

1 A Implementation Details

2 Our base model F_θ and the expert adaptor E_θ are constructed as 8-layer and 4-layer Multilayer
 3 Perceptrons (MLPs) with 256 channels, respectively, featuring ReLU activations at each layer. For
 4 every image, 1024 rays are sampled, with the network generating 64 coarse and 128 fine points per
 5 ray. During the conflict-aware self-distillation process, an additional 1024 rays are sampled to derive
 6 supervision signals from the teacher model. Our implementation utilizes the PyTorch framework [2],
 7 leveraging the Adam optimizer [1] with an initial learning rate set to $5e-4$, decaying exponentially
 8 down to $5e-5$. The initial network F_θ is trained for 200k iterations on the original scenes from each
 9 dataset. Subsequently, the expert adaptor E_θ undergoes fine-tuning for 10k iterations on each newly
 10 introduced scene, utilizing an initial learning rate of $5e-4$. All experimental procedures are carried
 11 out on a single NVIDIA GeForce RTX 3090 GPU.

12 B More Experiments

13 B.1 Synthetic Dataset Kitchen

14 **Dataset** We conduct a comprehensive evaluation of our algorithm using the *Dataset Kitchen*,
 15 simulating a room-level kitchen environment as depicted in Figure B.1. For the initial training phase,
 16 200-300 multi-view images are captured from six strategically planned camera trajectories to ensure
 17 complete scene coverage and robust training. The dataset is divided into an 80% training subset and a
 18 20% validation subset for performance evaluation. Modifications to the original scene are managed
 19 by capturing approximately 30 multi-view images per change, centered around the area of alteration.
 20 Among them, 10 images are uniformly sampled for training, while the remaining serve for evaluation
 21 purposes. The resolution of all captured images is 960×640 .

22 **Results** We compare two baseline algorithms and the quantitative results are presented in Table
 23 1, Table 2 and Table 3. From the tables, it can be concluded that direct fine-tuning causes drastic
 24 catastrophic forgetting, leading to a significant degradation in the rendering quality of the *old* scene.
 25 Although the naive memory replay mechanism mitigates this (e.g., with the lowest BTM value), it
 26 underperforms in the *new* scene and overall scene (e.g., FM value is not satisfied). Contrastingly, our
 27 proposed approach exhibits outstanding performance in the *new* scene, concurrently preserving the
 28 *old* scenes. The qualitative results, as depicted in Figure B.1, further demonstrate the ability of our
 29 method.

Table 1: Rendering results of atomic/composite/sequential operation in *Dataset Kitchen*. It showcases the PSNR on both new and old scenes following the corresponding operations. The best results are highlighted in **bold**.

Operations	ADD		DELETE		MOVE		REPLACE		SEQ	
	<i>Old</i>	<i>New</i>								
FT	25.48	26.83	23.71	31.06	23.33	23.02	25.13	28.97	23.06	28.58
MR	29.00	25.94	28.11	26.88	28.59	19.97	28.95	29.17	27.63	29.41
Ours	29.32	27.43	29.05	31.09	28.61	23.09	29.33	30.16	28.06	30.07

Table 2: Rendering result of the sequential operation in *Dataset Kitchen*. Tasks 1-4 are defined to evaluate the model based on camera poses corresponding to four operations: add, delete, move, and replace, respectively. Models are evaluated after being tuned for previously completed tasks. The best results are highlighted in **bold**.

Testing on		FT				MR				Ours			
		<i>Task1</i>	<i>Task2</i>	<i>Task3</i>	<i>Task4</i>	<i>Task1</i>	<i>Task2</i>	<i>Task3</i>	<i>Task4</i>	<i>Task1</i>	<i>Task2</i>	<i>Task3</i>	<i>Task4</i>
Training on	<i>Task1</i>	26.83	-	-	-	25.94	-	-	-	27.43	-	-	-
	<i>Task2</i>	23.18	31.06	-	-	23.69	26.88	-	-	26.77	31.09	-	-
	<i>Task3</i>	19.54	25.32	23.02	-	23.71	26.98	19.97	-	26.25	28.13	23.09	-
	<i>Task4</i>	18.36	24.33	18.83	28.97	23.60	26.92	18.83	29.17	25.97	27.92	22.29	30.16

Table 3: Forgetting metrics in sequential operations.

Metrics	FM \uparrow	BTM \downarrow
FT	22.623	6.463
MR	24.630	1.147
Ours	26.585	1.810

