

# Continuous Pose for Monocular Cameras in Neural Implicit Representation

## Supplementary Material

789 In this document, we provide additional details of our  
790 method and implementations. We further provide quali-  
791 tative examples corresponding to the main paper results.  
792 Please also refer to the supplementary video for additional  
793 qualitative visualizations. We have also attached an exam-  
794 ple code in a separate file.

### 795 6. PoseNet Coordinate Frames

796 In this section, we report the details of the PoseNet out-  
797 puts concerning the coordinate frame under different ap-  
798 plications. For the first two applications involving inac-  
799 curate pose and asynchronous events, we follow the work  
800 [29] and output refined pose from the noisy pose with re-  
801 spect to the  $i$ th camera frame.  $T_{c_i,w} = T_{init_i} \circ T_{refine_i}$ ,  
802  $T_{init_i} = T_{c_i,w} \circ T_{noise_i}$ , while  $T_{a,b}$  represents the rigid-  
803 body transformation matrix that transforms homogeneous  
804 points defined in frame  $b$  to the equivalent points in frame  
805  $a$ . Note  $c_i$  refers to pose estimation of camera  $i$  and  $c_i$   
806 denotes the ground truth. The target of PoseNet is to learn  
807 the cancellation of noise perturbations, essentially to serve as  
808 the inverse of  $T_{noise_i}$ . In the real experiment, we assume  
809 unknown initial pose so  $T_{c_i,w} = I \circ T_{refine_i}$  making the  
810 objective of PoseNet to directly estimate  $T_{c_i,w}$ .

811 In the RGB-D SLAM application in Table 5, we analyze  
812 the impact of varying reference coordinates on tracking. We  
813 denote the PoseNet output with respect to frames  $x$ . So the  
814 estimation of  $i$ th camera pose:  $T_{w,c_i} = T_{w,c_{i-1}} \circ T_{c_{i-1},x} \circ$   
815  $P(f(\theta_p, t_i))$ . The output with respect to different refer-  
816 ence frames is shown in Table 9. Note we get the random frame  
817 by perturbing the pose of  $c_{i-1}$ .

818 For the experiments of IMU, PoseNet outputs the pose of  
819 the agent which is fixed as the IMU sensor. Then we trans-  
820 form the pose to camera frame with  $T_{w,c_i} = T_{w,b_i} \circ T_{b_i,c_i}$   
821 while  $b_i, c_i$  is constant and read from the sensor extrinsic.

### 822 7. NeRF from Inaccurate Pose

823 **Implementation details.** Compared to [29] we make the  
824 following modifications and extensions: (1) BARF perturbs

Reference Frame	Transformation
Default	$T_{c_{i-1},c_i}$
World	$T_{w,c_i}$
Random	$T_{r,c_i}$
Intrinsic	$T_{I,c_i}$
IMU	$T_{b_{i-1},b_i}$

Table 9. PoseNet on different reference coordinates.

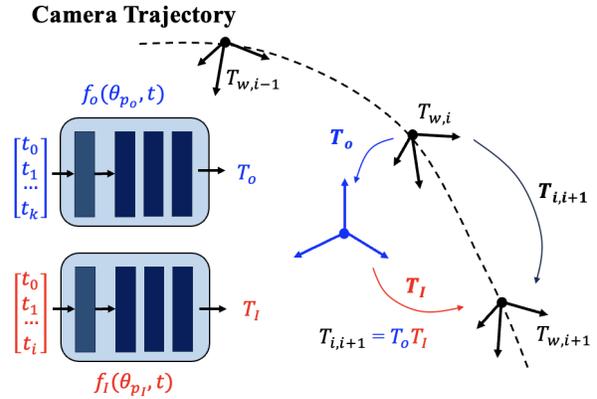


Figure 7. **Intrinsic Motion Frame.** We decompose the relative motion  $T_{i,i+1}$  with a slowly changing rigid transform  $T_o$  and a low dimensional frame-wise motion  $T_I$  using two separate PoseNets.

825 the ground truth pose in synthetic datasets by independently  
826 sampling 6 dimensions Gaussian noise in  $SE(3)$ . We in-  
827 troduce time-dependent noise which is closer to the real-  
828 world scenario, for monocular cameras, where the error of  
829 pose estimation increases with time due to drift and error  
830 accumulation. Furthermore, we also separate the rotation  
831 and translation perturbation and instead of sampling noise  
832 across all frames we only sample a subset of frames and  
833 interpolate the poses for the rest. By doing so we can ex-  
834 plicitly set the maximal deviation on translation or rotation.  
835 (2) Unlike BARF when we optimize one camera pose it also  
836 affects the surrounding poses, therefore a larger batch size  
837 is important for stable training. We use 4096 random rays  
838 for each iteration to optimize camera poses collectively.

839 For joint training with the radiance field, we use the  
840 Adam [23] optimizer for both translation and rotation net-  
841 works with different learning rates. We use a smaller learn-  
842 ing rate for rotation since quaternion rotation expression is  
843 highly nonlinear and difficult to train compared to transla-  
844 tion [28]. We use 1e-3 for TransNet and 2e-4 for RotsNet  
845 and exponentially decay the schedule to 1e-5 and 1e-6 re-  
846 spectively for stable training.

847 **More results on the synthetic dataset.** From Table 11  
848 we can see our method is robust to large translation noise  
849 of up to 40% of the whole scene and is also robust to large  
850 rotation deviations of up to 90 degrees. BARF fails to reg-  
851 ister the camera frame under 20% translation and 60-degree  
852 rotation perturbation and although the 3D object is correctly  
853 reconstructed with largely correct poses, certain novel view  
854 synthesis yields bad PSNR as the object deviates from the  
855 image centre. This can be clearly seen in qualitative results

856 comparison in Figure 8.

857 **More results of real dataset.** More results on other real  
858 scenes as well as qualitative results can be found in Table  
859 16 and Figure 9. Benefiting from neighboring temporal in-  
860 formation our proposed pose representation performs consi-  
861 stently well on different speeds of camera motion. Simi-  
862 lar to the above experiments we can find the novel view  
863 deviates from the image center in Fort/2(19) results. Fur-  
864 thermore, our method is robust to high-speed scenarios with  
865 slight artifacts while BARF diverges and provides very in-  
866 accurate results. Note that in the reported results we disable  
867 the test-time photometric optimization for better compari-  
868 son of camera pose registration performance.

869 **B-spline baseline experiments on 3D.** Similar to 2D pla-  
870 nar experiments we report also the results using classical  
871 continuous B-spline to enforce continuity between neigh-  
872 bouring poses. We experimented with various parameter  
873 configurations to illustrate the challenge of tuning classical  
874 methods in the context of neural radiance fields.

875 **Ablation on network size.** We report the performance  
876 evaluations with different network sizes. The reduction in  
877 network size affects camera localization performance. We  
878 use the 8-layer and 256 width model for other applications.

## 879 8. Continuous Pose for Asynchronous Events

880 **Implementation details.** [47] shows EventNeRF recon-  
881 struction quality cannot handle inaccurate camera poses  
882 over  $1^\circ$ . The real sequences angle offset reported by Event-  
883 NeRF however can reach up to  $2.85^\circ$ . Following its noise  
884 perturbation method, we introduce different magnitudes of  
885 pose inaccuracies in the real datasets. Furthermore, we  
886 also consider the pose inaccuracies due to unknown asyn-  
887 chronous event poses. [47] provides in total 1000 ground  
888 truth poses from Blender which describes a circumferen-  
889 tial movement. We uniformly sample different numbers of  
890 poses to linearly interpolate the whole circular path position  
891 and keep the orientation unchanged. Similarly to above, we  
892 use the Adam optimizer with an exponential learning rate  
893 schedule which decays from  $2e-4$  to  $2e-6$  for TransNet and  
894  $5e-5$  to  $5e-7$  for RotsNet.

Parameter	Rotation error ↓	Translation error ↓	PSNR ↑	SSIM ↑	LPIPS ↓
Regularize = $1e-3$	26.815	14.5	8.87	0.62	0.60
Regularize = $1e-2$	74.586	350.301	9.37	0.71	0.55
Regularize = $1e-1$	115.65	581.81	4.51	0.39	0.73
Regularize = 1	94.81	284.61	9.85	0.70	0.56
knots = 75	50.779	199.2	8.46	0.61	0.60
knots = 50	3.009	9.523	14.46	0.69	0.21
knots = 25	3.01	9.53	14.46	0.69	0.21

Table 10. **Quantitative results of BARF with B-Spline.** We use scipy B-spline interpolation implementation splrep. On top part we use knots = 25 and for bottom part we use s =  $1e-3$ .

