Eurographics Workshop on Visual Computing for Biology and Medicine (2023)
C. Hansen, J. Procter, and R. G. Raidou (Editors)

Supplementary material for NeRF for 3D Reconstruction from X-ray Angiography: Possibilities and Limitations

K.W.H. Maas[1], N. Pezzotti[1,2], A.J.E. Vermeer[1,3], D. Ruijters[1,2], and A. Vilanova[1]

[1]Eindhoven University of Technology, The Netherlands
[2]Philips Healthcare, The Netherlands
[3]Catharina Ziekenhuis Eindhoven, The Netherlands

1. Circular heatmap tool

Quality metrics, such as PSNR or SSIM, represent specific global aspects of the images. However, for navigation applications, perceived topological accuracy is considerably more important than precision. This perceived accuracy cannot be easily captured in metrics [RHM*11]. Therefore, we also evaluate the effects of the experimental parameters on the visual quality of the 2D projections guided by image quality metrics. An interactive tool was developed to evaluate quality based on the camera positions with interactive circular heatmaps. Our openly available tool is available at https://nerfforangiography.netlify.app/ and shown in Figure 1. The explanation of the circular heatmap is provided in Section 5.2 of the main paper. The circular heatmap ① contains regions where each region represents one spherical camera position. Camera positions are sampled for every 10° of the sphere, independently from the number of camera positions used for reconstruction. The regions of the heatmap are colored based on the user-defined quality metric score of the respective 2D predicted projections. For example, in the figure, the heatmap is colored based on the SSIM metric. This visual inspection overcomes the limitations of the individual metric scores, as users can identify the visual differences relevant to them. Hovering over the regions displays the ground-truth, predicted, and difference images for the camera position (θ, ϕ), indicated with ②. The difference image is computed as the absolute pixel-wise difference between the ground truth and the predicted image. To analyze the parameter-dependent trends, the settings represent the different parameters of the experiments, shown at ③. These settings include limited angle projections, number of sparse projections, low or high vessel sparsity, and binary or background imaging. The orientation parameters of the sphere can also be varied.

2. Projection experiments results

Figure 2 shows the Dice scores and LPIPS score for the projections experiment. Note that the y-axis is scaled differently for some of the metrics to make the differences between the lines more visible. The other metric scores for this experiment are reported in Figure 4 of the main paper.

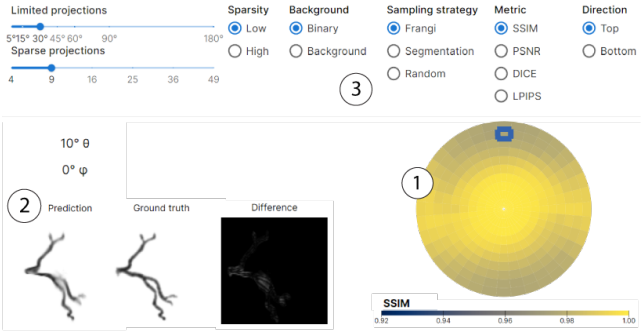

Figure 1: The interactive circular heatmap tool. ① the circular heatmap component, where one region is selected. ② the angles and difference predicted and the ground-truth image of the selected region. ③ the parameter option settings.

3. Vessel sparsity experiments

Table 1 displays the table of sparse and limited angle projections for the high-sparsity dataset. It supports the trend that is shown in Figure 6 of the main paper.

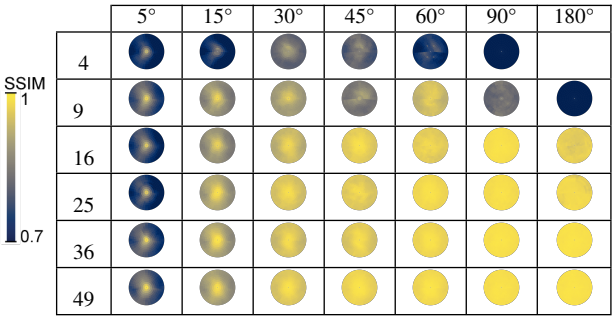

Table 1: Heatmaps for the high-sparsity dataset. The rows of the tables represent the sparse projections and the columns the limited angle projections.