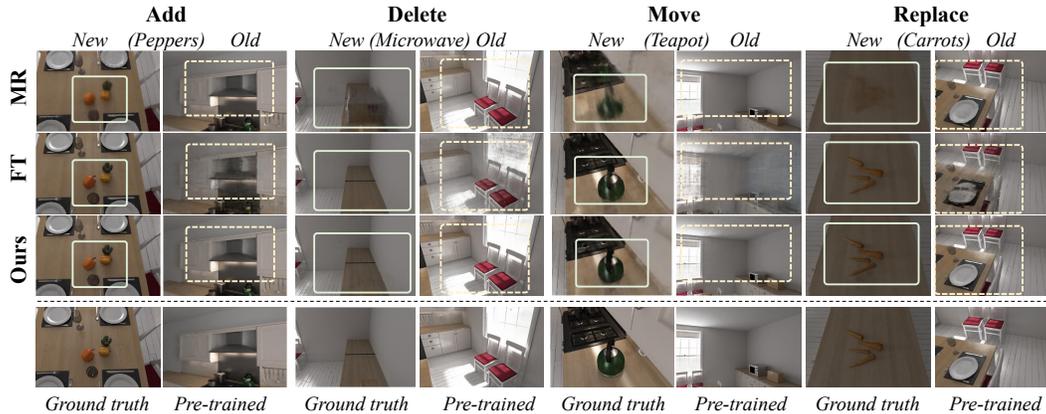


Figure 1: Visualization results after single-step operations in *Dataset Kitchen*. We evaluate 4 single-step operations: add, delete, move, and replace. The ground truth images and inference images from the pre-trained model are used to highlight the rendering quality on new tasks and quality degradation caused by forgetting severity.

30 B.2 City-level Dataset Rome

31 In many cases in the real world, we may only be able to capture a small number of images focused on
 32 the localized area that has been modified. This limitation often results from various factors such as
 33 time constraints, limited access to equipment and resources. For instance, acquiring data at the city
 34 level, whether through aerial or street-level photography, is a demanding and costly task. As a result,
 35 the costs of recapturing data of the whole scene after each scene change are substantial.

36 To illustrate the effectiveness of our model, we also construct a synthetic dataset *Dataset Rome* at the
 37 city level. We collect 200-300 aerial images around the Colosseum in Rome to reconstruct the initial
 38 scene, followed by two ATOMIC operations - ADD and DELETE. The ADD operation adds a black
 39 sports car at the center of the Colosseum, while the DELETE operation removes a tower outside.
 40 The original dataset is partitioned into training and validation sets, with the former accounting for
 41 80% of the data and the latter 20%, to evaluate the algorithm’s efficacy. After scene alterations,
 42 approximately 10 additional images are captured to fine-tune the models, all of which are resized to
 43 640×360 . Visualization results are shown in Figure 3. Our study has demonstrated that in scenarios
 44 where the cost of recapturing data is prohibitively high, our model can proficiently reconstruct new
 45 localized regions with limited amounts of new data and minimal time investment, while significantly
 46 preventing catastrophic forgetting phenomenon.

47 B.3 Real-scene Dataset Laboratory

48 To demonstrate the adaptability of our methods in diverse datasets, we also conduct data collection in a
 49 room-level dataset named *Dataset Laboratory* to establish the practical application of our algorithm in
 50 real-scene settings. We employ an iPhone 13 Pro to capture approximately 120 frames of a laboratory
 51 scene. The resulting images are resized to 640×360 and then divided into training/validation sets
 52 with a ratio of 50%/50% to assess our algorithm’s effectiveness. Following this, we conduct ADD
 53 and DELETE operations on objects within the scene and captured around 10 additional images after
 54 each manipulation to emphasize the modified regions. We compute intrinsic and extrinsic camera
 55 parameters for each frame using an off-the-shelf open-source software COLMAP [3]. The qualitative

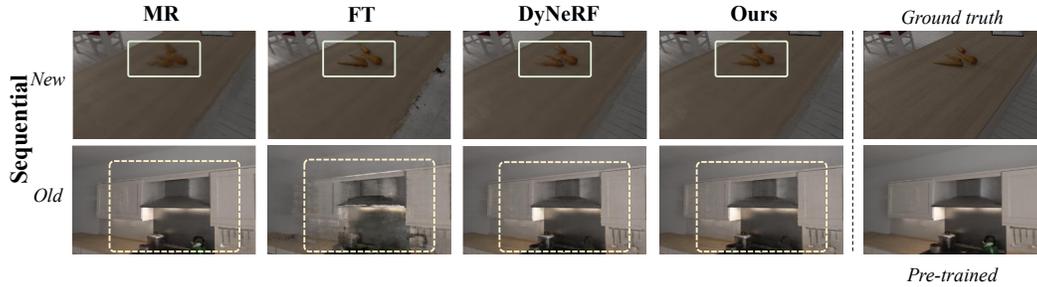


Figure 2: Visualization results after sequential operations in *Dataset Kitchen*. The images in the upper row are rendered after the model has been fine-tuned on four operations, while the images in the lower row test and demonstrate the forgetting issue on the old task..