TM		80 (10%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	0.06	0.254	27.72	0.92	0.04	
Ours	0.03	0.196	27.91	0.92	0.04	
TM		160 (20%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	24.76	57.342	9.79	0.61	0.52	
Ours	0.05	0.292	26.74	0.91	0.06	
TM		240 (30%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	19.77	95.631	6.97	0.50	0.73	
Ours	0.03	0.178	28.44	0.93	0.04	
TM		320 (40%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	18.66	127.3	7.39	0.53	0.71	
Ours	0.03	0.200	28.25	0.93	0.04	

### (a) Interpolated translational noise experiments.

RM		30°				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	0.067	0.265	27.75	0.92	0.05	
Ours	0.049	0.105	28.22	0.93	0.04	
RM		60°				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	0.101	0.378	26.82	0.91	0.06	
Ours	0.050	0.141	28.13	0.93	0.04	
RM		90°				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	12.103	37.380	10.40	0.61	0.42	
Ours	0.061	0.181	28.03	0.93	0.04	
RM		120°				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	40.526	122.454	6.62	0.54	0.66	
Ours	19.279	66.572	8.79	0.56	0.52	

### (b) Interpolated rotational noise experiments.

TM+RM $R t$		30° + 80(10%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	0.062	0.306	27.78	0.92	0.04	
Ours	0.064	0.266	28.97	0.93	0.04	
TM+RM $R t$		60° + 160(20%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	5.835	29.560	11.63	0.63	0.35	
Ours	0.077	0.293	26.64	0.91	0.06	
TM+RM $R t$		90° + 240(30%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	46.352	160.639	8.17	0.63	0.60	
Ours	0.378	2.813	22.10	0.83	0.09	
TM+RM $R t$		120° + 320(40%)				
Method	RE	TE	PSNR	SSIM	LPIPS	
BaRF[29]	55.640	195.134	7.7	0.63	0.63	
Ours	16.122	52.527	9.16	0.56	0.51	

### (c) Interpolated translational and rotation noise experiments.

Table 11. **Interpolated pose noise experiments.** TM refers to Translational maximal deviation and RM refers to Rotational maximal deviation. The diameter of the circular trajectory is 800, the maximal deviation of the translation perturbation is set to be 10%, 20%, 30%, and 40%.

Method	Rotation error ↓	Translation error ↓	PSNR ↑	SSIM ↑	LPIPS ↓
8-layer, width 256	<b>0.07</b>	<b>0.28</b>	<b>27.30</b>	<b>0.92</b>	<b>0.06</b>
8-layer, width 128	0.09	0.31	27.33	0.90	0.09
4-layer, width 256	0.10	0.32	27.13	0.90	0.10
4-layer, width 128	0.11	0.33	27.15	0.91	0.11

Table 12. **Ablation study on network sizes.** The performance of camera localization drops only slightly with decreased network size. The experiments is conducted using our synthetic dataset, consistent with Table 3.

895 **Qualitative results in interpolation error experiments.**  
 896 In Figure 10 we report the qualitative results of novel view  
 897 synthesis on synthetic sequences of chair and hotdog which  
 898 correspond to Table 4 of the main text. We can see Event-  
 899 NeRF suffers from strong fuzzy artifacts and the depth  
 900 seems to dilate around the object while our method correctly  
 901 learns the depth and reconstructs clearer 3D objects.

902 **Qualitative results on the synthetic datasets in angle off-  
 903 set calibration experiments.** In Figure 11 we report 3  
 904 more real dataset experiments on sequences of multimeter  
 905 and plant. Similar to Figure 5 in the main text, EventNeRF  
 906 suffers from trailing artifacts and at large angle offsets it  
 907 nearly reconstructs 2 separate objects around the image cen-  
 908 ter. In contrast, our method learns the offset angle and re-  
 909 positions the object back to the image center.

## 910 9. Visual SLAM with Depth and IMUs

911 **Full IMU fusion.** In the main text, we elaborate on har-  
 912 nassing gyroscope readings through both loose and tight  
 913 coupling methods. However, direct utilization of ac-  
 914 celerometer readings poses challenges as it provides accel-  
 915 eration instead of velocity in the body frame, resulting in  
 916 a significant error when integrating with an unknown ini-  
 917 tial speed. Additionally, effective processing of accelera-  
 918 tion data necessitates critical steps such as gravity removal  
 919 and denoising [5, 11, 35, 42]. Therefore we first show the  
 920 experiment with simulated IMU on the modified ScanNet  
 921 dataset with simulated IMU as shown in Table 13. Given  
 922 accelerator reading on time  $t$ ,  $\hat{\alpha}_t = (\hat{\alpha}_x, \hat{\alpha}_y, \hat{\alpha}_z)$ . We  
 923 first transform the reading to the last body frame with cap-  
 924 tured image,  $\hat{\alpha}_{t_{i-1},t} = R_{t_{i-1},t} \circ \hat{\alpha}_t$ . We calculate  $R_{t_i,t}$   
 925 from loosely coupled method mentioned above. We then  
 926 use auto-differentiation to calculate the second derivative of  
 927 TransNet with respect to input time and supervise it with L1  
 928 loss:

$$929 \mathcal{L}_{acc} = |\ddot{f}(\theta_p, t) - \hat{\alpha}_{t_{i-1},t}|; \quad (8)$$

930 **Implementation details.** In the NICE-SLAM experi-  
 931 ments, we follow the original work [64] and the bundle ad-  
 932 justment is disabled. Learning rate for TransNet is set to  
 933  $1e-3$  and for RotsNet  $2e-4$ . For IMU experiment we use  
 934

	With IMU				
	scan/059	scan/106	scan/181	scan/207	Average
Nice-SLAM [64]	37.28	174.27	71.94	80.00	89.75
Ours(Gyro)	14.51	12.78	43.98	18.23	22.36
Ours(Acceleration)	14.98	11.49	44.13	19.38	22.49
Ours(Combined)	<b>13.80</b>	<b>10.68</b>	<b>38.20</b>	<b>14.80</b>	<b>19.37</b>

Table 13. **Tracking performance on challenging Scannet [9].** Our PoseNet improves the tracking performance of NICE-SLAM significantly by fusing the IMU tightly. Using full IMU reading yields the best results over all experiment sequences.

Method	v101	v102	v103	v201	v202	v203	Avg
No IMU	2.17	N/A	5.82	7.76	5.04	N/A	N/A
Gyro	<b>1.98</b>	6.09	5.55	<b>4.99</b>	<b>3.03</b>	15.34	6.16
Accelerator	2.16	4.76	5.10	6.72	4.14	15.10	6.33
Combined	2.40	<b>5.33</b>	<b>3.63</b>	5.84	3.46	<b>13.63</b>	<b>5.71</b>

Table 14. **Tracking performance on EUROC [4].** Note that here we report only PoseNet based results. Utilizing both gyroscope and accelerometer data proves beneficial, particularly in challenging scenes, as compared to not using IMU.

Method	GFLOPs	Params[ $10^3$ ]	Time-cost[s/it]
BARF	65.60	514	0.133
Ours	65.62	791	0.138

Method	Tracking time-Cost[ms/iter]	Convergence rate[iter]
NICE-SLAM	27.1	11.96
Ours	31.5	13.21

Table 15. **Left-Computation & Runtime.** Computation of a 1024 batch ray using RTX 3090, with the negligible inclusion of extra computation and time expense. **Right-Runtime & Convergence rate.** We follow the default setting of Replica.yaml. We assume convergence when the tracking loss remains unchanged.

$\lambda_{gyro} = 1$  and  $\lambda_{acc} = 1$ . For IMU simulation we interpolate the ground truth from 20 Hz to 200 Hz and calculate the numerical derivatives. We use the cubic interpolator for translation and *SLERP* for rotation. We downsample the dataset from 20Hz to 5 Hz to highlight the importance of using IMU which is 100 Hz.

**More experiments on RGB-D SLAM with IMU.** As Table 13 shows, by fusing the acceleration and angular velocity we improve NICE-SLAM significantly and can maintain tracking to the end on challenging ScanNet. Taking advantage of both temporal information yields the best tracking performance. Qualitative results can be seen in Figure 12. We then use our method on EUROC [4]. We first use EKF-SLAM to denoise accelerator readings with sensor-fusion from gyroscope and Vicon Pose. As Table 14 demonstrates, fusing accelerator is beneficial especially under challenging scenes such as v103 and v203, and combining both sensor data yields the best results on average.

