Instant Continual Learning of Neural Radiance Fields - Supplementary Materials

Ryan Po Zhengyang Dong Alexander W. Bergman Gordon Wetzstein
Stanford University
{rlpo, awb, leozdong, gordonwz}@stanford.edu

1. Additional Results

1.1. Degradation of early task

We show additional qualitative results on the degradation of earlier views under the continual learning setting for each method in Supp. Fig. 2. Naive methods such as NeRF-Incre and EWC suffer from catastrophic forgetting. Although the early view is reconstructed well when the current view is part of the training views of the current task, reconstruction quality quickly degrades as the training task moves to other regions of the 3D scene. Our method and MEIL-NeRF [1] maintains relatively high quality throughout all stages of training.

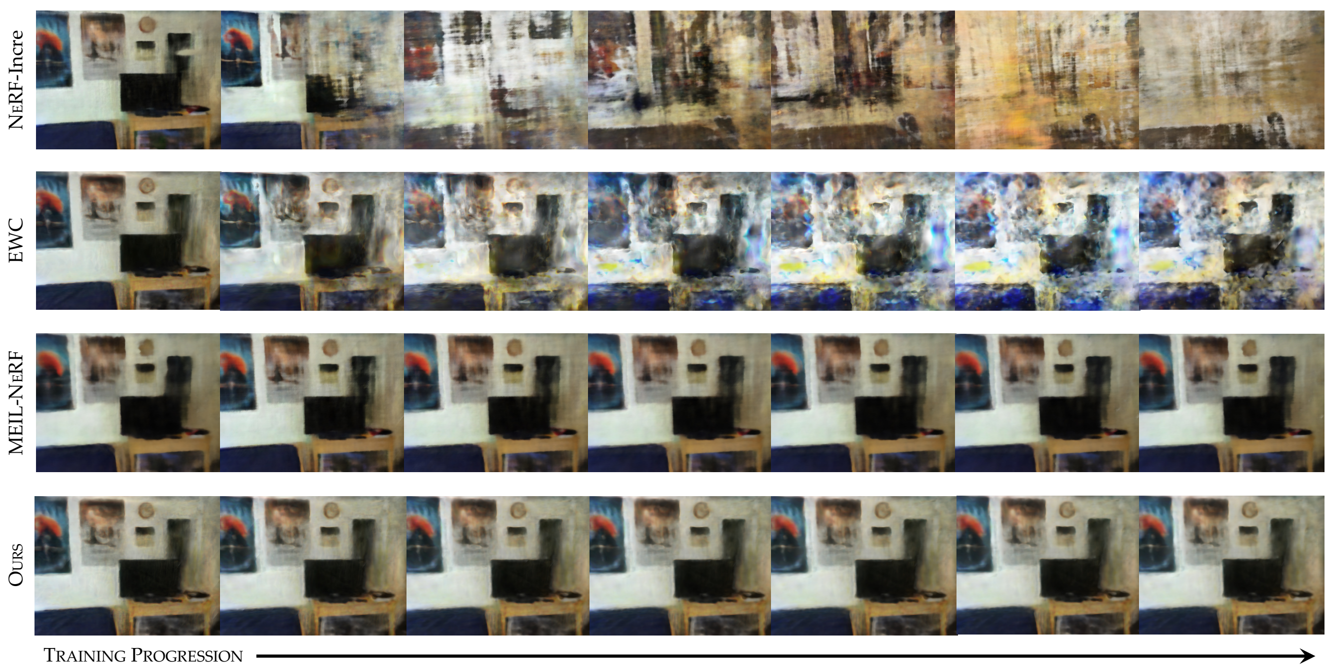

Figure 1. Reconstructed view of earlier task during different stages of training. Quantitative results shown in Figure 6 are obtained from evaluating PSNR of the right-most column.

1.2. Ablation on # of tasks

We perform an ablation over the number of tasks used for the continual learning setting, showing results in Table 1. As the number of tasks increases, the continual learning setting becomes more challenging, causing performance of our method to drop minimally as number of tasks increases. Naive method of NeRF-Incre performs poorly in the continual learning setting even for lower number of tasks.

