

Towards Abdominal 3-D Scene Rendering from Laparoscopy Surgical Videos using NeRFs

Khoa Tuan Nguyen^{1,2}, Francesca Tozzi^{4,6}, Nikdokht Rashidian^{5,6},
Wouter Willaert^{4,6}, Joris Vankerschaver^{2,3}, and Wesley De Neve^{1,2}

¹ IDLab, ELIS, Ghent University, Ghent, Belgium

² Center for Biosystems and Biotech Data Science, Ghent University Global Campus,
Incheon, Korea

³ Department of Applied Mathematics, Informatics, and Statistics, Ghent University,
Ghent Belgium

{khoatuan.nguyen,joris.vankerschaver,wesley.deneve}@ghent.ac.kr

⁴ Department of GI Surgery, Ghent University Hospital, Ghent, Belgium

⁵ Department of HPB Surgery & Liver Transplantation, Ghent University Hospital,
Ghent, Belgium

⁶ Department of Human Structure and Repair, Ghent University, Ghent, Belgium
{francesca.tozzi,nikdokht.rashidian,wouter.willaert}@ugent.be

Abstract. Given that a conventional laparoscope only provides a two-dimensional (2-D) view, the detection and diagnosis of medical ailments can be challenging. To overcome the visual constraints associated with laparoscopy, the use of laparoscopic images and videos to reconstruct the three-dimensional (3-D) anatomical structure of the abdomen has proven to be a promising approach. Neural Radiance Fields (NeRFs) have recently gained attention thanks to their ability to generate photorealistic images from a 3-D static scene, thus facilitating a more comprehensive exploration of the abdomen through the synthesis of new views. This distinguishes NeRFs from alternative methods such as Simultaneous Localization and Mapping (SLAM) and depth estimation. In this paper, we present a comprehensive examination of NeRFs in the context of laparoscopy surgical videos, with the goal of rendering abdominal scenes in 3-D. Although our experimental results are promising, the proposed approach encounters substantial challenges, which require further exploration in future research.

Keywords: 3-D reconstruction · Laparoscopy · Neural Rendering · View Synthesis.

1 Introduction

Laparoscopy, also known as keyhole surgery or minimally invasive surgery (MIS), is a surgical technique that enables a surgeon to access the inside of the abdomen without the need for making large incisions in the skin. The surgeon inserts a slender tool, with a light and a camera attached, through small skin incisions, which makes it then possible for the surgeon to see inside the abdomen, negating




Fig. 1. Illustration of NeRFs and 3-D reconstruction using estimated depth information. From left to right: a video frame extracted from a laparoscopic surgical video, an image rendered by NeRFs, depth estimation for the given video frame, and reconstructions in both mesh and voxel formats based on the estimated depth information.

the need for large skin cuts and a long recovery period. However, the view of the surgeon during MIS is limited to the perspective of the camera. To address this limitation, various approaches have been developed to provide surgeons with additional information such as depth, segmentation, 3-D reconstruction, and augmented surgery techniques [1,10,24]. Despite substantial advancements, current methods still face challenges in achieving high-quality and photorealistic 3-D reconstruction results, as shown in Fig. 1. Neural Radiance Fields (NeRFs) [13,15] have recently emerged as a popular method for view synthesis using natural images. However, their application in the medical field has thus far remained limited [4,8,26]. In this paper, we focus on the task of understanding the underlying scene in MIS through the utilization of NeRFs. Specifically, we investigate how well NeRFs can be applied to laparoscopic surgical videos, particularly during the preparatory phase of a surgical procedure where surgical tools, surgical actions, scenes with blood, and moving organs are absent. In summary, our study makes the following contributions:

- We present a workflow that analyzes laparoscopic surgical videos by employing state-of-the-art NeRFs-based methods to render a complete abdominal environment.
- We adopt mask-based ray sampling to better mitigate the presence of artefacts that are associated with the black outer areas commonly found in laparoscopic surgical videos.
- We investigate the performance of each adopted NeRFs-based method, identifying and discussing the challenges associated with their usage.