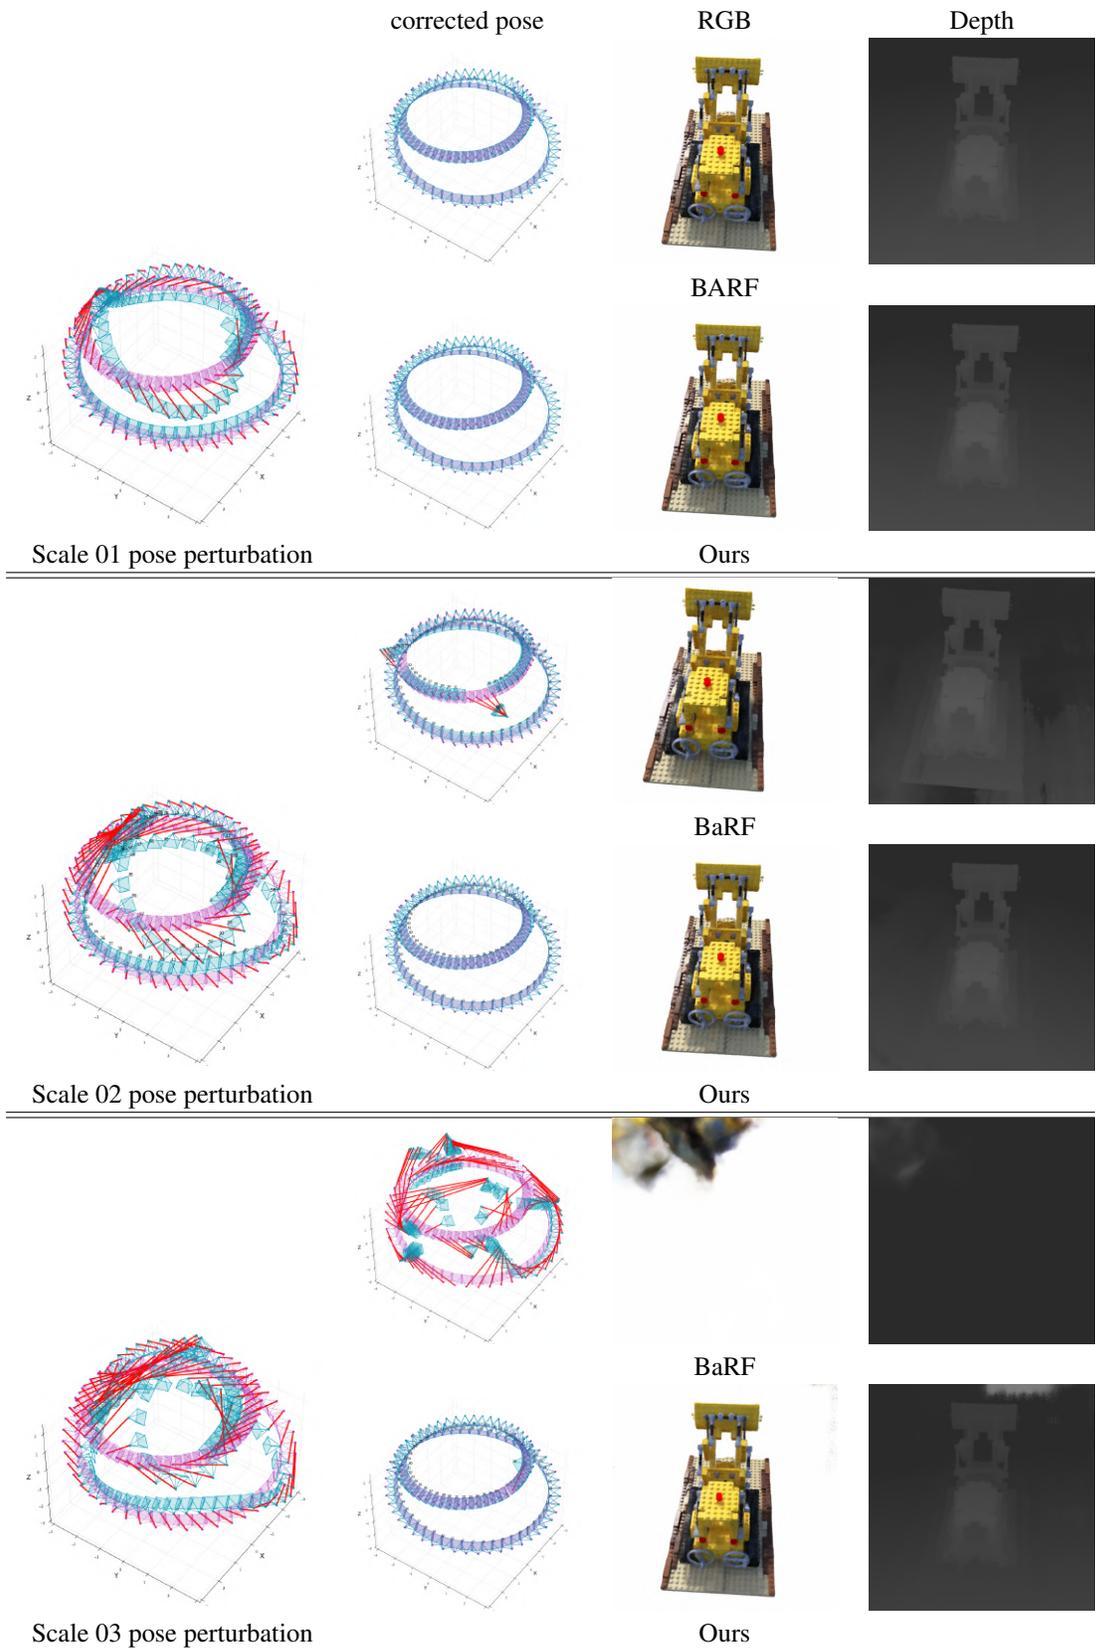
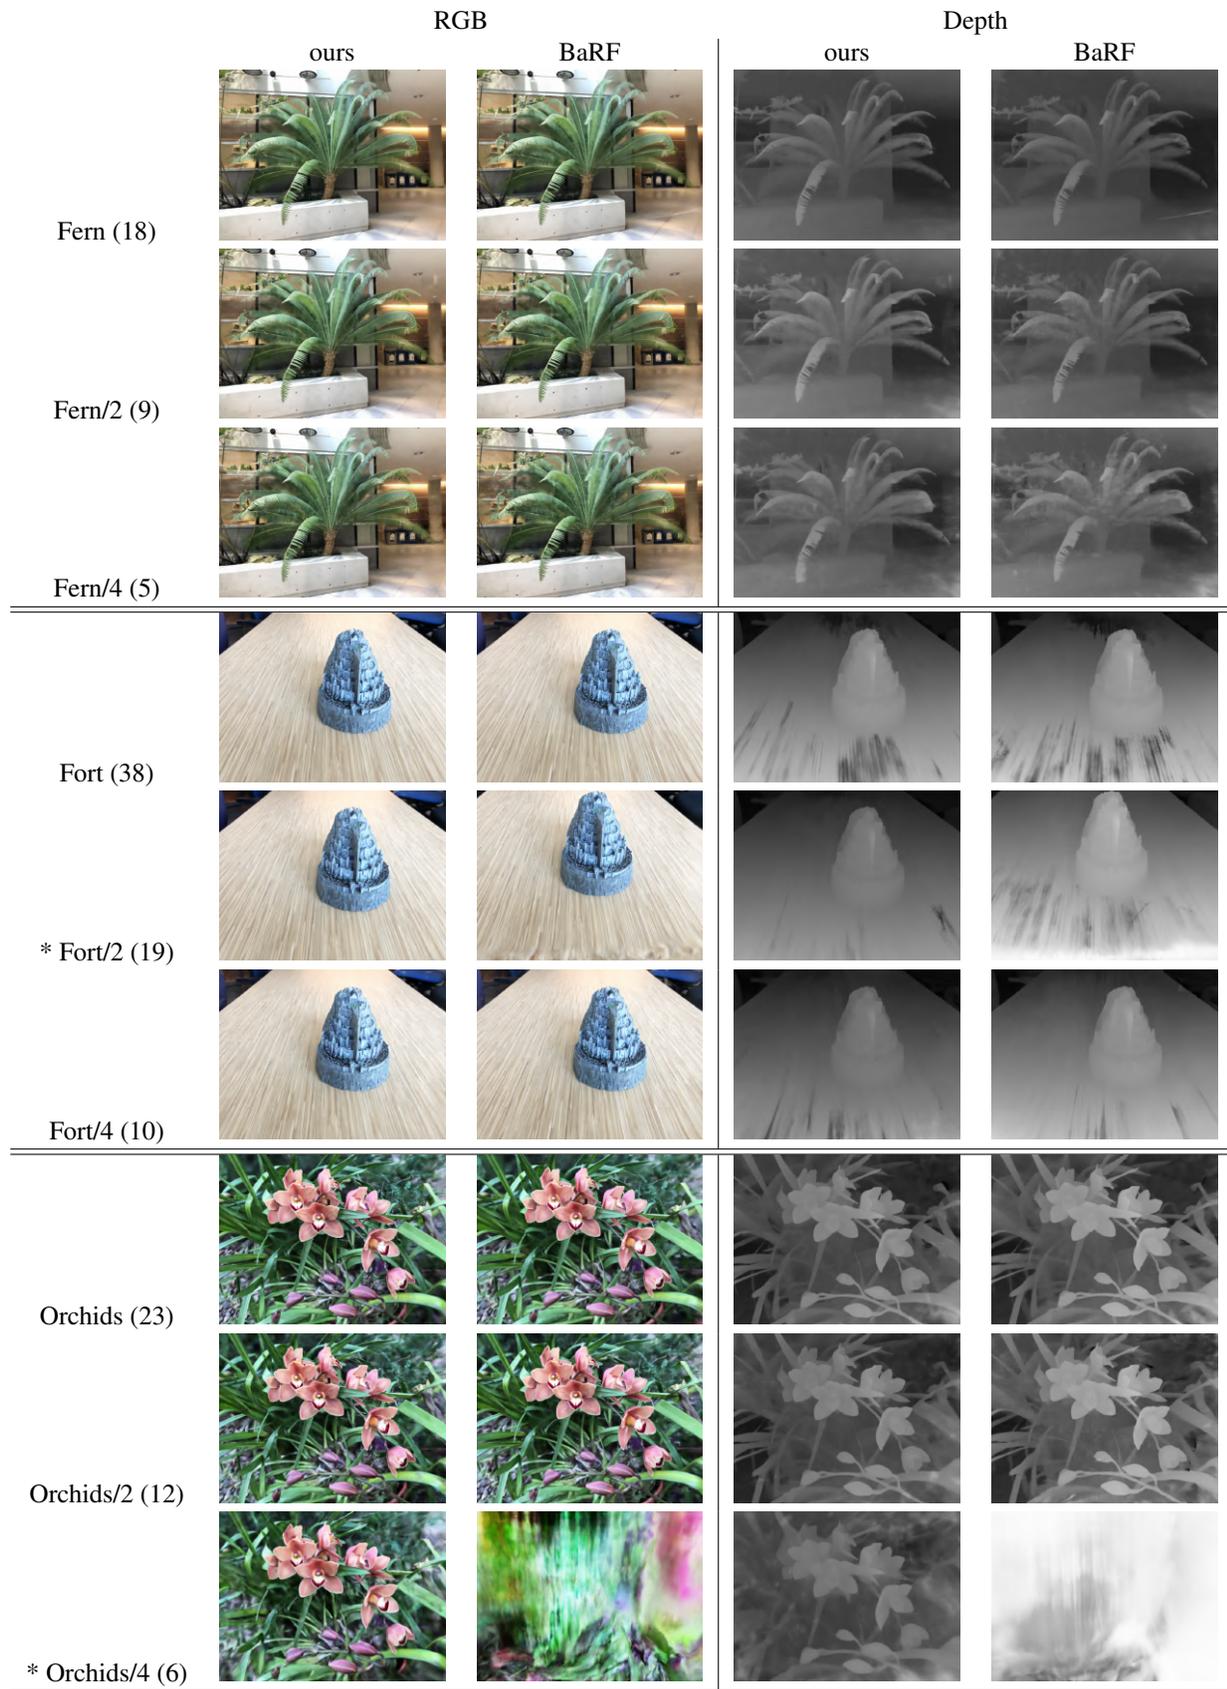
