Neural Light Field 3D Printing: Supplemental Document

QUAN ZHENG, Max Planck Institute for Informatics, Germany
VAHID BABAEI, Max Planck Institute for Informatics, Germany
GORDON WETZSTEIN, Stanford University, USA
HANS-PETER SEIDEL, Max Planck Institute for Informatics, Germany
MATTHIAS ZWICKER, University of Maryland, College Park, USA
GURPRIT SINGH, Max Planck Institute for Informatics, Germany

CCS Concepts: • **Applied computing** → **Computer-aided manufacturing**; • **Computing methodologies** → *Neural networks*; Volumetric models.

Additional Key Words and Phrases: Volumetric display, light field, neural networks, 3D printing, computational fabrication

ACM Reference Format:
Quan Zheng, Vahid Babaei, Gordon Wetzstein, Hans-Peter Seidel, Matthias Zwicker, and Gurprit Singh. 2020. Neural Light Field 3D Printing: Supplemental Document. *ACM Trans. Graph.* 39, 6, Article 207 (December 2020), 3 pages. https://doi.org/10.1145/3414685.3417879

1 TRAINING DYNAMICS

To assess the achieved quality with respect to time cost, we compare our method and Layered3D on all five scenes. For each scene, one view is held out for test. The maximum iteration of Layered3D is set to 15 and we run it with the same settings as in the original work. Our neural network is trained for 3000 epochs with 20 batches per epoch and the batch size is 1200. Figure 1 shows the PSNR and SSIM measurements on the test view of each scene. The dashed line stands for the final quality of Layered3D after optimization and the locations of solid dots on dashed lines denote the optimization time. Our approach runs for a longer time in training, but its performance exceeds Layered3D soon and improves further along with the training process.

2 ADDITIONAL SIMULATION RESULTS

In addition to the three scenes shown in Figure 4 of the main paper (see Section 6.3 of the main paper), we present additional results on the *Butterfly* scene and the *Car* scene in Figure 2. They are produced by the baseline grid-based approach, Layered3D and our approach. The grid resolution for the baseline approach is set to $512\times 384\times 45$. We solve its optimization with the constrained large-scale trust-region reflective linear [Coleman and Li 1996] and the SGD solver Adam [Kingma and Ba 2014] respectively. For Layered3D, we assign it with five layers and use the same settings as the original paper. Our approach utilizes a D8W512 neural network. The display thickness is 12.5 mm for the simulations. Both results of the baseline show blurriness and residual noise around local structures (see blue insets of the *Butterfly* scenes) and cause faint highlight on the roof of the car (see blue insets of the *Car* scenes), whereas Layered3D suffers from ringing artifacts in multiple closeup regions. Moreover, both the baseline and Layered3D miss the micro details on the wings of the butterfly (the orange insets). By contrast, our approach generally preserves local details and exhibits a better visual quality. We tabulate numerical quality measurements of all evaluated methods in Table 1, where our approach achieves a higher numerical performance.

REFERENCES

Thomas F Coleman and Yuying Li. 1996. A reflective Newton method for minimizing a quadratic function subject to bounds on some of the variables. *SIAM Journal on Optimization* 6, 4 (1996), 1040–1058.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Authors' addresses: Quan Zheng, Max Planck Institute for Informatics, Saarbrücken, Germany, qzheng@mpi-inf.mpg.de; Vahid Babaei, Max Planck Institute for Informatics, Saarbrücken, Germany; Gordon Wetzstein, Stanford University, USA; Hans-Peter Seidel, Max Planck Institute for Informatics, Saarbrücken, Germany; Matthias Zwicker, University of Maryland, College Park, USA; Gurprit Singh, Max Planck Institute for Informatics, Saarbrücken, Germany.