2 Related Work

One approach to achieve 3-D scene reconstruction is through the utilization of Simultaneous Localization And Mapping (SLAM), which enables the mapping of 3-D coordinates and the direct reconstruction of 3-D surfaces from a given camera view [10]. Guodong *et al.* proposed an enhanced SLAM method called MKCF-SLAM [27,28], incorporating multi-scale feature patch tracking to reconstruct laparoscopic scenes characterized by texture deficiency, high levels of reflection, and changes in target scale. In contrast, Haoyin *et al.* presented a 2-D

non-rigid SLAM system known as EMDQ-SLAM [31,32], capable of compensating for pixel deformation and performing real-time image mosaicking. Another approach for 3-D scene reconstruction involves depth estimation. Recasen *et al.* proposed Endo-Depth-and-Motion [17], a pipeline for 3-D scene reconstruction from monocular videos that combines self-supervised depth estimation, photometric odometry, and volumetric fusion. Additionally, Baoru *et al.* [7] leveraged stereo image pairs to generate corresponding point cloud scenes and minimize the point distance between the two scenes. Shuwei *et al.* [22] proposed a fusion of depth prediction, appearance flow prediction, and image reconstruction to achieve robust training.

Unlike the methods discussed previously, NeRFs are based on deep neural networks that focus on learning to represent 3-D scenes, as proposed by Mildenhall *et al.* [13]. To the best of our knowledge, the application of NeRFs to surgical videos, and endoscopic videos in particular, was first discussed in the work of Wang *et al.*, resulting in an approach known as EndoNeRF [26]. EndoNeRF specifically operates on prostatectomy data and leverages the D-NeRF method [16] to render a short dynamic scene and inpaint the surgical tools. However, whereas the focus of EndoNeRF is on rendering a specific area, the focus of our research effort is on rendering an entire abdominal environment. Moreover, our research effort primarily focuses on rendering static scenes, with a minimal presence of surgical tools, surgical actions, and moving organs. In particular, we generate static scenes through the use of two NeRFs-based methods that offer fast convergence: NerfAcc [9,25] and Instant-NGP [15].

3 Method

As depicted in Fig. 2, we present a workflow to apply NeRFs to laparoscopic surgical videos. We describe each step of this workflow in the following sections.

3.1 Video Frame Extraction

In this section, we describe the process of extracting video frames from a selected laparoscopic surgical video to create a dataset. In particular, we extract a total of n_{frame} equidistant frames from the given video. These extracted frames are then stored for subsequent analysis. By varying the value of n_{frame} , we generate multiple datasets, each containing a different number of video frames. This allows us to investigate the impact of the variable n_{frame} in our experiments, which encompass a broad spectrum of sparsity levels, ranging from sparsely sampled views (with only a few frames) to densely sampled views (with numerous frames).

3.2 Camera Pose Estimation

Upon obtaining a dataset consisting of n_{frame} extracted video frames, it becomes necessary to determine the camera pose for each individual frame. In order to accomplish this task, we employ Structure-from-Motion (SfM) algorithms that