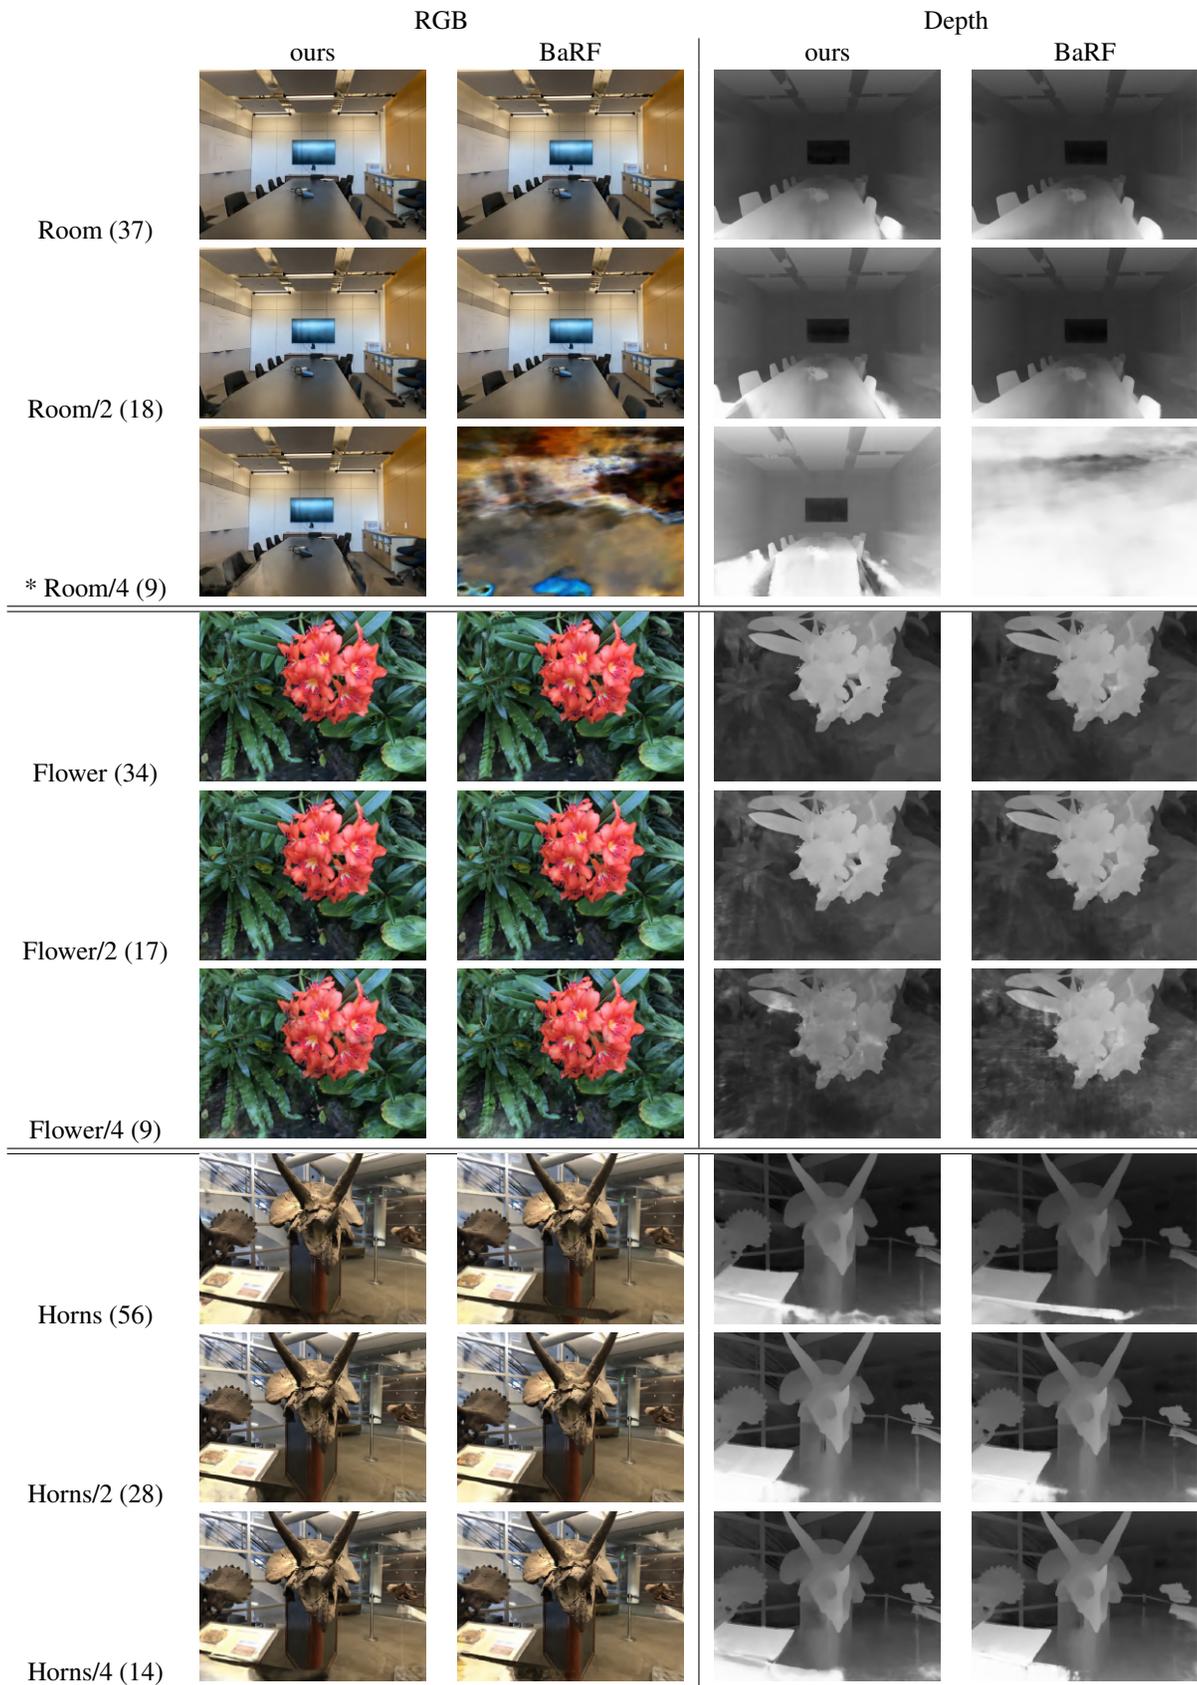


