



# Dynamic Depth-Supervised NeRF for Multi-view RGB-D Operating Room Videos

Beerend G. A. Gerats<sup>1,3</sup>(✉) , Jelmer M. Wolterink<sup>2</sup>,  
and Ivo A. M. J. Broeders<sup>1,3</sup>

<sup>1</sup> Centre for Artificial Intelligence, Meander Medisch Centrum,  
Amersfoort, The Netherlands

{bga.gerats, iamj.broeders}@meandermc.nl

<sup>2</sup> Department of Applied Mathematics and Technical Medical Center,  
University of Twente, Enschede, The Netherlands

j.m.wolterink@utwente.nl

<sup>3</sup> Robotics and Mechatronics, University of Twente, Enschede, The Netherlands

**Abstract.** The operating room (OR) is an environment of interest for the development of sensing systems, enabling the detection of people, objects, and their semantic relations. Due to frequent occlusions in the OR, these systems often rely on input from multiple cameras. While increasing the number of cameras generally increases algorithm performance, there are hard limitations to the number and locations of cameras in the OR. Neural Radiance Fields (NeRF) can be used to render synthetic views from arbitrary camera positions, virtually enlarging the number of cameras in the dataset. In this work, we explore the use of NeRF for view synthesis of dynamic scenes in the OR, and we show that regularisation with depth supervision from RGB-D sensor data results in higher image quality. We optimise a dynamic depth-supervised NeRF with up to six synchronised cameras that capture the surgical field in five distinct phases before and during a knee replacement surgery. We qualitatively inspect views rendered by a virtual camera that moves 180° around the surgical field at differing time values. Quantitatively, we evaluate view synthesis from an unseen camera position in terms of PSNR, SSIM and LPIPS for the colour channels and in MAE and error percentage for the estimated depth. We find that NeRFs can be used to generate geometrically consistent views, also from interpolated camera positions and at interpolated time intervals. Views are generated from an unseen camera pose with an average PSNR of 18.2 and a depth estimation error of 2.0%. Our results show the potential of a dynamic NeRF for view synthesis in the OR and stress the relevance of depth supervision in a clinical setting.

**Keywords:** Neural radiance fields · RGB-D imaging · Operating room videos

**Supplementary Information** The online version contains supplementary material available at [https://doi.org/10.1007/978-3-031-46005-0\\_19](https://doi.org/10.1007/978-3-031-46005-0_19).

# 1 Introduction

The operating room (OR) is an environment of interest for the development of sensing systems, with tasks ranging from person detection and human pose estimation [4] to domain modeling and role identification [9]. These sensing systems could automatically register adverse events and distractions [3] or monitor X-ray radiation exposure [10] to enhance patient and staff safety. Due to frequent occlusions by large devices, hanging monitors and a crowded space, detection systems often rely on video input from multiple cameras. Typically, the underlying algorithms perform better when the number of cameras is increased [2]. However, there are limitations to the number and locations of cameras in the OR due to sterility concerns and limited available space.

Neural Radiance Fields (NeRF) [7] is a powerful technology for the reconstruction of a 3D scene from a set of images that capture the scene from various camera positions. Although the introduction of this technology has caused an explosion of interest in the field of computer vision, with many follow-up studies [20], NeRF-based methods for clinical use remain largely unexplored [16]. The use of NeRF could help to overcome the limited availability of camera placement in the OR, by virtually increasing the number of cameras with view synthesis from new camera locations. Subsequently, renders of OR scenes from arbitrary camera angles could be used to improve detection algorithms.

In this paper, we explore the use of NeRF for view synthesis of dynamic scenes in the OR and we show that regularisation with depth supervision [1] increases the render quality and reduces the need for many camera positions. We find that a depth-supervised NeRF optimised with only six synchronised camera views is able to generate images of the surgical intervention from a range of camera angles. In contrast to existing depth supervision in NeRF, we directly optimise our model using measured RGB-D sensor data instead of estimated depth from a structure from motion (SfM) algorithm. Additionally, we extend the method with a notion of scene dynamics by optimising with an extra time variable enabling the reconstruction of the surgical scene at specific time intervals. We show how the dynamic NeRF could be used for region-of-interest localisation through unsupervised segmentation of objects and people in the OR.