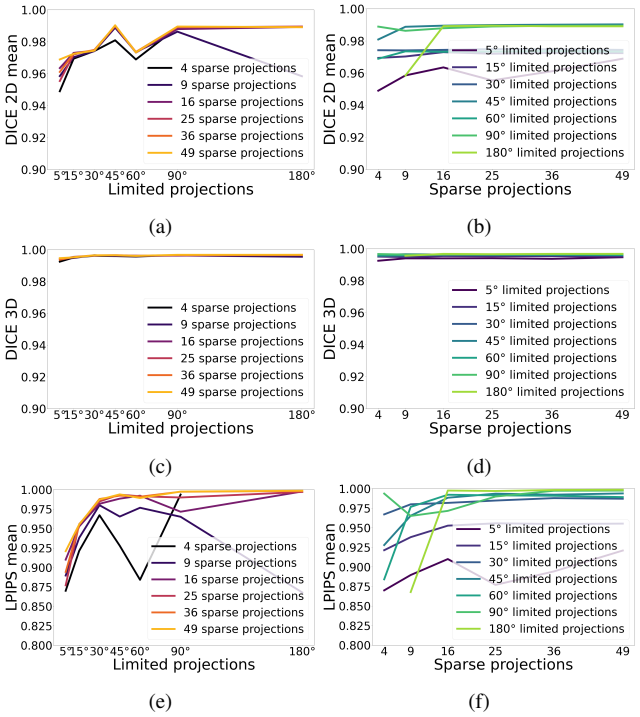

Figure 2: Relation between limited angle and sparse projections and specified metric scores for the Projections experiment. Limited angle projections in Dice 2D (a), Dice 3D (c), and LPIPS (e). Sparse projections in Dice 2D (b), Dice 3D (d) and LPIPS (f).

4. Sampling for low-sparsity and high-sparsity datasets results

Figure 3 shows the effects of the Frangi, segmentation and random sampling approaches on the low-sparsity dataset. Only the angular ranges below 30° are reported, as the vessel sparsity experiment (see Section 6.2 main paper) proved that larger angular ranges already provide mostly optimal performance for both datasets. As can be seen, no significant differences are there between the approaches, in contrast to the high-sparsity as shown in Figure 10 of the original paper. The Frangi and segmentation methods slightly outperform the random sampling approach. The average training time for the low-sparsity dataset is 52, 50, and 65 minutes for the Frangi, segmentation, and random sampling approaches, respectively. Overall, we can conclude that Frangi sampling can benefit the reconstruction of the low-sparsity dataset slightly, as it decreases the rendering time slightly, while maintaining quality.

The average rendering time per experiment for the high-sparsity dataset is 47, 46, and 54 minutes for the Frangi, segmentation, and random sampling approaches, respectively. In this case, we can conclude that the sampling approaches do not significantly influence the rendering time. Especially in the binary case, we rather expect the empty space skipping of NeRFAcc to be crucial for the rendering time. Figure 4 shows the qualitative examples for the Frangi and Random sampling approaches for the high-sparsity dataset. The setting is 4 sparse projections and 15° limited projections. It can be observed that the random sampling approach leads

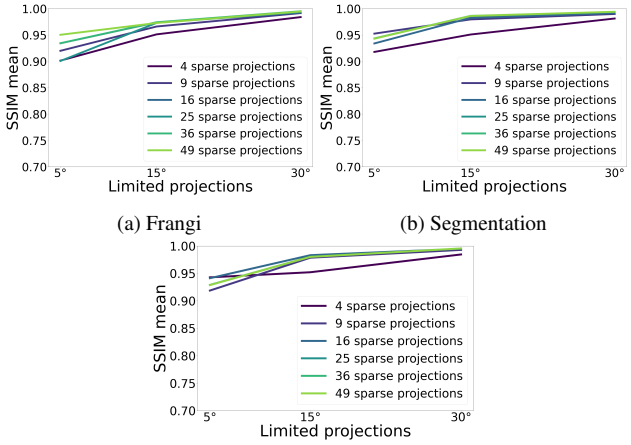

Figure 3: Relation between limited angle projections and SSIM for (a) Frangi, (b) segmentation, and (c) random sampling approaches for the low-sparsity dataset.

to a more accurate topological reconstruction, whereas the Frangi approach leads to more noise in the empty space.

5. Model architecture for vessel sparsity

This section discusses the exploration of different model architectures.