Figure 8. **Qualitative results of interpolated pose noise.** Our method can handle large pose noise and render images in the centre with the correct depths.





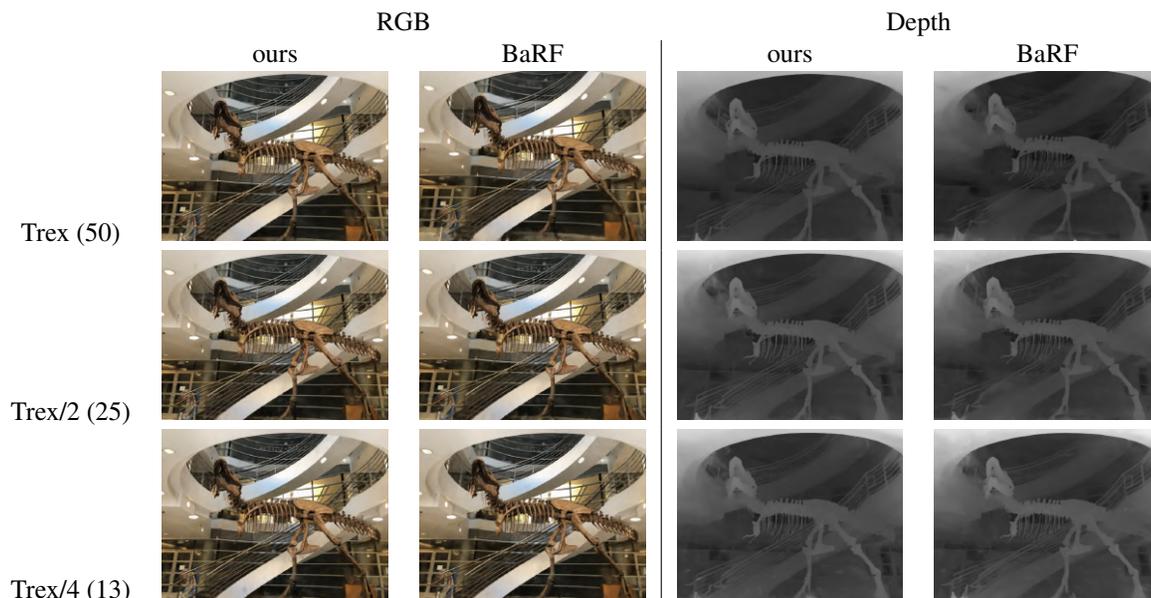
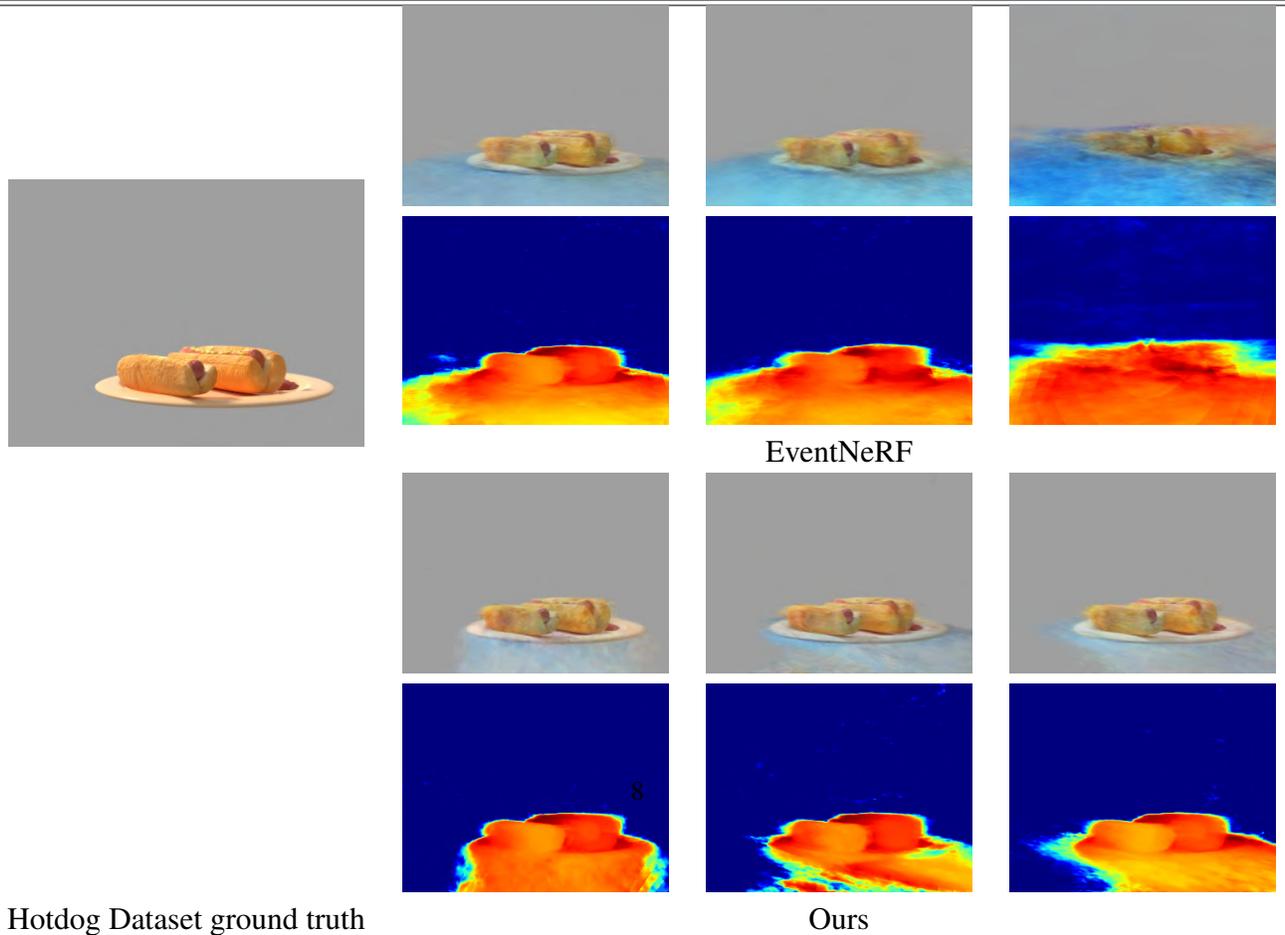
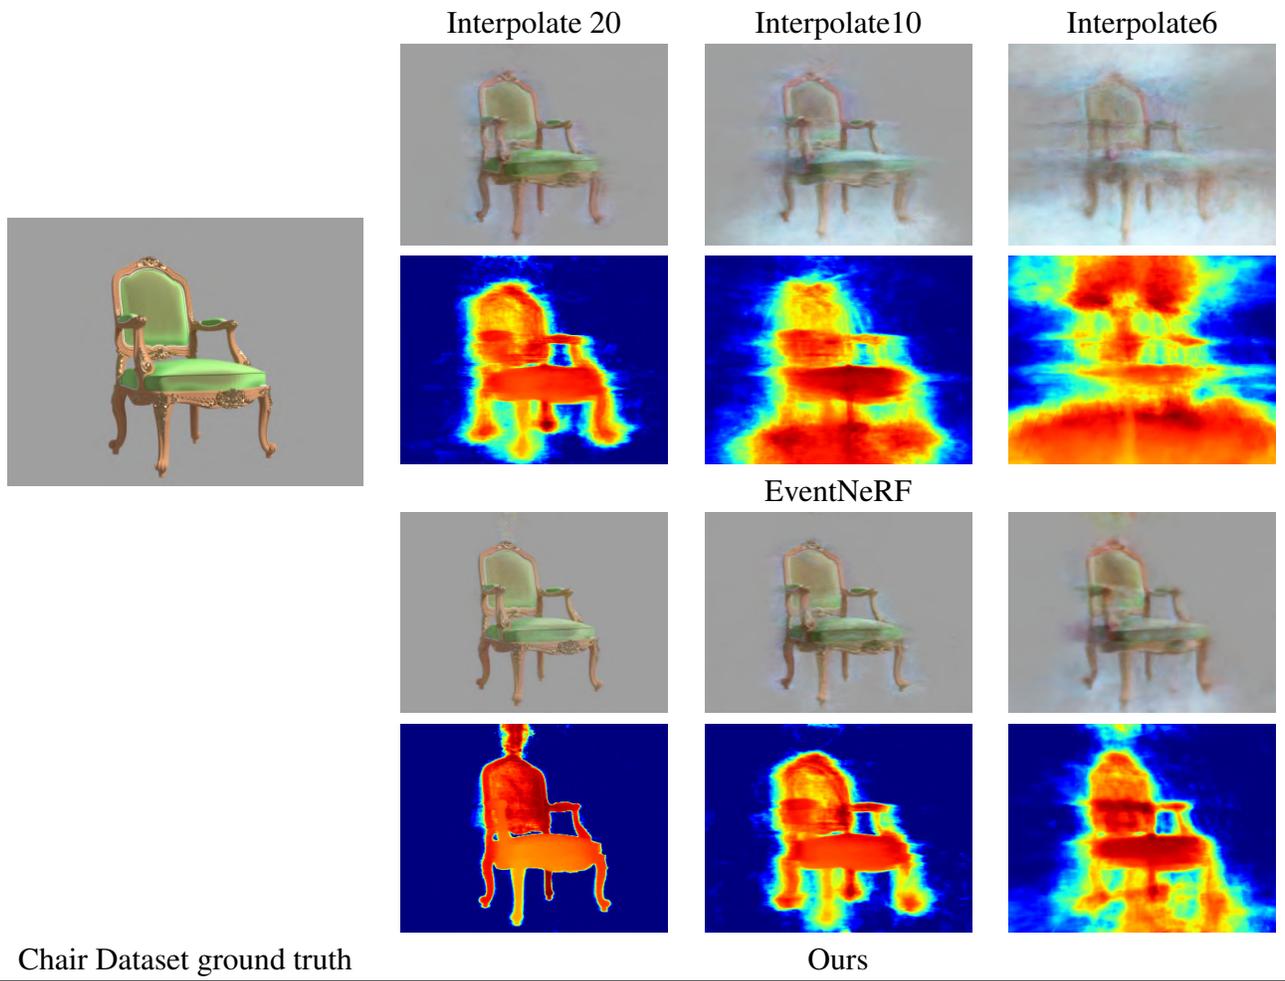
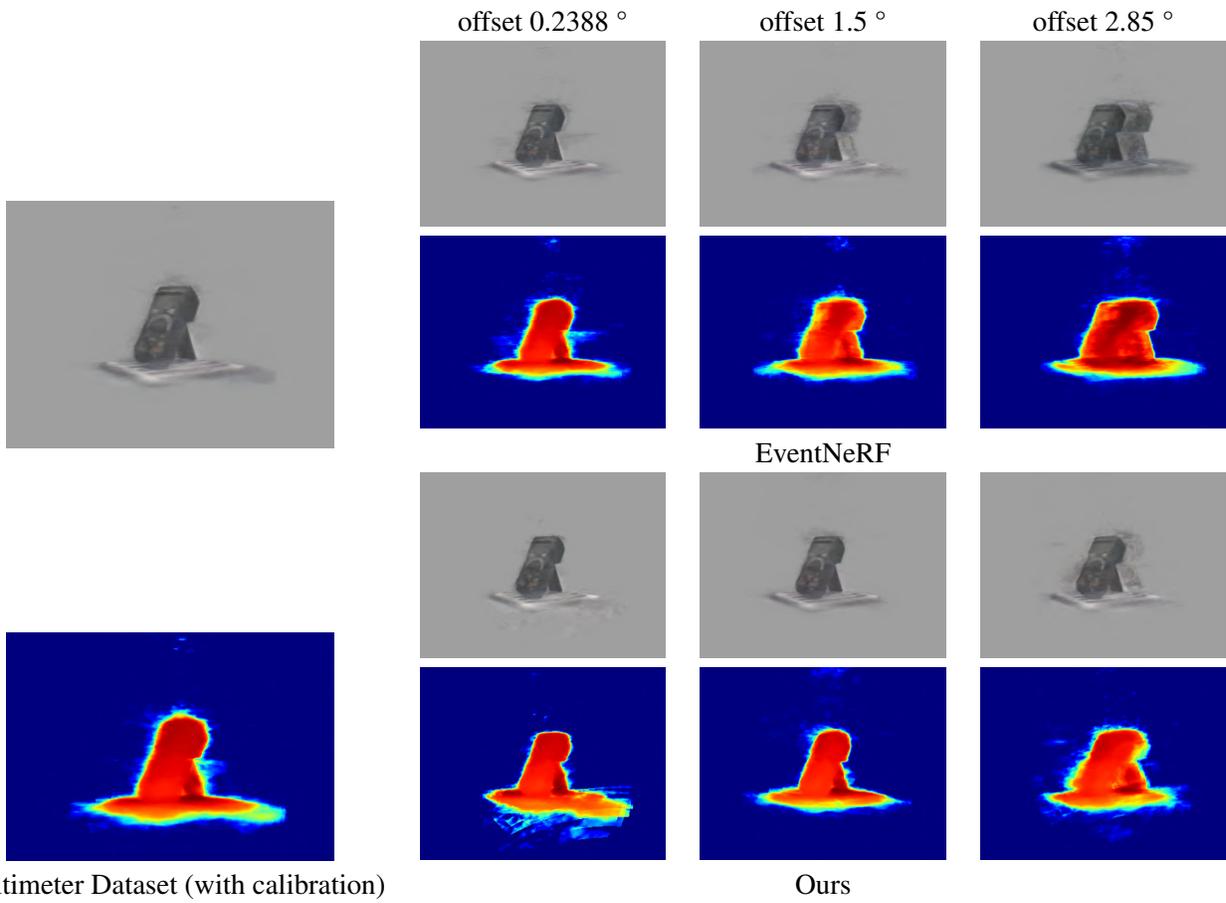