## 2 Related Work

### 2.1 Neural Radiance Fields

NeRF is a method for volume rendering, based upon the *implicit representation* of a 3D scene in the weights of a neural network  $F_{\Theta}$  [7]. This network is generally a standard multi-layer perceptron (MLP) that takes a 5D vector  $(x, y, z, \theta, \phi)$  as input and that outputs a 4D vector  $(RGB, \sigma)$ . An input vector consists of a 3D location  $(x, y, z)$  in the captured scene and an orientation  $(\theta, \phi)$  from which this location is viewed.  $F_{\Theta}$  returns for each vector a colour  $RGB$  and a volume density  $\sigma$ . With this simple setup, NeRF can reconstruct images by casting a viewing ray from each pixel, sampling points along that ray, asking the MLP to find

the colours and densities for these points and to sum over these results. In this way, it is possible to use a discrete set of sampled points in the 3D scene, while representing the scene in continuous form. Reconstructed images are compared with ground truth images that are taken from the same camera positions. The rendering function that sums over the colours and densities is differentiable such that the MLP can be optimised by stochastic gradient descent. The loss function is often a mean squared error between the colours of the rendered and the ground truth images.

## 2.2 NeRF with Depth Priors

Although NeRF has the ability to synthesise photo-realistic images from unseen camera perspectives, the method does not guarantee to capture 3D geometry accurately. This limitation becomes visible when rendering poorly textured areas that often occur in indoor scenes [18]. It is likely that this problem will play a role when NeRF is applied to the OR.

Several solutions are proposed that involve regularisation with depth priors. Nerfing MVS [18] provides a guided optimisation scheme, where points are sampled along a viewing ray only around depth values found earlier. A sparse set of depth values is found by applying the COLMAP SfM algorithm [14] on multi-view images. The sparse sets are used to train a depth completion network that provides full sets of depth values. A more common approach is to use a depth loss that enforces NeRF to represent a large amount of volume density around ground truth depth values [12], or that enforces NeRF to terminate rays close to depth observations, provided by a depth completion network [13], the COLMAP algorithm or RGB-D data [1].

## 2.3 NeRF for Dynamic Scenes

While the original NeRF has been developed to represent static scenes, several adaptations have been proposed to give NeRF a sense of scene dynamics. D-NeRF [11] uses an additional deformation network that models point translation to a canonical configuration. With this network, it is possible to render 3D objects that change in shape over time. However, it is less suitable for the clinical setting, where people and objects can appear and disappear in the reconstructed scene. Li *et al.*, [5] add temporal latent codes to the input vectors for video synthesis. To speed up training, they propose hierarchical frame selection and importance sampling of camera rays.

## 2.4 NeRF for Clinical Interventions

Although NeRF has inspired several works in medical image computing [15, 19], it has not been widely adopted in the field of computer-assisted interventions. To the best of our knowledge, the technology has thus far only been used for 3D

reconstruction of soft tissues in robotic surgery videos [16]. Because the laparoscopic view is a single view, their method uses time-dependent modelling of neural radiance and displacement fields, based on D-NeRF. Using the stereoscopic camera of the surgical robot, the method finds depth maps along the coloured image frames. The depth information is used to constrain NeRF optimisation with an additional loss function.

### 3 Methods

#### 3.1 Dynamic Depth-Supervised NeRF

We use the depth-supervised NeRF (DS-NeRF) by Deng *et al.* [1] for building 3D reconstructions of OR scenes. This method regularises the training with an additional depth loss such that a model can be optimised with relatively few camera positions. The key idea in DS-NeRF is that most viewing rays terminate at the closest surface, which is often opaque. Therefore, most volume density should be found close to the distance of this surface along the viewing ray. DS-NeRF enforces such a distribution of volume density by minimising the KL divergence between the volume density distribution  $h_i(s)$  and a normal distribution around the ground truth depth  $d_i$  of keypoint  $x_i \in X$ :

$$\text{KL}[\mathbb{N}(d_i, \hat{\sigma}_i) \| h_i(s)], \quad (1)$$


where  $X$  is the set of all keypoints in an image for which the depth is known and  $s$  is the far endpoint of the viewing ray. The variance  $\hat{\sigma}_i$  is set to the uncertainty of the depth estimation for keypoint  $x_i$ . When depth is estimated with COLMAP, the uncertainty is calculated by re-projecting the keypoint to and from another camera position in which the keypoint is visible. In RGB-D data, however, the depth values are measurements rather than estimations. Therefore, we set  $\hat{\sigma}_i = 1.0$  for all keypoints such that each depth value is weighted equally. We sample 100K depth values in each ground truth image at random positions, where pixels that have zero depth are never included.

