Supplementary Material: Gravity-aligned Rotation Averaging with Circular Regression

Linfei Pan[1], Marc Pollefeys[1,2], Dániel Baráth[1]
[1]ETH Zurich, [2]Microsoft
{linfei.pan, marc.pollefeys, danielbela.barath}@inf.ethz.ch

This supplementary material provides more results on the LaMAR datasets, an alternative solution to the problem based on changing the standard 3-DoF optimization, analysis of the period estimation, comparisons with deep learning-based alternatives, results with gravity-aligned essential matrix, implementation details for Theia baselines, and detailed experiment results.

A Additional Results on LaMAR

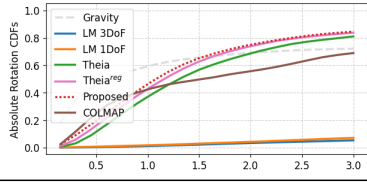

	AVG (°)	MED (°)	AUC@0.5°	@1°	@2°	Time (h)
LM 3DoF	52.30	32.36	0.07	0.35	1.26	4.08 $\times 10^{-2}$
LM 1DoF	44.90	31.61	0.28	0.73	1.81	1.69 $\times 10^{-2}$
Theia [12]	**3.46**	1.28	3.28	13.92	34.62	3.88 $\times 10^{-2}$
Theia^{reg} [4]	8.14	1.13	5.36	17.65	39.42	4.31 $\times 10^{-2}$
Proposed	7.19	**1.06**	7.63	20.46	**42.02**	**1.21** $\times 10^{-2}$
COLMAP [11]	34.18	1.42	**9.63**	**22.59**	36.09	98.52

Fig. 1: Results for on LaMAR [10]. Only images taken by Hololens data are with gravity, thus [2, 5] are not applicable.

In the main paper, we consider gravity directions only in the data captured by smartphones to test the mixed scenario. Here, we provide additional results for when HoloLens directions are used instead of the ones captured by smartphones. Moreover, we show results when both data sources have known gravity directions.

The results considering gravity only in the HoloLens sequences are in Fig. 1. Planar pose graph optimization methods [2, 5] are excluded since they are not applicable in the mixed scenario. A similar trend can be observed as in the case when only gravity captured by smartphones is known (Table 4 in the main paper). The proposed method achieves higher AUC scores than Theia-based ones and is about 3-4 times faster. While the proposed method has a higher average error than the algorithm implemented in Theia [12], it is important to note that the average error is usually not representative in such problems since it is

not a robust measure. In *all* other metrics, the proposed method significantly outperforms Theia. The proposed method is on par with COLMAP [11], having slightly lower AUC scores at $0.5°, 1°$ and a higher AUC score at $2°$. The proposed method also achieves lower mean and median errors and is faster by several orders of magnitude. The cumulative distribution functions (CDF) of the absolute rotation errors are also shown in Fig. 1. The proposed method approaches the top left corner most, indicating its accuracy.

Table 1: Average and median rotation errors (°), the AUC score at 0.5°, 1°, and 2° and average run-time (hour) on LaMAR [10]. Images without gravity are excluded.

	AVG (°)	MED (°)	AUC@0.5°	@1°	@2°	Time (h)
LM 3DoF	52.31	32.20	0.08	0.36	1.25	4.05×10^{-2}
LM 1DoF	46.25	35.76	0.26	0.78	1.98	$\mathbf{1.28} \times 10^{-3}$
LAGO [2]	51.28	37.17	0.33	0.94	2.22	5.18×10^{-3}
CPL-Sync [5]	11.57	7.95	0.92	2.63	6.27	3.09×10^{-3}
Theia [12]	3.52	1.28	3.26	13.82	34.36	4.03×10^{-2}
Theia$^{\text{reg}}$ [4]	4.26	1.04	7.12	20.91	41.99	3.79×10^{-2}
Proposed	**2.90**	0.94	9.32	23.85	**46.24**	3.88×10^{-3}
COLMAP [11]	9.68	**0.88**	**12.71**	**28.38**	44.22	98.52

The results for the case when gravity is available for both types of images are in Table 1. We exclude approximately 150 images in total from this experiment for which the dataset does not provide gravity directions, and we keep only the largest connected component. This is necessary to test [2,5]. Despite our tuning efforts, LAGO [2] and CPL-Sync [5] do not achieve reasonable accuracy on this dataset. The proposed method achieves higher AUC scores at 0.5°, 1°, and 2° compared with Theia-based methods and is about an order of magnitude faster.

The cumulative distribution functions (CDF) of the absolute rotation errors can be found in Figure 2c. The curve of the proposed method (red) is closest to the top left, implying that it obtains the most accurate estimation.

B Alternative Representation

As mentioned in the main text, an alternative way to achieve 1DoF optimization is to constrain the update step in the numerical optimization so that it only changes the rotation around the y-axis. More precisely, in each optimization step, the updates on rotations are projected to be around the y-axis. Though minimizing the same loss function as the proposed method, the alternative formulation requires solving a larger linear system. As a result, it achieves similar accuracy at the cost of a significantly longer runtime compared to the proposed method. Furthermore, it is not applicable when gravity is partially known. A detailed analysis of the relationship between the proposed method and such alternative formulation can be found in Section C. We will call this alternative formulation Proposed$^{\text{alt}}$.