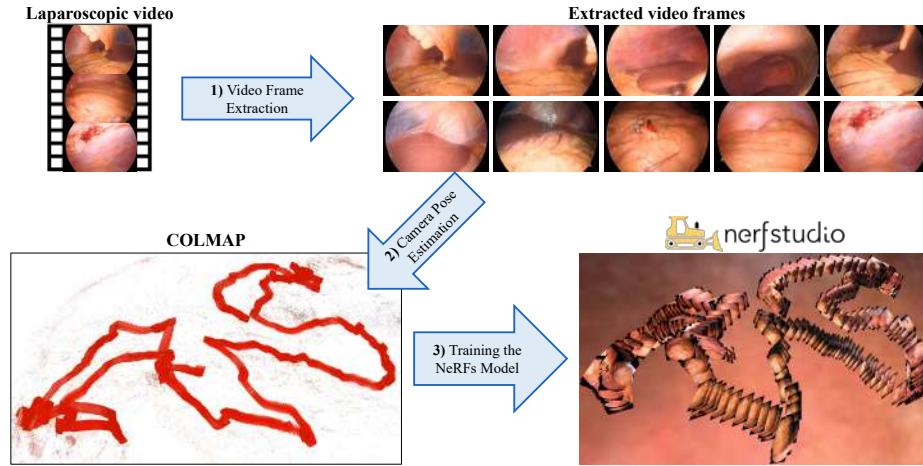


Fig. 2. NeRFs-based workflow for rendering of the abdomen from a laparoscopic surgical video. The workflow consists of three major steps, including two pre-processing steps and a training step. First, a dataset is created by extracting a total of n_{frame} video frames from the laparoscopic surgical video. Second, we estimate the camera pose for each extracted video frame, resulting in n_{pose} camera poses. The figure displays n_{pose} camera poses obtained using COLMAP after the second step. Finally, NeRFs-based models are trained using the extracted images and their corresponding camera poses, utilizing the Nerfstudio library.

make use of COLMAP⁷ [20,21] with Scale-Invariant Feature Transform (SIFT) features, as well as the hierarchical localization toolbox (HLOC)⁸ [18]. Specifically, the HLOC toolbox utilizes COLMAP with SuperPoint [2] for feature extraction and SuperGlue [19] for image matching. In addition, we also make use of LocalRF [12] as an alternative camera pose estimation method. This method is capable of handling long-range camera trajectories and large unbounded scenes. By leveraging these tools, from the initial n_{frame} frames, we can estimate camera poses for a total of n_{pose} frames. Ideally, by using LocalRF [12], we would have $n_{frame} = n_{pose}$.

3.3 Training and Mask-based Ray Sampling

For training the NeRFs-based models, we utilize the Nerfstudio [25] library, which includes the NerfAcc [9] and Instant-NGP [15] models. As depicted in Fig. 3, our approach involves sampling rays from the input image and its corresponding camera pose, with each ray passing through a pixel on the image. It is important to note that we exclude the red lines in order to prevent the occurrence of artifacts at the intersections between these red lines and the blue lines from

⁷ <https://colmap.github.io/>

⁸ <https://github.com/cvg/Hierarchical-Localization>

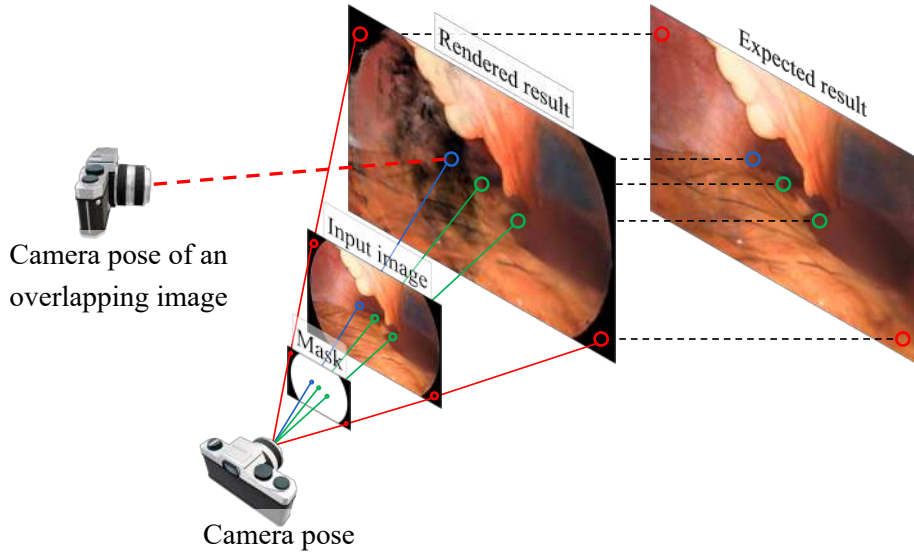


Fig. 3. Ray sampling for NeRFs training. Each line represents a ray. Green lines indicate the required sampled rays. Red lines represent rays in the black area that should not be sampled. Artifacts occur in the rendered result when an image overlaps with the black area of other images. This is demonstrated by the existence of artifacts within the blue circle, which illustrates the point where the blue line intersects with the red lines from other images (represented by the thick dashed red line). To obtain the desired result, we create a mask that only samples rays from pixels with color values.

overlapping images. To that end, we propose sampling rays with masks. Instead of creating a mask for each individual image, we opted to use a single mask that applies to all images. This one-for-all mask takes the form of a circle located at the center of the image, with the radius calculated based on the circular view of the laparoscope. This mask allows us to selectively sample rays only from pixels that contain color values. By employing this mask, the model is able to infer the content of the black areas solely based on the color information from the overlapping regions in other images.