The OR is a dynamic environment, where objects and people move around in a controlled manner. Training NeRF with a notion of these dynamics over time is a natural choice for the synthesis of video. Inspired by DyNeRF [5], we extend DS-NeRF to generate views dependent on a time variable  $t$ . Our Dynamic DS-NeRF takes a 6D vector  $(x, y, z, \theta, \phi, t)$  as input and outputs the same 4D vector  $(RGB, \sigma)$ . In contrast to the other input variables, we do not encode the time variable. We concatenate the scalar to the positional embeddings of the locations.


#### 3.2 Dataset

The 4D-OR dataset [9] from the Technical University of Munich contains RGB-D images and camera poses from ten simulated knee replacement surgeries. We use this dataset to train the method in the reconstruction of five surgical phases (see

Fig. 1). In this way, we can evaluate the quality of synthesised views for different OR activities. The dataset contains synchronised images from six cameras that have fixed locations in the OR. Three cameras are located above the surgical field, each rotated approximately  $90^\circ$  around the yaw axis (the red dots in Fig. 2). Two of the three other cameras capture the OR from a wide perspective, while the sixth camera records from a position that is closer to the ground (the green dots).



**Fig. 1.** Five distinct phases during or before the surgery.



**Fig. 2.** Locations of the RGB-D cameras are indicated by a red or green dot. The viewing angle is directed by the arrows. Red dots indicate three cameras located above the surgical field, whereas the green dots capture the OR from very different perspectives. The scene is a coloured point cloud formed by camera projection using depth values. (Color figure online)

### 3.3 Experimental Setup

First, we train Dynamic DS-NeRF with images from all six camera positions and qualitatively inspect a set of synthesised views. To obtain this set, we ask the

algorithm to synthesise views from the three camera locations above the surgical field as well as from interpolated poses that together form a  $180^\circ$  rotation around the surgical field. Second, we train the method with five camera positions and evaluate view synthesis from the remaining unseen sixth position (the circled camera pose in Fig. 2). We evaluate the resulting images in terms of peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM) and learned perceptual image patch similarity (LPIPS). These metrics indicate image reconstruction quality in comparison to real images, where PSNR is derived from a pixel-wise mean squared error, SSIM indicates the inter-dependency of pixels that are spatially close [17], and LPIPS is based upon similarity in the deep representations of a convolutional neural network [21]. We evaluate the quality of depth perception in terms of mean absolute error (MAE) and in percentage of the ground truth depth value. Depth values equal to zero in the ground truth images are not evaluated, since these values are inaccurate measurements.

For each surgical phase, we optimise the method with  $t = \{-2, -1, 0, 1, 2\}$  and evaluate the reconstructed scene at  $t = 0$  such that we can compare the results with non-dynamic methods. All models are trained to synthesise images of  $512 \times 384$  pixels with the following hyperparameters: 768 neurons in each network layer, 4096 selected rays per batch, 192 points sampled per ray, 50K iterations and depth loss weighting factor  $\lambda_D$  set to 0.1. This number of iterations takes around 20h of training time on a single NVIDIA A40 (48 GB) GPU. At inference, the rendering of a single frame requires 22s.

