Supplementary Material
NeuRoRA: Neural Robust Rotation Averaging

Pulak Purkait[0000−0003−0684−1209], Tat-Jun Chin[0000−0003−2423−9342], and Ian Reid[0000−0001−7790−6423]

The University of Adelaide, Adelaide SA 5005, Australia
https://github.com/pulak09/NeuRoRA

Table of Contents

List of Figures

1 Results on ScanNet datasets

The dataset [3] consists of 2.5M images from 1.5K indoor scenes. The images are collected by capturing videos with a tablet. An RGBD SLAM system (BundleFusion [4]) is then applied to get the ground-truth camera poses (along with the absolute orientations). The dataset however does not provide the relative poses of the images. We employed a publicly available toolbox VisualSFM [7] to obtain the relative orientations of consecutive images. Note that only 50 neighboring images for each image are used for the comparison. For each sequence only the largest connected component was considered to form our dataset. Note that the pair-wise comparison is slow, and so far we have collected 480 sequences (out of 1500) of which 40% were used for training, 10% for validation and remaining 50% were used for testing. The average number of #cameras is 1401 and average number of edges is #edges is 5.8%.

The results are displayed in Figure 1. Note that these sequences are very tough sequences as these are very sparsely connected forming chains. The proposed CleanNet does a very good job cleaning the network and FineNet fine-tunes it. The combined network NeuRoRa produces best results compared to the baselines. Although the performance gap w.r.t. Chatterjee [2] is small, NeuRoRa is much faster. We will release the entire dataset (once built) public to foster research in this direction.