Figure 9. **Qualitative results of novel view synthesis real datasets [32] with unknown pose.** Corresponding to Table 2 in the main text we report the real dataset results in which the camera moves in a smooth trajectory. We denote BaRF failure cases with \* and the number in the parentheses is the number of frames used in training.

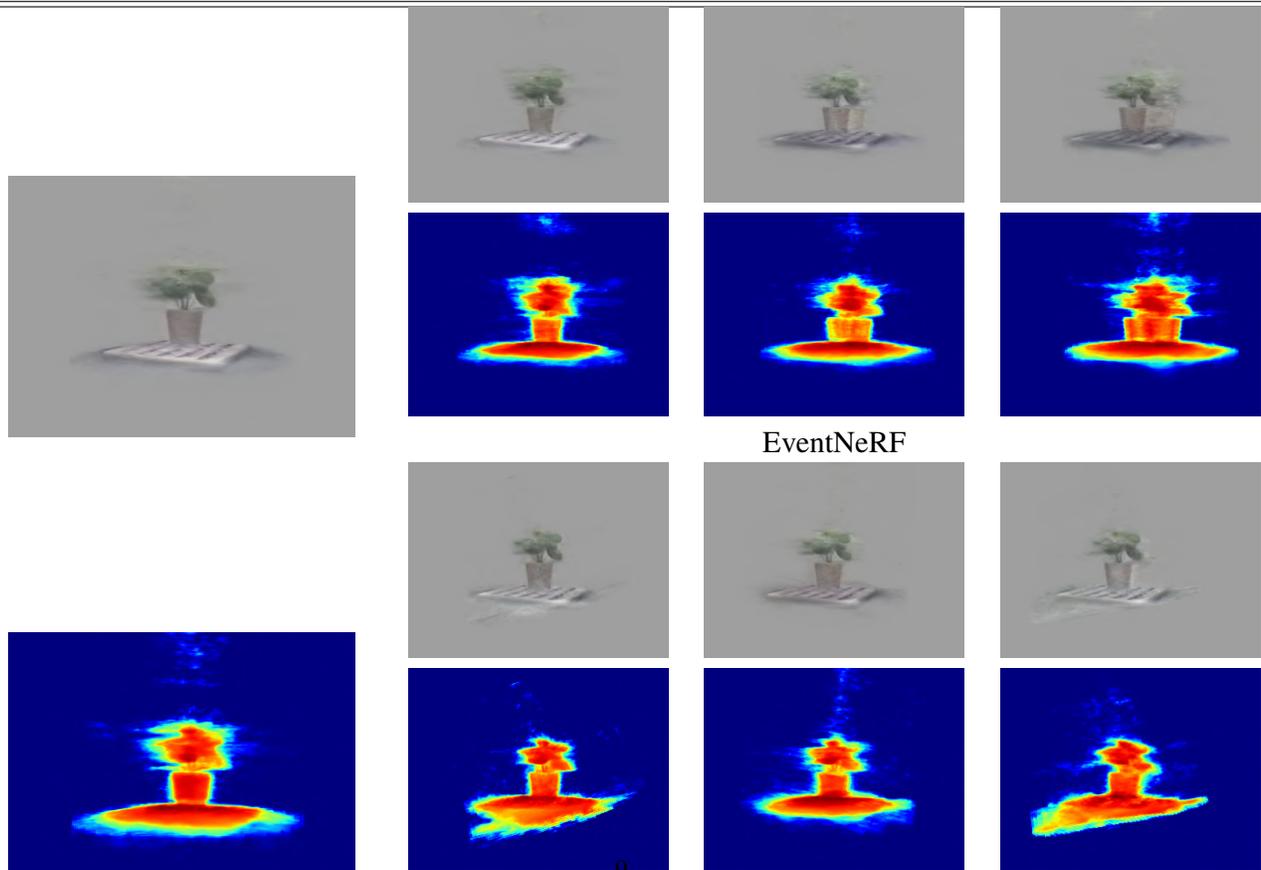
Scene	Rotation ↓		Translation ↓		PSNR ↑		SSIM ↑		LPIPS ↓	
	BARF	ours	BARF	ours	BARF	ours	BARF	ours	BARF	ours
Flower	0.64	0.49	0.27	0.25	17.18	17.93	0.34	0.41	0.27	0.21
Flower/2	0.62	0.46	0.28	0.25	17.18	17.94	0.34	0.36	0.27	0.23
Flower/4	0.59	0.54	0.31	0.29	17.07	17.38	0.33	0.38	0.29	0.27
Horns	0.18	0.19	0.18	0.18	19.58	18.89	0.59	0.55	0.32	0.27
Horns/2	0.27	0.33	0.20	0.17	16.24	16.09	0.49	0.45	0.31	0.28
Horns/4	0.21	0.24	0.16	0.17	16.91	16.85	0.54	0.53	0.32	0.32
Trex	0.49	0.41	0.38	0.35	16.53	17.04	0.42	0.45	0.21	0.19
Trex/2	0.56	0.26	0.43	0.29	16.37	18.96	0.40	0.61	0.23	0.16
Trex/4	0.19	0.20	0.24	0.26	21.62	20.74	0.73	0.70	0.17	0.15
Average	0.42	<b>0.34</b>	0.27	<b>0.24</b>	17.63	<b>17.95</b>	0.46	<b>0.49</b>	0.27	<b>0.23</b>

Table 16. **More quantitative results on real datasets.** In addition to Table 2 we report more results on LLFF [32] dataset. Note that in this dataset images are captured in a top-down, left-right manner rather than following a continuous trajectory. Consequently, our method may not be fully leveraged. Nevertheless, when considering average values, our approach outperforms the baseline.