## 4 Results

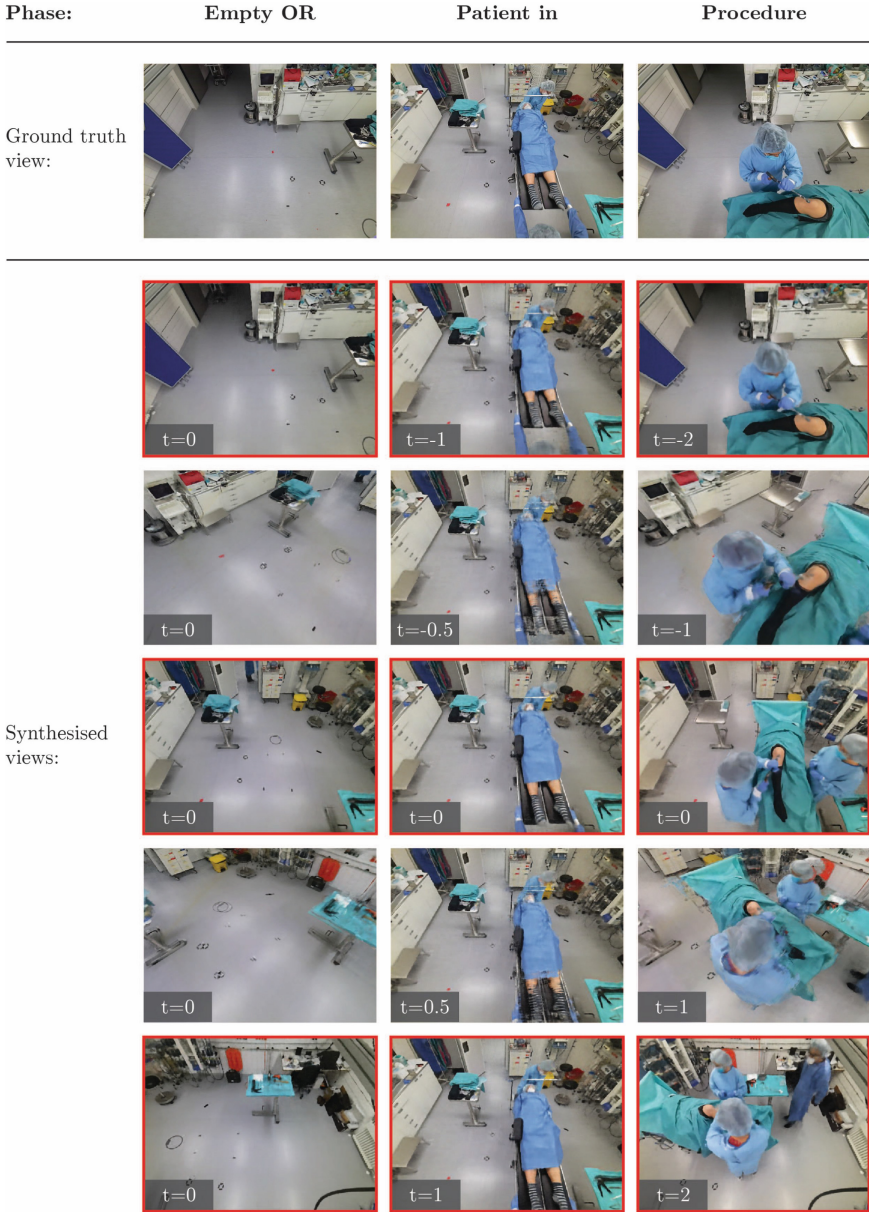
### 4.1 Qualitative Analysis

The synthesised views for three surgical phases are given in Fig. 3. When comparing the top-row synthesised views with the ground truth images, it can be seen that Dynamic DS-NeRF is able to reconstruct the scenes independently of the surgical phase. Lighting conditions, such as reflections on the floor and shadows, are realistically rendered. However, the reconstructed scene looks smoothed with missing details, e.g., in the keyboard of the mobile monitor.

Synthesised views from the interpolated camera poses (images without red borders) correctly grasp the geometry of the scenes and find realistic lighting conditions. For the “patient in” in phase, we interpolate the time values at  $t = -0.5$  and  $t = 0.5$ . The moving patient bed is correctly positioned in these frames, at an interpolated location. However, the interpolated camera positions and time values induce a number of artefacts. First, various objects seem to be misaligned. For example, the tape on the floor and the instrument table are not always represented correctly. Second, fine-grained details, such as the surgical instruments in the hands of the left physician, are missing.

