaced to the upper-right corner in the last frame, where only a small part of the instrument remains available. Also in this scenario, the tracks follow the target pixels accurately. Instrument movements are visualized in Figure 3 for two example clips. In the 3D plots, we show the median location of the instrument over time. It can be seen that the method captures movements over three axes.

Table 1 provides the quantitative measurements in tracking performance of anatomies and instruments. With a mean accuracy of 92.4%, the method is able to track a large portion of the anatomy pixels accurately. Although most instrument pixels




FIGURE 2 Examples of tracked pixels in four videos (top: two gallbladder videos, bottom: two grasper videos). For each video, we display the first, middle and last frame here. The full videos are available via: <https://vimeo.com/920225544>.

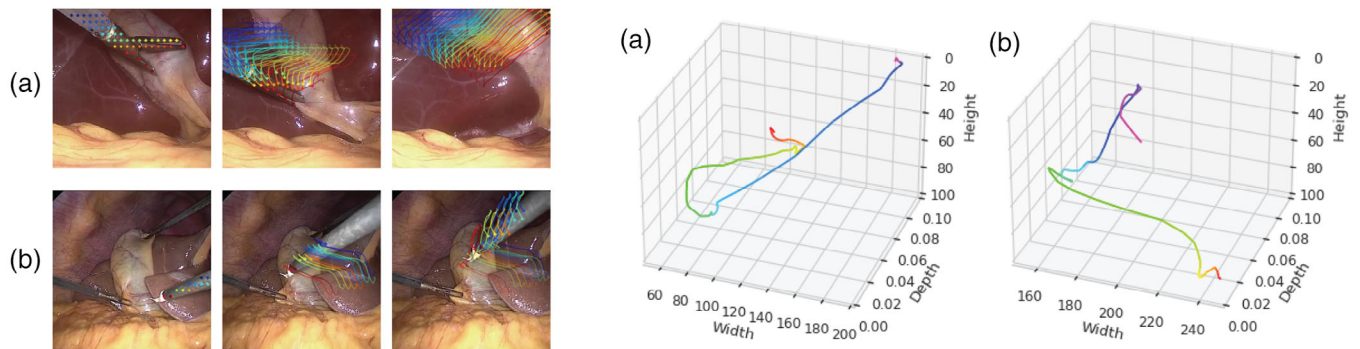


FIGURE 3 Two examples of surgical instruments, with a grasper (a) and an L-hook (b), tracked throughout an 80-frame video. Left: first, middle and last frame of the video, with colored trails indicating our methods predictions for pixel displacement. Right: predicted instrument tracks visualized in 3D, where the colored lines correspond to instrument movement over time.

TABLE 1 Mean accuracy (%) and standard deviations for 2D tracking of anatomical structures and surgical instruments in CholecSeg8k dataset. Instruments are tracked with different class weights w_{class} . Note that $w_{\text{class}}=1$ means that no class weighting is used.

	$w_{\text{class}} = 1$	$w_{\text{class}} = 5$	$w_{\text{class}} = 10$	$w_{\text{class}} = 20$
Gallbladder	91.7 (± 11.8)	—	—	—
Liver	93.6 (± 8.7)	—	—	—
Gastroint. tract	90.9 (± 15.0)	—	—	—
Grasper	83.6 (± 23.0)	84.3 (± 21.3)	85.1 (± 19.8)	85.9 (± 20.0)
L-hook	83.6 (± 28.8)	92.0 (± 21.7)	92.5 (± 20.6)	91.5 (± 20.3)
All anatomies	92.4 (± 11.4)	—	—	—
All instruments	83.6 (± 24.7)	86.3 (± 21.7)	87.1 (± 20.3)	87.4 (± 20.3)

are properly tracked as well, the tracking performance is lower. The drop in performance can be explained by the small image area that the instruments typically cover. The effect can be counterbalanced with class weighting, where instrument tracking performance gradually increases from 83.6% to 87.4% for larger weights.

In our results, class weighting is more effective for L-hooks than for graspers. A possible cause is that graspers display more displacement through the video frames than L-hooks, making tracking more difficult. To illustrate this, the average DICE overlap of tool masks in two consecutive frames is 0.88 for “grasper” and 0.92 for “l-hook electrocautery”. For comparison, the average overlaps of anatomical structures are 0.94, 0.97 and 0.92 for “gallbladder”, “liver” and “gastrointestinal tract”, respectively.

We use class weighting for the surgical instruments only, as those pixels are largely underrepresented in the video frames. Typically, anatomical structures fill larger parts of the image. In the dataset used, an anatomy class takes 12.2% of the pixels in a video clip, on average. For a surgical tool, the average is 3.6%.

Tracking accuracy for different frame rates is given in Figure 4. Here, we decrease the temporal resolution of 25 fps gradually to 50, 25, 12.5, and 6.25 percent of the original. Where the tracking performance for the anatomical structures remains largely unaffected, the quality of the instrument tracks decreases with fewer frames per second. This effect could be explained by the more drastic movements that instruments generally make in comparison to the anatomy.