4 Experiments

To experiment with the use of NeRFs for rendering an abdominal scene, we chose a laparoscopic surgical video from a collection of videos recorded at Ghent University Hospital. This chosen video captures the preparatory phase of a laparoscopic surgery and does not involve the use of surgical tools. It provides a laparoscopic view recorded during an examination of the abdomen, captured at a resolution of 3840×2160 pixels. To optimize the video for our purposes, we cropped out 6.6% from the left and 9.1% from the right, removing redundant

Table 1. Results of camera pose estimation. We conducted experiments with different settings and recorded the number of frames, denoted as n_{pose} , for which camera poses were successfully estimated. The highest n_{pose} (excluding LocalRF [12]) is shown in **bold**, and the second best is underscored for each n_{frame} dataset. LocalRF is capable of estimating all camera poses, albeit with a trade-off in terms of running time. Similarly, HLOC suffers from long running times. Consequently, we did not use HLOC and LocalRF for $n_{frame} = 2784$.

n_{frame}	Method	n_{pose}	Time (hh:mm:ss)
310	COLMAP	62	00:08:26
	HLOC	<u>43</u>	01:08:39
	LocalRF	310	23:20:52
1392	COLMAP	<u>376</u>	00:49:29
	HLOC	647	21:51:34
	LocalRF	1392	75:10:54
2784	COLMAP	2667	04:32:53

black areas. This resulted in video frames having a resolution of 3237×2160 pixels. The recorded footage has a duration of 93 seconds, and the frame rate is approximately 29.935 fps (frames per second), leading to a total of 2784 available video frames. All execution times were obtained on a machine equipped with a single RTX A6000 GPU and 128 GB RAM.

4.1 Camera Pose Estimation

We extract video frames from a selected laparoscopic surgical video, with n_{frame} having the following values: 310, 1392, and 2784 frames. This results in three datasets, each having a different sampling interval. By doing so, we can assess the effectiveness of SfM methods on datasets with varying levels of sparsity. The best results of COLMAP and HLOC can be found in Table 1. Furthermore, we employ LocalRF [12] to ensure we can estimate all camera poses. We have made modifications to the code provided in the LocalRF paper to restrict the sampling of rays solely to valid color pixels. Table 1 indicates that COLMAP exhibits a substantially higher speed than HLOC, with a speedup factor of 8 and 26.5 for $n_{frame} = \{310, 1392\}$, respectively. However, COLMAP only performs well in the case of a dense view with $n_{frame} = 2784$, whereas HLOC demonstrates better effectiveness in other scenarios with $n_{frame} = \{310, 1392\}$. Although LocalRF [12] is capable of estimating all camera poses, its extensive running time presents a drawback. Consequently, LocalRF is only suitable for application in sparse views, where the number of images that needs to be determined is limited.

4.2 Qualitative and Quantitative Rendering Results

Based on the results of the previous experiment, we created individual image sets for five possible n_{pose} values: $\{62, 310, 647, 1392, 2667\}$. We then trained

NeRFs-based models on each of these sets. For training, we took 90% of n_{pose} , comprising equally spaced images, as the training set. The remaining images were allocated to the evaluation set. To avoid memory limitations when working with a large number of 4K images, we down-scaled the input images by a factor of four. We trained the NerfAcc [9] and Instant-NGP [15] models using the default hyperparameter settings provided by the Nerfstudio library [25]. We observed that training with a sampling mask increased the processing time because of the underlying Nerfstudio implementation. Therefore, we adjusted the number of training iterations to be {3000, 15000, 30000, 50000, 100000}, corresponding to the number of images in each n_{pose} set, respectively.