### 4.2 Quantitative Analysis

Results of the quantitative analysis can be found in Table 1. On average, Dynamic DS-NeRF is able to synthesise views from the unseen camera pose



**Fig. 3.** Dynamic DS-NeRF synthesised views for three phases in the OR at time intervals  $t$ . For “empty OR”, the virtual camera rotates  $180^\circ$  around the surgical field. For “patient in”, the camera position is static, but the time intervals are interpolated. For “procedure”, the images differ in camera location and time value. The top row displays the ground truth images for the starting camera pose. Views with a red border are generated from camera poses with which the algorithm is trained, corresponding to the red camera positions in Fig. 2. (Color figure online)

with 18.2 PSNR, 0.61 SIMM and 0.41 LPIPS. The image quality of the rendered views on the 4D-OR dataset differs per phase and ranges in PSNR from 17.5 for “patient in” to 19.6 for “empty OR”. This shows that the image quality is dependent on the complexity of the surgical scene. Table 2 provides a comparison with a naive baseline method for view synthesis that uses the projection of pointclouds to reconstruct the images. In this method, the five training images are projected into a single pointcloud, which is reprojected to the sixth camera position. Figure 4 shows a left-right comparison between the baseline images and Dynamic DS-NeRF for two surgical phases. It can be seen that our method is able to generate proper colours for locations where the pointcloud has no presence. This results in more realistic-looking images and better image quality scores.

**Table 1.** Evaluation metrics comparing Dynamic DS-NeRF synthesised views with unseen ground truth RGB-D images.


Surgery phase	Colour Image			Depth Map	
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MAE (in cm) $\downarrow$	Error (in %) $\downarrow$
Empty OR	19.6	0.70	0.35	2.5	0.86
Patient in	17.5	0.59	0.39	3.2	1.28
Draping	18.4	0.62	0.43	4.3	1.76
Cleaning	17.7	0.58	0.42	5.6	2.66
Procedure	17.7	0.58	0.44	5.5	3.43
<b>Average</b>	<b>18.2</b>	<b>0.61</b>	<b>0.41</b>	<b>4.2</b>	<b>2.00</b>

**Table 2.** Ablation of Dynamic DS-NeRF where dynamics and depth supervision are disabled consecutively. The last row gives the performance of a baseline method.


Method	Colour Image			Depth Map	
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MAE (in cm) $\downarrow$	Error (in %) $\downarrow$
Dynamic DS-NeRF	<b>18.2</b>	<b>0.61</b>	0.41	<b>4.2</b>	2.00
DS-NeRF	18.1	<b>0.61</b>	<b>0.39</b>	4.3	<b>1.98</b>
NeRF	13.9	0.45	0.59	48.4	18.91
Pointcloud (baseline)	12.0	0.41	0.59	10.8	4.38

An ablation study is provided in Table 2, where dynamics and depth supervision are disabled consecutively. It can be seen that the dynamic extension does not negatively affect view synthesis performance. The presence of depth supervision has a large positive impact on image quality, with PSNR increasing from 13.9 to 18.1. The depth estimation error decreases drastically from 18.91

to 1.98%. Figure 5 displays the estimated depth for an unseen camera position in the “procedure” phase in comparison with the ground truth depth channel captured from the same camera pose. It can be seen that Dynamic DS-NeRF is able to grasp the geometry of the captured scene accurately. Moreover, the algorithm is able to generate depth values that are not present in the ground truth image due to the depth sensor’s hexagon shape or sensor artefacts (e.g., the zero-valued “shadows”). These results show that depth supervision helps NeRF to reconstruct the scene’s 3D geometry accurately, resulting in a higher quality of synthesised images.



**Fig. 4.** Left-right comparison between the pointcloud projection as baseline method (left) and output of Dynamic DS-NeRF (right) for an unseen camera pose in two surgical phases.