Multimeter Dataset (with calibration)



Plant Dataset (with calibration)

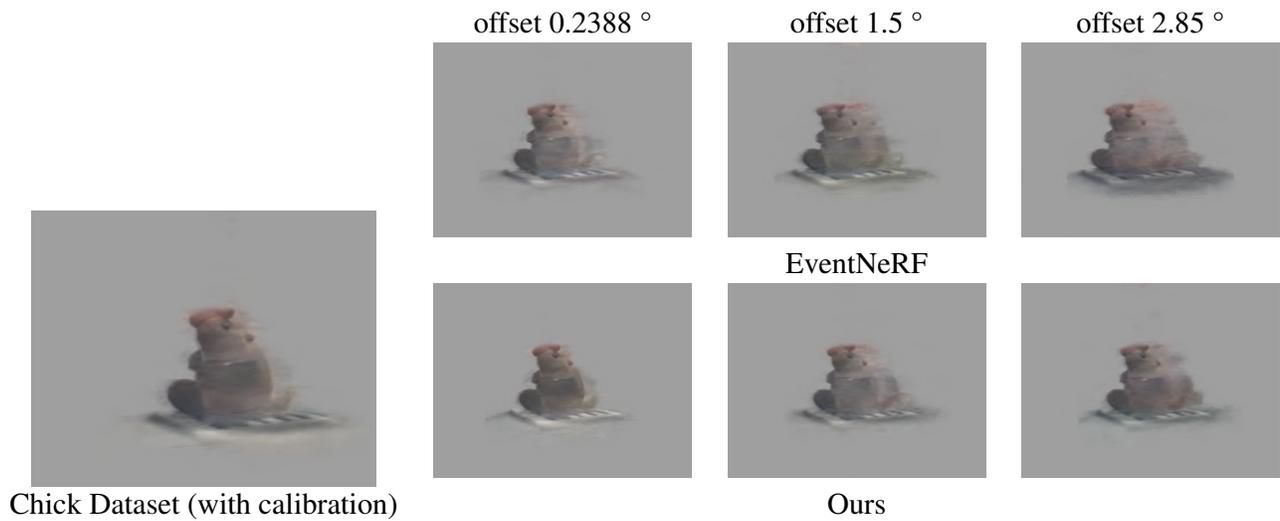


Figure 11. **Qualitative results of novel view depth and rgb synthesis in angle offset calibration experiments.** Our method improves EventNeRF significantly in all six experimental setups.

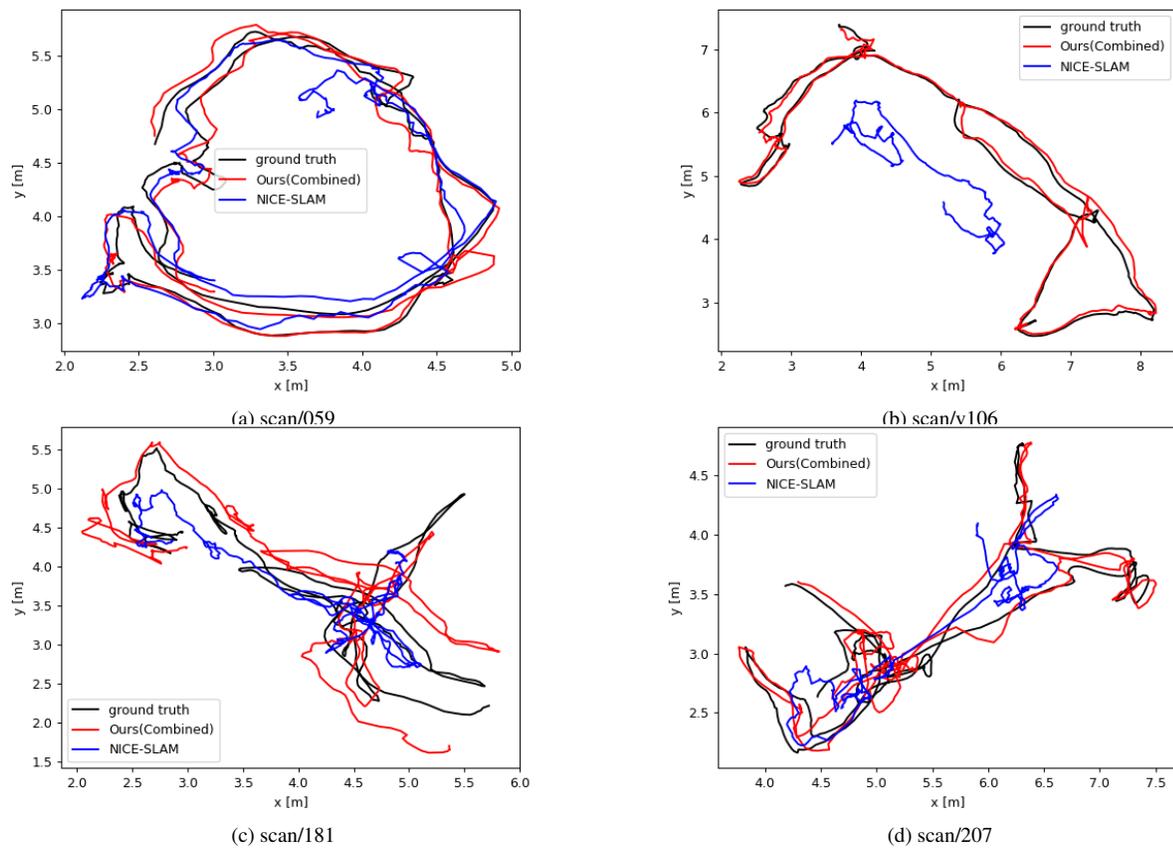
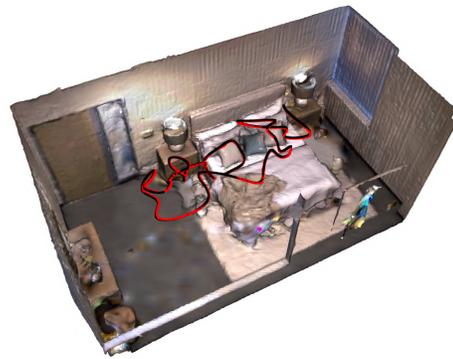


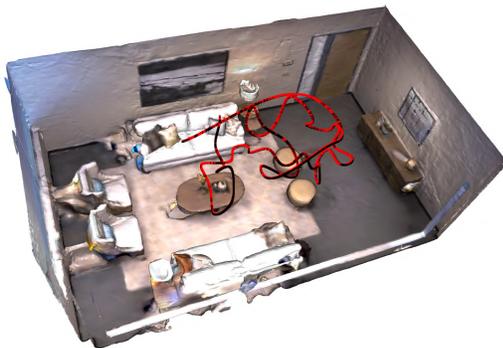
Figure 12. **Qualitative results of tracking on challenging ScanNet** With the assistance of simulated IMU information, our method maintains robust tracking and preserves scale accuracy.



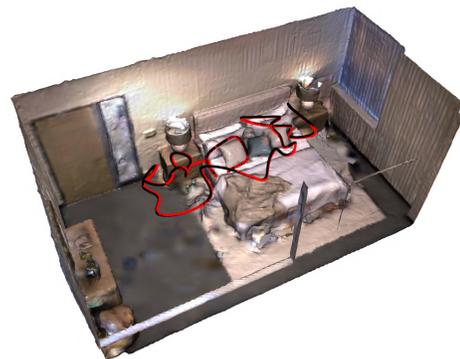
(a) NICE-SLAM



(a) NICE-SLAM



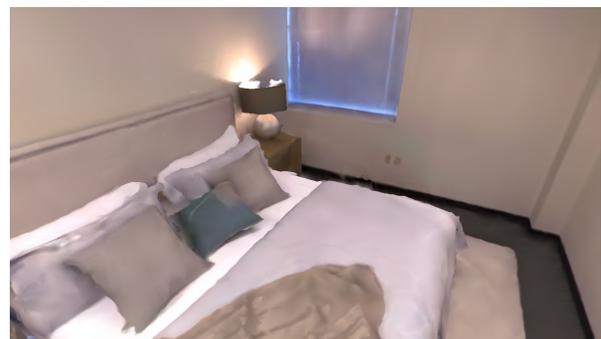
(b) Ours



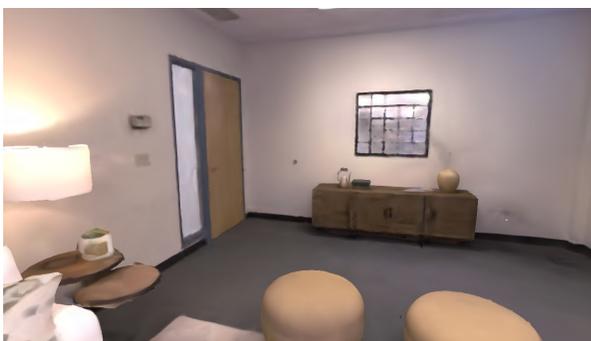
(b) Ours



(c) NICE-SLAM PSNR: 33.9



(c) NICE-SLAM PSNR: 32.7



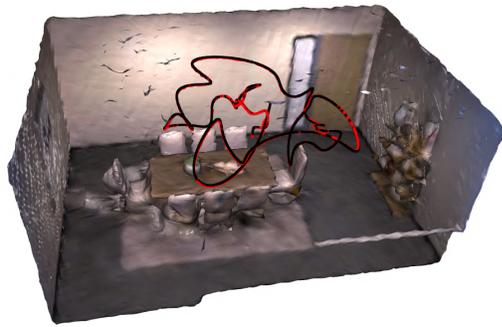
(d) Ours PSNR: 36.9



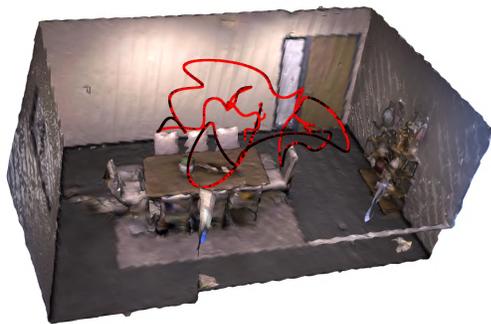
(d) Ours PSNR: 33.3

**Figure 13. Reconstruction and Rendering of Replica Room0.** Thanks to the improvement of tracking performance, our method is able to substantially increase the fidelity of the renderings. This is also supported by the quantitative results PSNR. We reconstruct clean details compared to NICE-SLAM.