For qualitative evaluation, Fig. 3 shows the render results obtained for a select number of images in the evaluation set. As explained in Section 3.3, training without a mask resulted in black patterns, while training with a mask improved the image quality and expanded the field of view. However, it is important to note that training with a mask currently presents a trade-off between image quality and processing time. Overall, the rendered frames from Instant-NGP [15] exhibited higher quality and less blurriness compared to NerfAcc [9]. Furthermore, among the different values of n_{pose} , the set with $n_{pose} = 2667$ (dense views) yielded the highest quality rendered frames compared to the other sets. This finding is consistent with the notion that a NeRFs-based method generally performs better when this method is applied to dense and overlapping views.

To evaluate the objective quality of the synthesized novel views, we measure the PSNR, SSIM, and LPIPS [30] scores, which are conventional image quality metrics [13,26], between the synthesized views and the corresponding ground truth video frames. The results obtained for training with a mask are presented in Table 2. We can observe that these results are in line with our qualitative findings, suggesting that Instant-NGP [15] outperforms NerfAcc [9].

5 Conclusions

In this paper, we introduced a workflow for applying state-of-the-art NeRFs-based methods to laparoscopic surgical videos in order to render a complete abdominal environment, identifying and overcoming several challenges.

A first challenge we encountered is the presence of artifacts in recorded video frames (e.g., blurriness caused by laparoscope movement), leading to failures in camera pose estimation using COLMAP. To mitigate these failures, we utilized LocalRF [12], which enabled us to estimate all camera poses but required a substantial amount of computational time (e.g., 75 hours for 1392 images).

A second challenge we faced was determining the appropriate frame resolution and the number of extracted frames (n_{frame}). Opting for frames with a higher resolution, such as 4K or even 8K [29], resulted in rendered images of a higher quality but also increased training time. This is due to the need to randomly sample rays (pixels) in each training iteration to ensure full pixel coverage. Also, larger n_{pose} values substantially increased the training time. Consequently, a trade-off exists between training time and rendered image quality.

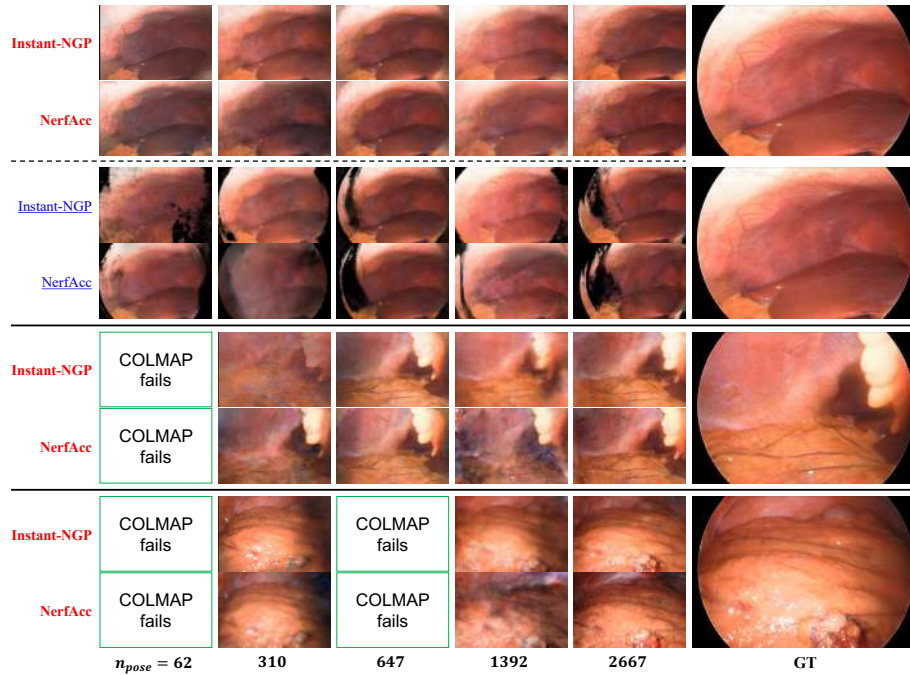


Fig. 4. Novel view synthesis for a select number of unseen video frames in the evaluation set. Training with a mask is shown in **bold**, whereas training without a mask is shown as underscored. Black patterns occur when training without a mask. The green box indicates a failure of COLMAP to estimate the camera pose of the video frame in the n_{pose} set (rendering a video frame is not possible without the correct camera perspective). The frames generated by Instant-NGP [15] are better in subjective quality and less blurry than those produced by NerfAcc [9].