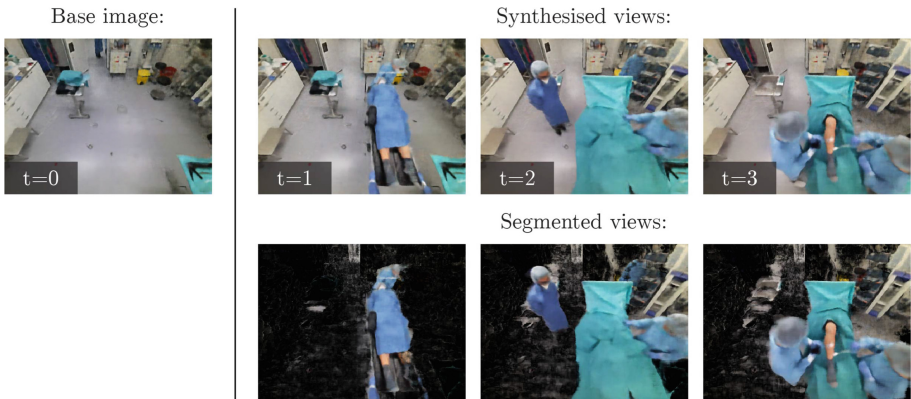


**Fig. 5.** Depth estimation output of Dynamic DS-NeRF for the “procedure” phase (right) in comparison with the ground truth depth channel (left). Color bar displays distance in cm. (Color figure online)

### 4.3 Dynamic NeRF for Segmentation

Besides the modelling of consecutive video frames, a dynamic NeRF could also represent the scene at distinct moments in time. We use the algorithm to reconstruct an empty OR at  $t = 0$  and the other phases at  $t = \{1, 2, 3, 4\}$ . Hence, a

single model is optimised to represent these five distinct phases jointly. The top row in Fig. 6 displays the model output at four time intervals. To demonstrate that Dynamic DS-NeRF could be used for tasks other than view synthesis, we use this model configuration for unsupervised segmentation of objects and people. The representation of NeRF at  $t = 0$  is considered the base image. This image consists of an empty OR, and we assume that its materials are not of interest as long as they do not move. At this time interval, we sample points along camera rays as usual and store the densities  $\sigma^{(0)}$ . For the other phases, we sample the densities  $\sigma^{(t)}$  of the same point locations at the corresponding time interval. To segment the point locations, we subtract the density of the base image from the density of each other phase:  $\sigma^{(t)} - \sigma^{(0)}$ . The resulting densities correspond to material that is present at  $t$ , but that was not present in the empty OR. We used these resulting densities to render the segmented views. It can be seen that the static floor and left counter disappear from the segmented view, while the surgeons, patient and the anaesthetic tower remain present. Note that we do not simply obtain a difference image, but use the 3D estimated density of the Dynamic DS-NeRF model.



**Fig. 6.** Top row: synthesised views by Dynamic DS-NeRF, where each surgical phase is reconstructed at a specific time interval  $t$ . Bottom row: segmented views, constructed by subtracting the volume densities at  $t = 0$  from the densities at  $t$ .

## 5 Discussion

In this work, we explored the use of NeRF for reconstructing surgical scenes in the OR with multi-view RGB-D images from the 4D-OR dataset [9]. We showed that the original NeRF does not provide optimal reconstruction results and that the use of depth supervision benefits image quality and removes the necessity to train the algorithm with tens of camera positions, which are difficult

to obtain during a surgical procedure. With an additional time variable at the input vector, we showed the possibility of rendering views with a notion of scene dynamics without affecting image quality. When virtually rotating around the surgical field, the synthesised images remain geometrically consistent, even at interpolated camera locations and time values. On the other hand, the images miss fine details and contain artefacts. Training with larger image resolutions or geometric priors can potentially help to produce views with higher quality.

Dynamic DS-NeRF can generate views from an unseen position above the surgical field with 18.2 PSNR, 0.61 SSIM, 0.41 LPIPS and an average depth error of 2.0%. The method provides better results in comparison to a baseline method that is based on pointcloud projection. Besides a baseline model, we compared three configurations of NeRF: the original NeRF, DS-NeRF and a dynamic DS-NeRF. From the wide range of NeRF architectures available [20], we hypothesise that depth supervision and dynamics are most relevant to the clinical environment. Nevertheless, the exploration of other models, such as pixelNeRF and RegNeRF, would be an interesting follow-up. Other developments in this field are of interest as well, such as methods for speeding up the training of NeRF algorithms [8].