FIGURE 4 Mean accuracy (%) for 2D tracking with decreasing temporal resolutions. The performance decreases for lower frame rates with instrument tracks, while the performance for anatomical structures is stable over all temporal resolutions.

TABLE 2 Error of pseudo-depth estimation on surgical scenes in SCARED dataset, in comparison with estimations by a pre-trained dense depth estimator (DPT) [27]. Although our method is not specifically trained for depth estimation, its depth reconstructions are of comparable quality those from a state-of-the-art-method.

Scene no.	MAE (mm) ↓		AbsRel (%) ↓		$\delta < 1.25$ (%) ↑	
	OmniMotion	DPT	OmniMotion	DPT	OmniMotion	DPT
1	2.9	2.2	7.8	6.2	95.5	99.0
2	1.4	1.5	4.8	4.9	97.8	99.8
3	5.3	5.0	14.6	13.0	77.0	85.4
4	1.7	2.0	7.4	8.1	97.1	96.3
5	2.1	4.2	7.8	17.3	95.5	77.8
6	2.5	2.2	7.2	6.4	96.2	98.8
7	5.1	4.5	13.9	13.0	80.0	84.6
8	2.2	2.1	7.5	7.3	96.7	96.0
9	2.7	1.6	12.0	6.8	90.3	96.5
Average	2.9 (± 1.3)	2.8 (± 1.3)	9.2 (± 3.2)	9.2 (± 4.0)	91.8 (± 7.4)	92.7 (± 7.5)

3.2 | Evaluation of reconstructed pseudo-depth

To assess the quality of the third pixel-tracking axis, we compare the estimated depth values with ground-truth measurements. Table 2 shows the error in depth estimation for all nine videos in the SCARED dataset [25]. On average, the estimates have a mean average error of 2.9 mm and a relative error of 9.2% to the actual depth. This performance is comparable to depth estimates by DPT [27], a state-of-the-art pre-trained transformer for monocular depth estimation. The comparison shows that the depth axis of the pixel tracks is relatively accurate as well. Figure 5 shows an example frame, with depth estimates by DPT and OmniMotion. Although the method's depth values are less smooth, they represent an accurate geometry of the underlying scene. Where monocular depth estimators, such as DPT,

are computationally faster, they do not provide pixel tracking through video.

4 | DISCUSSION AND CONCLUSION

We combined tissue deformation modeling and surgical instrument tracking by tracking all pixels in a laparoscopic video at once. We used neural fields for 3D scene reconstruction, which could be trained without the need for camera calibration, camera locations, stereoscopic vision, or depth maps. In contrast to previously proposed methods for 3D scene reconstructions, this makes the method applicable as-is to standard minimally invasive procedures. In our experiments, we show that the combination of OmniMotion [20] and segmentation masks can be used for constructing accurate 3D anatomy and instrument




FIGURE 5 Left: an example frame from the SCARED dataset. Middle: disparity estimated by DPT [27]. Right: disparity calculated from the OmniMotion scene reconstruction. Although the depth values in the latter result are less smooth, the geometry of the underlying scene is accurately represented.

tracks. The quality of the instrument tracks can be improved by using class weighting, where instrument pixels are weighted more heavily during optimization.

A limitation of the used method is the relatively slow training speed. Ideally, the scene reconstructions are built in real-time, such that the resulting tracks are available during surgery. Yang et al. [18] have shown that drastic speed-ups, over a 100× faster, are possible by leveraging multi-plane fields. Moreover, the use of Gaussian Splatting [28] could potentially result in even faster training and real-time rendering.

A direction for future work is a deeper exploration of uses for neural fields from surgical videos. For example, when scene reconstructions are accurate enough, anatomical structures in pre-operative scans could potentially be mapped to the neural representation. Because neural fields can model tissue deformation, virtual overlays could be rendered that follow the underlying structures. Another example is the use of 3D anatomical tracks for the calculation of stress applied to the tissue, which could be visualized or used for haptic feedback. In our work, we have shown the feasibility of using neural fields for 3D monocular scene reconstructions, which could form the basis for these further explorations.

AUTHOR CONTRIBUTIONS

Berend G.A. Gerats: Conceptualization; data curation; investigation; methodology; software; visualization; writing—original draft; writing—review and editing. **Jelmer M. Wolterink:** Conceptualization; methodology; supervision; writing—review and editing. **Sep P. Mol:** Data curation; investigation; software. **Ivo A.M.J. Broeders:** Conceptualization; funding acquisition; supervision.

ACKNOWLEDGEMENTS

Research at the Meander AI & Data Science Center was sponsored by Johnson & Johnson MedTech. Jelmer M. Wolterink was supported by NWO domain Applied and Engineering Sciences VENI Grant (18192)

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

The CholecSeg8k dataset is available at <https://www.kaggle.com/datasets/newslab/cholecseg8k>. The SCARED dataset

is available at <https://endovissub2019-scared.grand-challenge.org>. Restrictions apply to the availability of these data, which were used under license for this study.