0730-0301/2020/12-ART207 $15.00
https://doi.org/10.1145/3414685.3417879

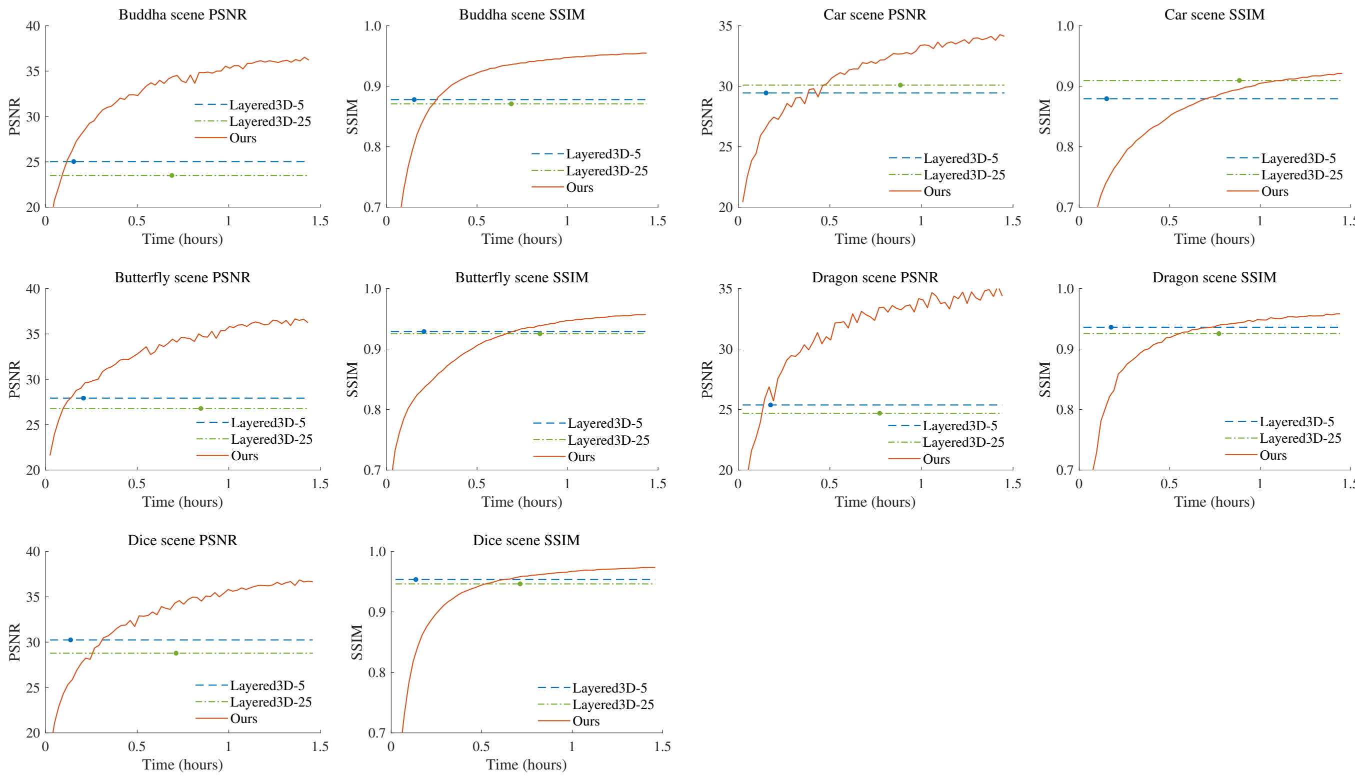

Fig. 1. PSNR and SSIM measurements of our approach and Layered3D on test views of five scenes. Dashed lines of Layered3D with 5 layers and 25 layers are the final quality measurements after optimization. Note Layered3D runs on CPU and our approach runs on GPU.

Table 1. Numerical quality measurements of all evaluated methods on the *Butterfly* and the *Car* scenes. The evaluation is conducted on the rendered test views. Images are evaluated by PSNR, SSIM (higher is better), and MAPE (lower is better).