A third challenge we came across is the limited view provided by a laparoscope, resulting in recorded videos that typically contain black outer areas. To address this challenge, we employed sampling with a mask as a workaround. As shown by our experimental results, the proposed one-for-all mask approach enhanced the quality of the rendered outcomes and broadened the perspective.

A fourth challenge is related to the camera trajectory followed, which is—in the context of laparoscopy surgery—usually free in nature, rather than being a regular forward-facing trajectory that captures extensive overlapping views. Such an elongated and unconstrained trajectory poses difficulties for NeRFs-based models to fully render an abdominal scene.

In future work, we will keep focussing on overcoming the primary challenge of the limited surgical view provided by a laparoscope. To that end, we aim at combining NeRFs with diffusion models [3,6,23], such as DiffRF [14], NerfDiff [5], and RealFusion [11]. Additionally, we plan to employ NeRFs to predict

Table 2. Quantitative results obtained for video frames in the evaluation set. We report the average PSNR/SSIM scores (higher is better) and LPIPS [30] scores (lower is better). The highest score for each n_{pose} evaluation set is indicated in **bold**. Overall, the Instant-NGP [15] model outperforms the NerfAcc [9] model across all n_{pose} sets.

n_{pose}	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
62	NerfAcc	12.760 \pm 0.424	0.764 \pm 0.018	0.448 \pm 0.020
	Instant-NGP	12.769 \pm 0.231	0.763 \pm 0.008	0.470 \pm 0.028
310	NerfAcc	12.520 \pm 2.439	0.674 \pm 0.084	0.619 \pm 0.097
	Instant-NGP	13.558 \pm 1.620	0.714 \pm 0.051	0.609 \pm 0.099
647	NerfAcc	13.817 \pm 1.447	0.730 \pm 0.044	0.510 \pm 0.084
	Instant-NGP	14.136 \pm 1.350	0.744 \pm 0.034	0.472 \pm 0.053
1392	NerfAcc	13.602 \pm 1.598	0.704 \pm 0.044	0.548 \pm 0.055
	Instant-NGP	14.694 \pm 1.244	0.746 \pm 0.025	0.538 \pm 0.051
2667	NerfAcc	14.226 \pm 1.526	0.713 \pm 0.039	0.529 \pm 0.058
	Instant-NGP	14.512 \pm 1.549	0.757 \pm 0.025	0.450 \pm 0.050

intraoperative changes during laparoscopic surgery, for instance addressing the issue of organ deformation between preoperative and intraoperative states.

References

1. Ali, S.: Where do we stand in AI for endoscopic image analysis? Deciphering gaps and future directions. *npj Digital Medicine* **5**(1), 184 (2022)
2. DeTone, D., Malisiewicz, T., Rabinovich, A.: Superpoint: Self-supervised interest point detection and description. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 224–236 (2018)
3. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)
4. Gerats, B.G., Wolterink, J.M., Broeders, I.A.: Depth-Supervise NeRF for Multi-View RGB-D Operating Room Images. *arXiv preprint arXiv:2211.12436* (2022)
5. Gu, J., Trevithick, A., Lin, K.E., Susskind, J., Theobalt, C., Liu, L., Ramamoorthi, R.: NerfDiff: Single-image View Synthesis with NeRF-guided Distillation from 3D-aware Diffusion. In: *International Conference on Machine Learning* (2023)
6. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
7. Huang, B., Zheng, J.Q., Nguyen, A., Xu, C., Gkouzionis, I., Vyas, K., Tuch, D., Giannarou, S., Elson, D.S.: Self-supervised depth estimation in laparoscopic image using 3D geometric consistency. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 13–22. Springer (2022)
8. Huy, P.N., Quan, T.M.: Neural Radiance Projection. In: *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5 (2022). <https://doi.org/10.1109/ISBI52829.2022.9761457>
9. Li, R., Gao, H., Tancik, M., Kanazawa, A.: NerfAcc: Efficient Sampling Accelerates NeRFs. *arXiv preprint arXiv:2305.04966* (2023)