5.1. Parameter description

The decision of the architecture of the MLP for NeRF can contribute largely to the reconstruction quality and time [MESK22]. This parameter explores the impact of architecture design on these two factors. A practical example is the trade-off between quality and time when reconstructing vascular morphologies with sparse vessels, as not all vessels may be of interest in clinical applications, such as navigation. For example, major vessels may be of higher importance than smaller vessels for coronary angiography [CC00]. Three types of architectures are defined. The first architecture is the default architecture with 4 layers and 128 hidden units, which is

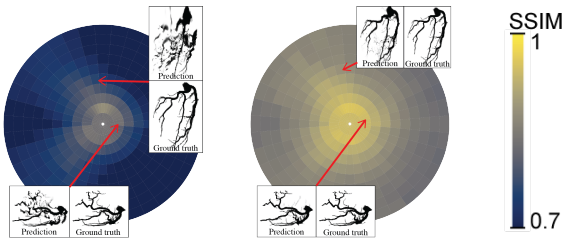

Figure 4: Heatmaps for the high-sparsity vessel morphology for sparse 4 and limited 15° for the Frangi (a) and random (b) sampling techniques.

commonly used across different applications for NeRFs [MST*21]. To identify whether the width of the architecture can improve reconstruction time, we also define the lower-width architecture with 4 layers and 64 hidden units. Following the previous works, we expect this network with lower width to provide lower-quality results with a slight decrease in computation time. Lastly, we inspect a lower-depth architecture with 2 layers and 128 hidden units. Similarly to the width architecture, this less deep architecture will improve computation time. In contrast, we only expect a small decrease in quality with respect to the default architecture. Generally, it is in our interest to analyze these different architectures to identify possible acceleration approaches, as well as their effect on reconstruction quality.

5.2. Results

We analyze the size of the MLP architecture with respect to the quality of the reconstructions, both for the low and high-sparsity datasets. A lower width architecture (2 layers and 64 neurons), default architecture (4 layers and 128 neurons), and lower depth architecture (2 layers and 128 neurons) are optimized for both sparse and limited projection settings.

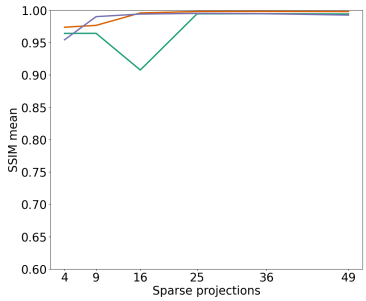

(a) Low sparsity

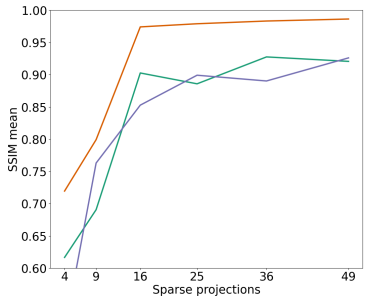

(b) High sparsity

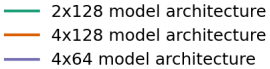

Figure 5: Relation between sparse projections and specified metric scores for the Model architecture experiment. The metrics are computed based on the average score of the 2D projections.

The metric scores in relation to the sparse projections for the low- and high-sparsity datasets are shown in Figure 5. The scores are averaged for the different limited projection parameters. For the low-sparsity dataset (Figure 5a), it can be observed that the decrease in width or depth of the architecture does not majorly affect performance. This may indicate that the simple vessel structures, such as the major vessels in the low-sparsity dataset, can still be learned by smaller architectures. The architecture size for the high-sparsity structure does impact the reconstruction quality, as can be seen in Figure 5b. Once more, a decrease in width or depth shows similar decreases in performance. One qualitative case for sparse projections 49 and 180°, reported in Figure 6, shows that, indeed, the smaller vessels of the structures are not reconstructed with the smaller architectures. It also shows that the lower-width architecture (see Figure 6b) leads to more noise in space than the lower-depth architecture (see Figure 6c) even though their performance scores are similar. This case may indicate that width indeed has a larger effect than depth on reconstruction quality as known from literature [MESK22]. Overall, we can report that, for high-precision applications, larger MLP architectures may be preferred for the reconstruction of sparse vessels or precise details. However, for visualization applications, smaller architectures may be preferred to decrease rendering time.

(a) Default architecture

(b) Lower-width architecture

(c) Lower-depth architecture

Figure 6: Comparison of reconstructions of high-sparsity dataset with different architectures sizes.

References

[CC00] CHEN S. J., CARROLL J. D.: 3-d reconstruction of coronary arterial tree to optimize angiographic visualization. *IEEE transactions on medical imaging 19*, 4 (2000), 318–336. 2

[MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG) 41*, 4 (2022), 1–15. 2, 3

[MST*21] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHI R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM 65*, 1 (2021), 99–106. 3

[RHM*11] RUIJTERS D., HOMAN R., MIELEKAMP P., VAN DE HAAR P., BABIC D.: Validation of 3d multimodality roadmapping in interventional neuroradiology. *Physics in Medicine & Biology 56*, 16 (2011), 5335–5354. 1