**Figure 14. Reconstruction and Rendering of Replica Room1.** In this relatively easier scene, we perform slightly better than NICE-SLAM in rendering and reconstruction with less artifacts.



(a) NICE-SLAM



(b) Ours

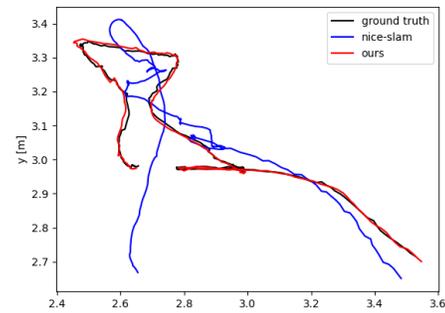


(c) NICE-SLAM PSNR: 33.3

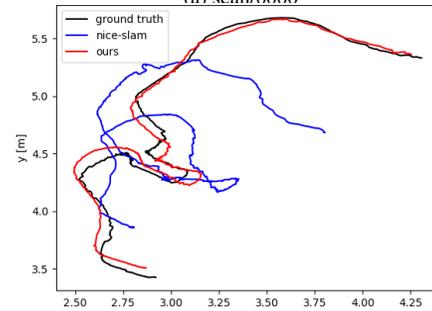


(d) Ours PSNR: 36.8

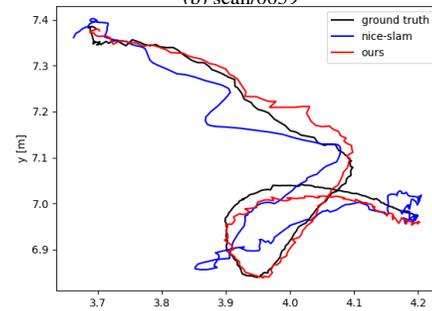
Figure 15. **Reconstruction and Rendering of Replica Room2.** While the reconstruction demonstrates that the NICE-SLAM trajectory is highly aligned with the ground truth, it adversely affects rendering performance, resulting in lower fidelity. In contrast, our method maintains high-fidelity rendering.



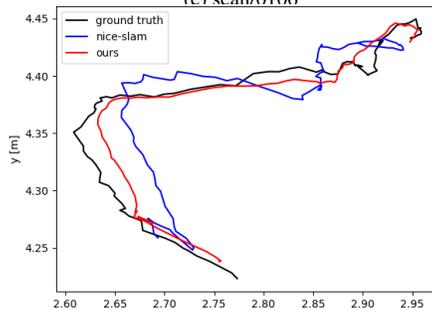
(a) scan/0000



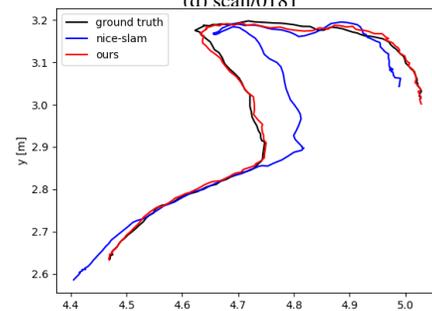
(b) scan/0059



(c) scan/0106



(d) scan/0181



(e) scan/0207

Figure 16. **Qualitative results of tracking on ScanNet[9].**The initial trajectories diverge in the NICE-SLAM trajectory from the ground truth, while ours align with it.

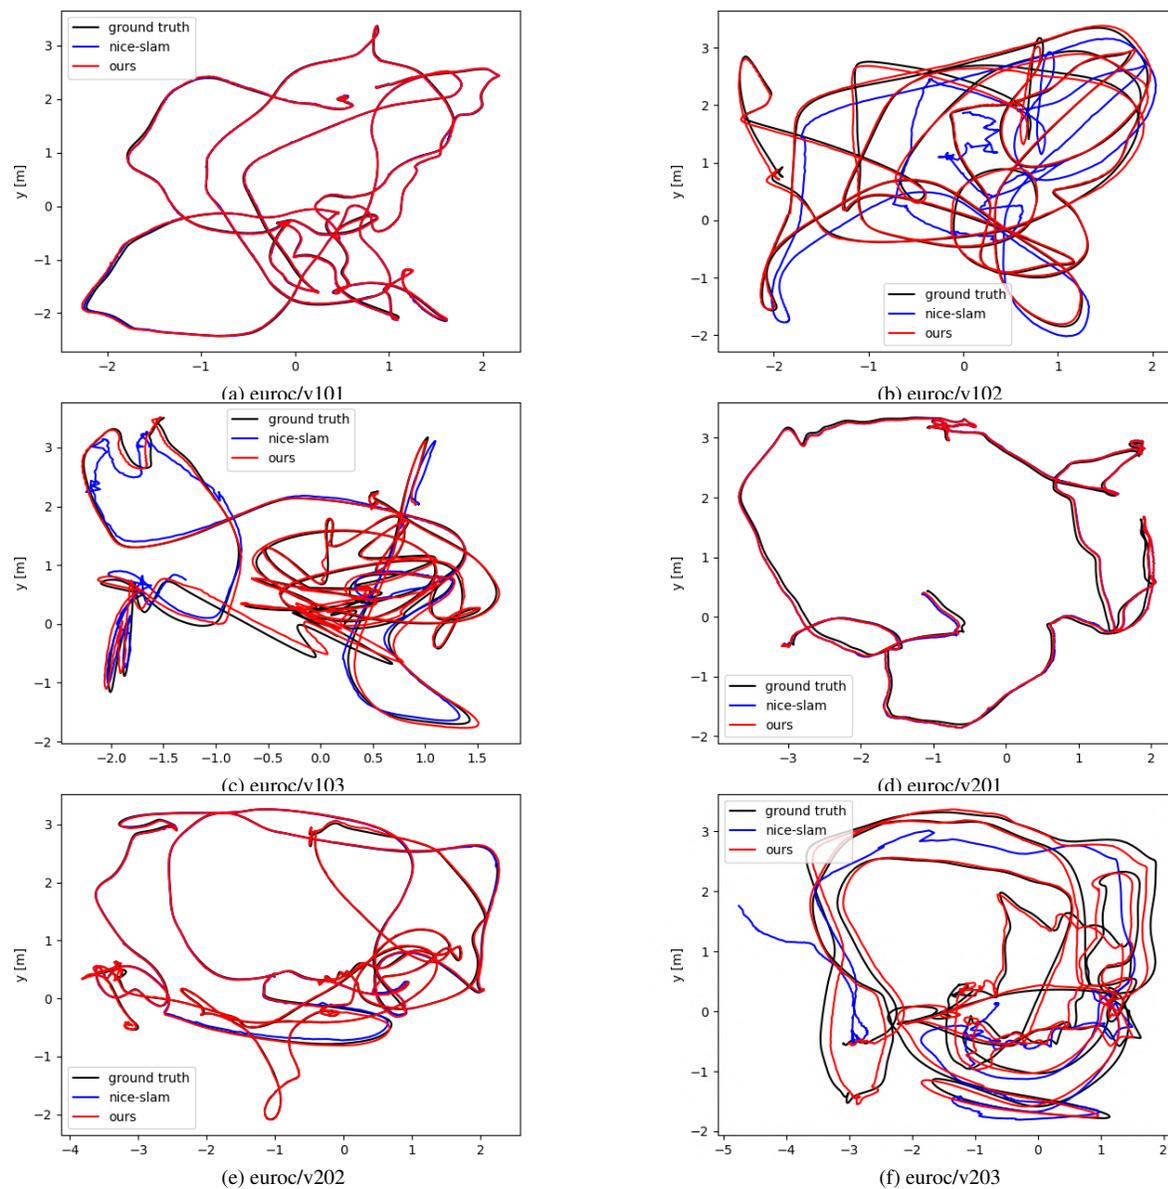


Figure 17. **Qualitative results of tracking on EUROCC4**. We compare the trajectories of our method to NICE-SLAM. Notably, NICE-SLAM encounters failures at v102, v202, and v203, so only part of trajectories are displayed. The results indicate that our method significantly aligns with the ground truth trajectory.