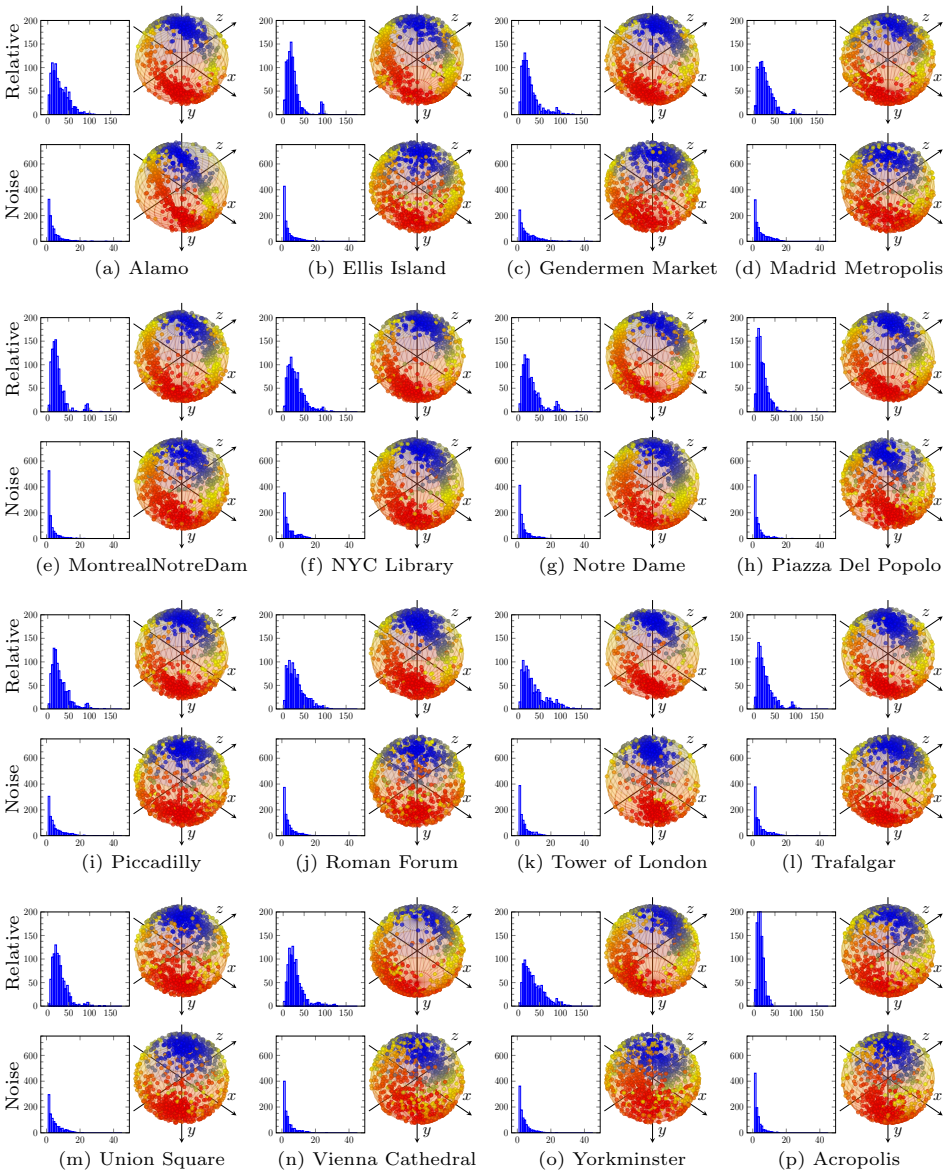

Fig. 1. More examples of Figure 1(a)-(c) are displayed. The angle and axes of sampled observed relative orientations (first row) and the same of noise (second row) in real datasets (for clarity only 1000 random samples) are displayed. The noise orientation is calculated from the ground-truth absolute orientations and the observed relative orientations. We plotted histograms of the magnitudes of the angles in degrees and the axes of the orientations. Notice that the axes of the sampled relative and noise orientations are distributed mostly along a vertical ring rather than uniformly on a unit ball.

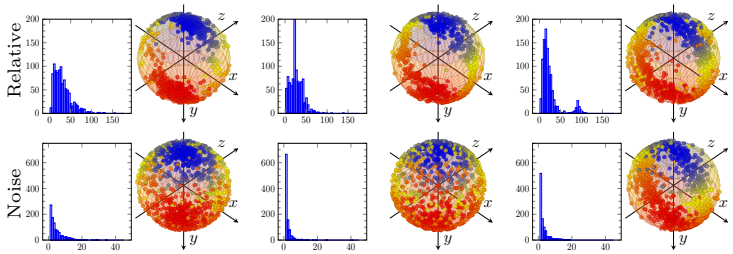

(q) Arts Quad (r) San Francisco (s) TNotre Dame

Fig. 2. Same as Figure 1 for the other view-graphs

Table 1. Results of Rotation averaging on ScanNet dataset [3]. The average angular error on all the view-graphs in our dataset is displayed. The proposed method NeuRoRA is remarkably faster than the baselines while producing better results. Note that for these sequences, NeuRoRA takes only **0.028*s*** on average on a GPU.

Methods	mn	md	cpu		mn	md	cpu
Chatterjee [2]	12.08°	7.80°	2.01s	Arrigoni [1]	32.44°	18.26°	30.20*s*
Weiszfeld [5]	22.93°	12.30°	82.92*s*	Wang [6]	18.68°	8.53°	6.52s
CleanNet-SPT	11.37°	7.41°	**0.45*s***	NeuRoRA	**11.02°**	**6.92°**	**0.92*s***

mn: mean of the angular error; md: median of the angular error;
cpu: the average runtime of the method;

2 The angle and axes of observed relative orientations

More examples of samples of Figure 2 in the main paper are plotted in Figure 1.

References

1. Federica Arrigoni, Beatrice Rossi, Pasqualina Fragneto, and Andrea Fusiello. Robust synchronization in so(3) and se(3) via low-rank and sparse matrix decomposition. *CVIU*, 174:95–113, 2018. 3
2. Avishek Chatterjee and Venu Madhav Govindu. Robust relative rotation averaging. *TPAMI*, 40(4):958–972, 2017. 1, 3
3. Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2017. 1, 3
4. Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (ToG)*, 36(4):1, 2017. 1
5. Richard Hartley, Khurrum Aftab, and Jochen Trumpf. L1 rotation averaging using the weiszfeld algorithm. In *Proceedings of CVPR*, pages 3041–3048, 2011. 3

6. Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2(2):145–193, 2013. 3
7. Changchang Wu et al. Visualsfm: A visual structure from motion system. 2011. 1