	# of Tasks			
Method	2	5	10	20
NeRF-Incre	12.30	16.40	13.70	13.43
Ours	24.29	24.27	24.13	23.87

Table 1. Ablation on # of tasks in continual learning setting. Results obtained from 0101 scene of the ScanNet dataset [2].

1.3. Reconstruction quality for individual tasks

Table 2 shows reconstruction PSNR for scene ScanNet-0101 for every single task after training has concluded. As expected, non-replay baselines (NeRF-Incre, iNGP-Incre, EWC) suffer from catastrophic forgetting, where earlier tasks are reconstructed poorly, but having high reconstruction PSNR for the final task. Replay based methods such as MEIL-NeRF and ours maintain high PSNR across all tasks, with our method achieving better performance across all tasks.

Method	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6	Task 7	Task 8	Task 9	Task 10
NeRF-Incre	13.92	13.17	11.01	10.94	12.04	13.84	15.96	18.73	17.43	25.01
iNGP-Incre	14.86	15.02	12.89	15.15	16.50	18.21	18.40	21.07	18.68	25.37
EWC	13.79	14.07	12.66	14.66	16.50	18.09	18.46	20.64	17.98	25.34
MEIL-NeRF	23.84	24.29	22.95	22.85	24.10	24.72	25.47	26.31	25.57	24.68
Ours	24.19	26.33	24.74	24.32	26.12	26.08	26.78	27.15	25.96	24.92

Table 2. Reconstruction PSNR of individual tasks after completion of training over all tasks in the continual learning setting.

2. Code and Implementation Details.

We implement our method as a modular component of the Nerfstudio pipeline [4]. Our code includes custom data parsers for processing datasets to be suited for the continual learning setting described in the paper. All code and custom data formats will be made available upon acceptance of the paper.

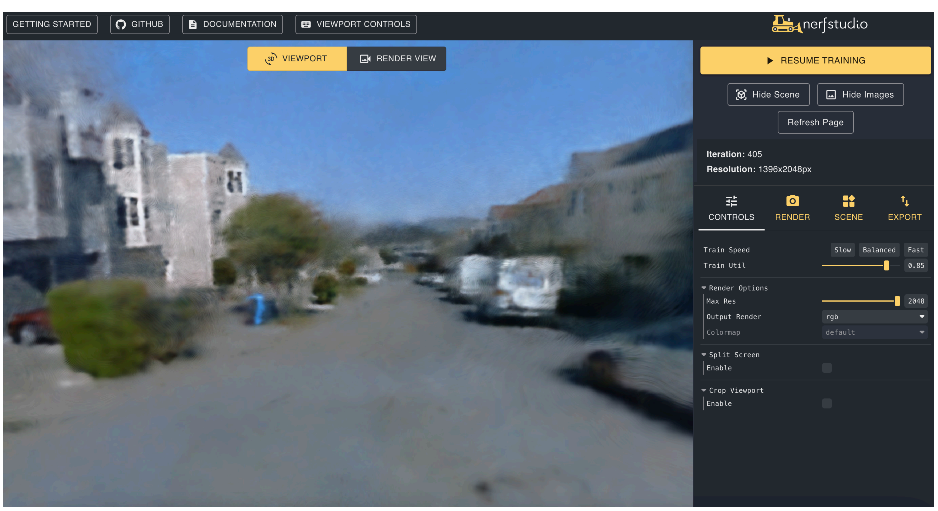

Figure 2. Nerfstudio interface for continual learning of NeRFs. Current view shows training for an initial task from the Waymo open dataset [3].

References

[1] Jaeyoung Chung, Kang-Ho Lee, Sungyong Baik, and Kyoung Mu Lee. Meil-nerf: Memory-efficient incremental learning of neural radiance fields. *ArXiv*, abs/2212.08328, 2022.

[2] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2443, 2017.

[3] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott M. Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2443–2451, 2019.

[4] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, David McAllister, and Angjoo Kanazawa. Nerfstudio: A modular framework for neural radiance field development. *ArXiv*, abs/2302.04264, 2023.