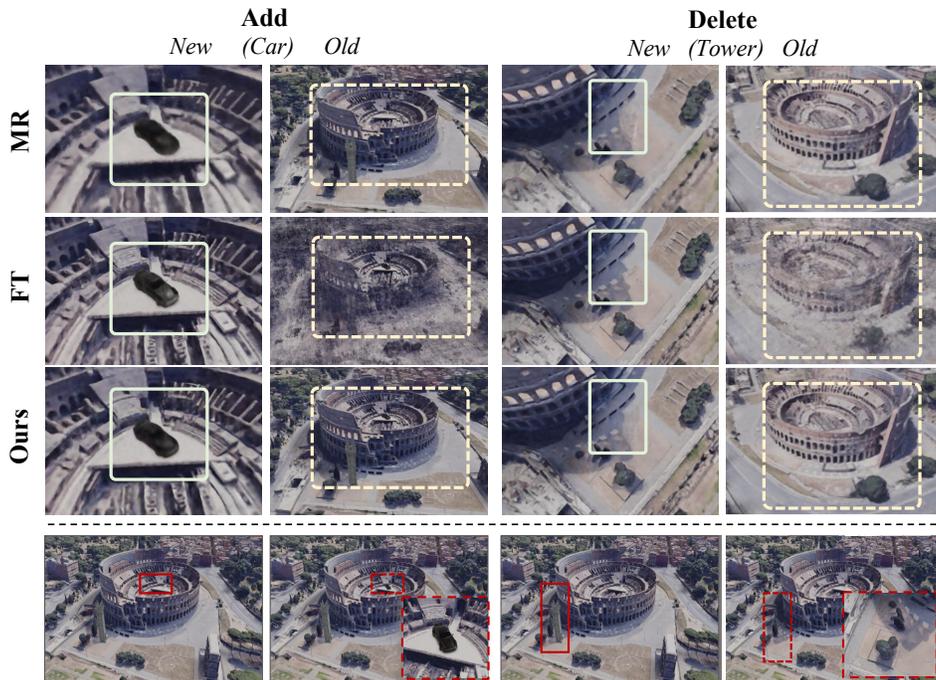


Figure 3: Visualization results after two ATOMIC operations in *Dataset Rome*.

56 results depicted in Figure 4 affirm that our model successfully adapts to scene alterations, such
 57 as introducing a yellow teapot or eliminating a doll, with minimal novel data while maintaining
 58 performance in unmodified regions.

59 B.4 More analysis about DyNeRF

60 It should be noted that in the main paper, training DyNeRF incorporates old images from previous
 61 stages while our method only trains on new images. To ensure a fair comparison, we also conduct
 62 an experiment to train DyNeRF solely with new images in an additional experiment. From Table 4,
 63 not using old images in DyNeRF’s training results in a significant decline in the previous task and
 64 problems of forgetting.



Figure 4: Visualization results in *Dataset Laboratory*.

Table 4: Comparison with DyNeRF without old data in *Dataset Whiteroom* and *Dataset Kitchen*. It showcases the PSNR on both new and old scenes following the corresponding operations. The best results are highlighted in **bold**.

Datasets	Methods	ADD		DELETE		MOVE		REPLACE		SEQ	
		Old	New								
Whiteroom	DyNeRF w/o old	27.40	23.54	27.65	34.68	21.72	30.21	26.31	29.76	22.21	29.06
	Ours	32.33	25.29	32.43	34.77	33.30	29.97	33.33	29.82	33.38	29.74
Kitchen	DyNeRF w/o old	27.33	27.19	27.62	31.28	25.48	22.13	27.69	30.10	26.50	30.00
	Ours	29.32	27.43	29.05	31.09	28.61	23.09	29.33	30.16	28.06	30.07

65 C Limitations

66 We have incorporated an additional expert network into our system. However, the dynamic nature
 67 of the world without constraints may give rise to scalability issues such as storage problems for the
 68 expert network over time. Despite this, we have trained an expert network for each scenario, which
 69 allows us to facilitate the model’s rollback at any given moment, thereby enhancing its practicality.
 70 Additionally, the experts are equipped with the ability to adapt to changes from the original scenarios.
 71 As a result, during model inference, each scenario is still matched with only one expert network,
 72 effectively eliminating scalability concerns during inference.

73 D Boarder Impact

74 This research work exhibits potential for widespread impact across numerous domains, leveraging the
 75 inherent temporal dynamism in real-world scenes. By pioneering an efficient method to adapt NeRFs
 76 to changes in the environment with minimal new data and retaining the memory of unaltered areas,
 77 we increase the accessibility and applicability of NeRFs in various scenarios. CL-NeRF’s ability to
 78 offer high-quality novel views of both changed and unchanged regions could prove instrumental in
 79 fields as diverse as augmented reality, robotics, surveillance, and architecture, where understanding
 80 and adapting to scene changes is paramount. Moreover, by introducing a new benchmark and making
 81 our code available, we invite and facilitate further research and collaboration, driving innovation in
 82 this field. We anticipate that our work, by reducing the retraining burden, will pave the way for more
 83 applications of NeRFs, thus enriching the interaction between artificial intelligence and the physical
 84 world.

85 E More Visualization Results

86 For additional visualization results, please refer to the supplementary video (*supplementary*
 87 *video.mp4*) that we provide.

88 **References**

- 89 [1] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
90 *arXiv:1412.6980*, 2014.
- 91 [2] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen,
92 Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep
93 learning library. *Advances in neural information processing systems*, 32, 2019.
- 94 [3] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the*
95 *IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016.