Scene	PSNR				SSIM				MAPE			
	Grid, Linear	Grid, SGD	Layered3D	Ours	Grid, Linear	Grid, SGD	Layered3D	Ours	Grid, Linear	Grid, SGD	Layered3D	Ours
Butterfly	25.2485	27.4109	27.9284	**36.1149**	0.9203	0.9387	0.9290	**0.9518**	0.1103	0.0912	0.0943	**0.0457**
Car	28.4675	30.9863	29.4445	**34.2322**	0.8991	0.9023	0.8793	**0.9240**	0.0674	0.0533	0.0627	**0.0362**

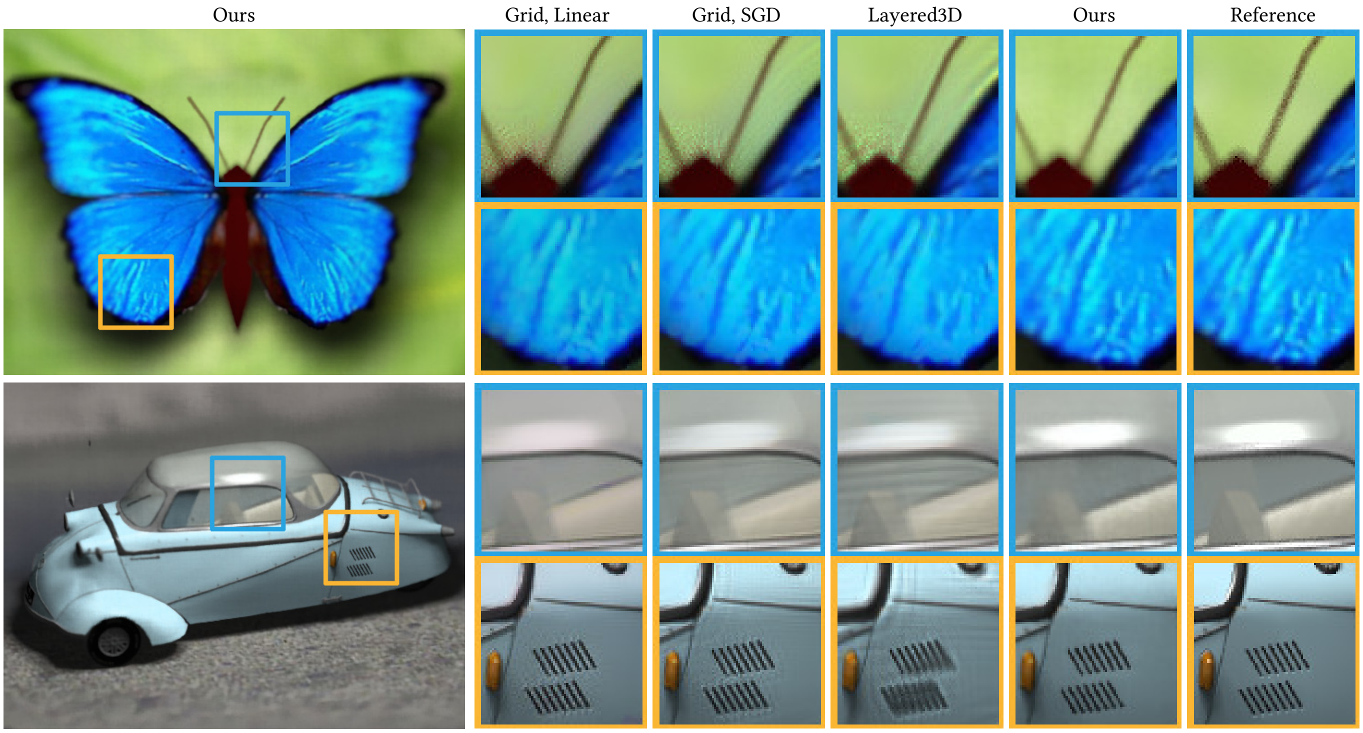

Fig. 2. We compare the visual quality of the rendered test views produced by the baseline, Layered3D and our approach on the *Butterfly* scene and the *Car* scene. We show the results of the baseline obtained by the linear solver and the SGD solver respectively. Our approach and the baseline using the SGD solver are trained with Adam for 3000 epochs. The batch size is 1200 for our method and 500K for the baseline. Also. the maximum iteration number for Layered3D and the baseline using the linear solver is 15.