ORCID

Berend G.A. Gerats  <https://orcid.org/0000-0003-2570-1834>

REFERENCES

- Fathollahi, M., et al.: Video-based surgical skills assessment using long term tool tracking. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 541–550. Springer, Cham (2022)
- Lavanchy, J.L., et al.: Automation of surgical skill assessment using a three-stage machine learning algorithm. *Sci. Rep.* 11(1), 5197 (2021)
- Pelani, E., et al.: Evaluation of a novel navigation platform for laparoscopic liver surgery with organ deformation compensation using injected fiducials. *Med. Image Anal.* 69, 101946 (2021)
- Du, X., et al.: Articulated multi-instrument 2-d pose estimation using fully convolutional networks. *IEEE Trans. Med. Imaging* 37(5), 1276–1287 (2018)
- Kurmann, T., et al.: Simultaneous recognition and pose estimation of instruments in minimally invasive surgery. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2017: 20th International Conference, Proceedings, Part II 20, pp. 505–513. Springer, Cham (2017)
- Robu, M., et al.: Towards real-time multiple surgical tool tracking. *Comp. Meth. Biomech. Biomed. Eng.: Imag. Vis.* 9(3), 279–285 (2021)
- Cartucho, J., et al.: An enhanced marker pattern that achieves improved accuracy in surgical tool tracking. *Comp. Meth. Biomech. Biomed. Eng.: Imag. Vis.* 10(4), 400–408 (2022)
- Liu, X., et al.: Laparoscopic stereoscopic augmented reality: Toward a clinically viable electromagnetic tracking solution. *J. Med. Imag.* 3(4), 045001 (2016)
- Du, X., et al.: Combined 2d and 3d tracking of surgical instruments for minimally invasive and robotic-assisted surgery. *Int. J. Comp. Assist. Radiol. Surg.* 11, 1109–1119 (2016)
- Zhao, Z., et al.: Tracking-by-detection of surgical instruments in minimally invasive surgery via the convolutional neural network deep learning-based method. *Comp. Assist. Surg.* 22(sup1) 26–35 (2017)
- Golse, N., et al.: Augmented reality during open liver surgery using a markerless non-rigid registration system. *J. Gastrointest. Surg.* 25, 662–671 (2021)
- Long, Y., et al.: E-dssr: efficient dynamic surgical scene reconstruction with transformer-based stereoscopic depth perception. In: Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Proceedings, Part IV 24, pp. 415–425. Springer, Cham (2021)
- Beilei, C., et al.: Surgical-dino: Adapter learning of foundation model for depth estimation in endoscopic surgery. *arXiv preprint, arXiv:240106013* (2024)

14. Shao, S., et al.: Self-supervised monocular depth and ego-motion estimation in endoscopy: Appearance flow to the rescue. *Med. Image Anal.* 77, 102338 (2022)
15. Li, Y., et al.: Super: A surgical perception framework for endoscopic tissue manipulation with surgical robotics. *IEEE Rob. Autom. Lett.* 5(2), 2294–2301 (2020)
16. Song, J., et al.: Dynamic reconstruction of deformable soft-tissue with stereo scope in minimal invasive surgery. *IEEE Rob. Autom. Lett.* 3(1), 155–162 (2017)
17. Wang, Y., et al.: Neural rendering for stereo 3d reconstruction of deformable tissues in robotic surgery. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 431–441. Springer, Cham (2022)
18. Yang, C., et al.: Neural lerplane representations for fast 4d reconstruction of deformable tissues. *arXiv preprint*, arXiv:230519906 (2023)
19. Zha, R., et al.: 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, Cham (2023)
20. Wang, Q., et al.: Tracking everything everywhere all at once. In: *International Conference on Computer Vision*. IEEE, Piscataway (2023)
21. Mildenhall, B., et al.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65(1), 99–106 (2021)
22. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: *Computer Vision—ECCV 2020: 16th European Conference, Proceedings, Part II 16*, pp. 402–419. Springer, Cham (2020)
23. Grammatikopoulou, M., et al.: A spatio-temporal network for video semantic segmentation in surgical videos. *Int. J. Comp. Assist. Radiol. Surg.* 19(2), 375–382 (2024)
24. Hong, W.Y., et al.: Cholecseg8k: a semantic segmentation dataset for laparoscopic cholecystectomy based on cholec80. *arXiv preprint*, arXiv:201212453 (2020)
25. Allan, M., et al.: Stereo correspondence and reconstruction of endoscopic data challenge. *arXiv preprint*, arXiv:210101133 (2021)
26. Ranftl, R., et al.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans. Pattern Anal. Mach. Intell.* 44(3), 1623–1637 (2020)
27. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12179–12188. IEEE, Piscataway (2021)
28. Kerbl, B., et al.: 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* 42(4), 139 (2023)

How to cite this article: Gerats, B.G.A., Wolterink, J.M., Mol, S.P., Broeders, I.A.M.J.: Neural fields for 3D tracking of anatomy and surgical instruments in monocular laparoscopic video clips. *Healthc. Technol. Lett.* 11, 411–417 (2024).
<https://doi.org/10.1049/htl2.12113>