We envision several potential uses for NeRF in the synthesis of OR images or videos. First, the technology could be used for rendering virtual environments displaying real surgeries in 3D. A spectator could watch a video-recorded procedure while moving the virtual camera to arbitrary camera positions in order to obtain a better view. This could enhance the reviewing or learning process. Second, when combining the 3D reconstruction with virtual reality, one could construct a training exercise for new surgeons or OR staff that shows a photo-realistic instead of simulated environment. This could increase the immersiveness and effectiveness of the training. Third, the implicit neural representations of the NeRF-model can be used for further processing. We showed briefly how these can be used for unsupervised segmentation of relevant objects and people. The resulting pointcloud could help to reduce the number of candidate locations for the detection of objects and human poses. Last, renders from new camera positions could be used to virtually increase the number of cameras in a dataset. This is particularly relevant for OR video datasets, due to limitations for camera placement.

In conclusion, Dynamic DS-NeRF is able to synthesise views of a dynamic surgical field in which the 3D geometry is captured accurately. Depth supervision with RGB-D sensing data increases render quality drastically while requiring fewer camera positions, making the technology applicable to clinical environments.

**Declarations.** This work was sponsored by Johnson & Johnson MedTech. Jelmer M. Wolterink was supported by NWO domain Applied and Engineering Sciences VENI grant (18192).

## References

1. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised NeRF: fewer views and faster training for free. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12882–12891 (2022)
2. Gerats, B.G., Wolterink, J.M., Broeders, I.A.: 3D human pose estimation in multi-view operating room videos using differentiable camera projections. *Comput. Methods Biomech. Biomed. Eng. Imaging Visual.* 1–9 (2022)
3. Goldenberg, M.G., Jung, J., Grantcharov, T.P.: Using data to enhance performance and improve quality and safety in surgery. *JAMA Surg.* **152**(10), 972–973 (2017)
4. Hansen, L., Siebert, M., Diesel, J., Heinrich, M.P.: Fusing information from multiple 2D depth cameras for 3D human pose estimation in the operating room. *Int. J. Comput. Assist. Radiol. Surg.* **14**(11), 1871–1879 (2019)
5. Li, T., et al.: Neural 3D video synthesis from multi-view video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5521–5531 (2022)
6. Mildenhall, B., et al.: Local light field fusion: practical view synthesis with prescriptive sampling guidelines. *ACM Trans. Graph.* **38**(4), 1–14 (2019)
7. 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: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12346, pp. 405–421. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58452-8\\_24](https://doi.org/10.1007/978-3-030-58452-8_24)
8. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* **41**(4), 1–15 (2022)
9. Özsoy, E., Örnek, E.P., Czempel, T., Tombari, F., Navab, N.: 4D-or: semantic scene graphs for or domain modeling. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. LNCS, vol. 13437, pp. 475–485. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-16449-1\\_45](https://doi.org/10.1007/978-3-031-16449-1_45)
10. Padoy, N.: Machine and deep learning for workflow recognition during surgery. *Minimally Invasive Therapy Allied Technol.* **28**(2), 82–90 (2019)
11. 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)
12. Rematas, K., et al.: Urban radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12932–12942 (2022)
13. Roessle, B., Barron, J.T., Mildenhall, B., Srinivasan, P.P., Nießner, M.: Dense depth priors for neural radiance fields from sparse input views. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, June 2022
14. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4104–4113 (2016)
15. Sun, Y., Liu, J., Xie, M., Wohlberg, B., Kamilov, U.S.: Coil: coordinate-based internal learning for tomographic imaging. *IEEE Trans. Comput. Imaging* **7**, 1400–1412 (2021)
16. Wang, Y., Long, Y., Fan, S.H., Dou, Q.: Neural rendering for stereo 3D reconstruction of deformable tissues in robotic surgery. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. LNCS, vol. 13437, pp. 431–441. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-16449-1\\_41](https://doi.org/10.1007/978-3-031-16449-1_41)

17. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
18. Wei, Y., Liu, S., Rao, Y., Zhao, W., Lu, J., Zhou, J.: Nerfingmvs: guided optimization of neural radiance fields for indoor multi-view stereo. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5610–5619 (2021)
19. Wolterink, J.M., Zwienenberg, J.C., Brune, C.: Implicit neural representations for deformable image registration. In: *International Conference on Medical Imaging with Deep Learning*, pp. 1349–1359. PMLR (2022)
20. Xie, Y., et al.: Neural fields in visual computing and beyond. *Comput. Graph. Forum* **41**, 641–676 (2022). Wiley Online Library
21. 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)