10. 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. *The International Journal of Medical Robotics and Computer Assisted Surgery* **12**(2), 158–178 (2016)
11. Melas-Kyriazi, L., Rupprecht, C., Laina, I., Vedaldi, A.: RealFusion: 360° Reconstruction of Any Object from a Single Image. In: *Arxiv* (2023)
12. Meuleman, A., Liu, Y.L., Gao, C., Huang, J.B., Kim, C., Kim, M.H., Kopf, J.: Progressively Optimized Local Radiance Fields for Robust View Synthesis. In: *CVPR* (2023)
13. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: *ECCV* (2020)
14. Müller, N., Siddiqui, Y., Porzi, L., Bulo, S.R., Kotschieder, P., Nießner, M.: DiffRF: Rendering-guided 3d radiance field diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4328–4338 (2023)
15. Müller, T., Evans, A., Schied, C., Keller, A.: Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* **41**(4), 102:1–102:15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>, <https://doi.org/10.1145/3528223.3530127>
16. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10318–10327 (2021)
17. Recasens, D., Lamarca, J., Fácil, J.M., Montiel, J., Civera, J.: Endo-Depth-and-Motion: Reconstruction and Tracking in Endoscopic Videos Using Depth Networks and Photometric Constraints. *IEEE Robotics and Automation Letters* **6**(4), 7225–7232 (2021)
18. Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From Coarse to Fine: Robust Hierarchical Localization at Large Scale. In: *CVPR* (2019)
19. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: SuperGlue: Learning Feature Matching with Graph Neural Networks. In: *CVPR* (2020)
20. Schönberger, J.L., Frahm, J.M.: Structure-from-Motion Revisited. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
21. Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise View Selection for Unstructured Multi-View Stereo. In: *European Conference on Computer Vision (ECCV)* (2016)
22. Shao, S., Pei, Z., Chen, W., Zhu, W., Wu, X., Sun, D., Zhang, B.: Self-supervised monocular depth and ego-motion estimation in endoscopy: Appearance flow to the rescue. *Medical image analysis* **77**, 102338 (2022)
23. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International Conference on Machine Learning*. pp. 2256–2265. PMLR (2015)
24. Soler, L., Hostettler, A., Pessaux, P., Mutter, D., Marescaux, J.: Augmented Surgery: An Inevitable Step in the Progress of Minimally Invasive Surgery. *Robotic Surgery* pp. 217–226 (2021)
25. Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Kerr, J., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., McAllister, D., Kanazawa, A.: Nerfstudio: A Modular Framework for Neural Radiance Field Development. In: *ACM SIGGRAPH 2023 Conference Proceedings. SIGGRAPH '23* (2023)
26. Wang, Y., Long, Y., Fan, S.H., Dou, Q.: 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 (2022)
27. Wei, G., Feng, G., Li, H., Chen, T., Shi, W., Jiang, Z.: A novel SLAM method for laparoscopic scene reconstruction with feature patch tracking. In: 2020 International Conference on Virtual Reality and Visualization (ICVRV). pp. 287–291. IEEE (2020)
 28. Wei, G., Yang, H., Shi, W., Jiang, Z., Chen, T., Wang, Y.: Laparoscopic Scene Reconstruction Based on Multiscale Feature Patch Tracking Method. In: 2021 International Conference on Electronic Information Engineering and Computer Science (EIECS). pp. 588–592 (2021). <https://doi.org/10.1109/EIECS53707.2021.9588016>
 29. Yamashita, H., Aoki, H., Tanioka, K., Mori, T., Chiba, T.: Ultra-high definition (8K UHD) endoscope: our first clinical success. SpringerPlus **5**(1), 1–5 (2016)
 30. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
 31. Zhou, H., Jayender, J.: EMDQ-SLAM: Real-time high-resolution reconstruction of soft tissue surface from stereo laparoscopy videos. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24. pp. 331–340. Springer (2021)
 32. Zhou, H., Jayender, J.: Real-Time Nonrigid Mosaicking of Laparoscopy Images. IEEE Transactions on Medical Imaging **40**(6), 1726–1736 (2021). <https://doi.org/10.1109/TMI.2021.3065030>