The cumulative distribution functions (CDF) of the absolute rotation errors are shown in Fig. 2. As expected, the curves of the alternative formulation and

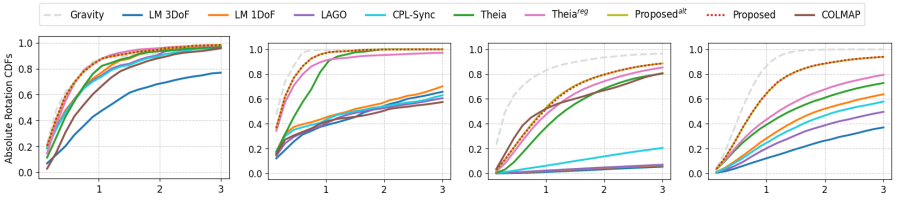

(a) EuRoC [1] **(b)** KITTI [6] **(c)** LaMAR [10] **(d)** 1DSfM [14]

Fig. 2: The cumulative distribution functions (CDFs) of the absolute rotation errors (°). Estimated by the Levenberg-Marquardt [8] method solving the 3-DoF (LM 3DoF) and 1-DoF (LM 1DoF) problems, by LAGO [2], by CPL-Sync [5], by the rotation averaging in the Theia library [3, 12], by Theia with an additional penalty term (Theia$^{\text{reg}}$ [4]), by projecting updates to y-axis (Proposed$^{\text{alt}}$), by COLMAP [11], and by the proposed method. Curve "Gravity" stands for the approximate upper bound achievable by using gravity direction. A method being accurate is interpreted by its curve close to the top-left corner.

Table 2: Runtime (s) for the projected method and direct optimization on the manifold. The latter is 4-9 times faster.

	EuRoC [1]	KITTI [6]	LaMAR [10]	1DSfM [14]
Proposed$^{\text{alt}}$	0.84	0.74	113.76	26.54
Proposed	**0.20**	**0.20**	**13.98**	**4.38**

the proposed methods nearly overlap, indicating that they obtain similar results. The runtimes are summarized in Table 2. The proposed formulation, directly optimizing on the manifold, leads to a consistent 4-9 times speedup compared to the alternative method.

C Spectrum of Constraints

In this section, we further inspect the formulation that incorporates gravity as an extra penalty term. We show how it allows a spectrum of problems, varying from unconstrained to hard-constrained, and how it can be converted into the proposed method under a specific condition.

Recall the objective function with an extra penalty term (Eq. 21 in the main paper) is formulated as

$$\arg\min_{\{\boldsymbol{R}_i\}} \sum_{(i,j)\in\mathcal{E}} \rho\left(d(\boldsymbol{R}_j^\top \tilde{\boldsymbol{R}}_{ij} \boldsymbol{R}_i, \boldsymbol{I})^p\right) + \sum_i \lambda\rho\left(d(\theta \cdot (\boldsymbol{v}_i - (\boldsymbol{v}_i^\top \boldsymbol{g}_i)\boldsymbol{g}_i))^p\right). \quad (1)$$

Since $d_{\text{geod}}(\boldsymbol{R}, \boldsymbol{S}) \approx \|\theta_R \boldsymbol{v}_R - \theta_S \boldsymbol{v}_S\|_2$, $\|\theta \cdot (\boldsymbol{v}_i - (\boldsymbol{v}_i^\top \boldsymbol{g}_i)\boldsymbol{g}_i)\|_2$ approximates the angle distance of $\boldsymbol{R}_i$ from the closest upright rotation. Thus, the proposed formulation is similar to the penalty in [4].

Parameter λ in Eq. (1) controls the weighting between consistency with the relative pose and the measured gravity. When λ is 0, the second part of the

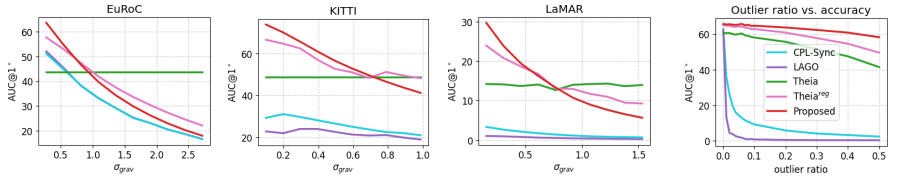

(a) median gravity error vs AUC@1°. $\alpha = 1$ corresponds to σ_{grav} equals $0.54°$, $0.20°$, and $0.31°$ respectively **(b)** outlier vs AUC@1°

Fig. 3: Sensitivity analysis w.r.t. gravity noise and outlier ratio in the measured relative poses. For the left plots, gravity are synthesized by interpolating between noisy measurements and ground truth with different weights. For the right plot, synthetic experiments with grid setting are conducted. In these experiments, different level of outlier are tested with fixed noise level on gravity and relative poses.

objective function can be factored out, and Eq. (1) restores the original 3DoF optimization. Naturally, as λ increases, more emphasis is put on the consistency with measured gravity. As a result, varying the level of λ forms a spectrum of problems, ranging from soft-constrained to hard-constrained ones. Worth noting is the case when $\lambda = \infty$. The optimal solution in this case should possess the following property:

$$\left\| \theta \cdot \left(\boldsymbol{v}_i - (\boldsymbol{v}_i^\top \boldsymbol{g}_i) \boldsymbol{g}_i \right) \right\| = 0. \tag{2}$$

To fulfill the above condition, the downward direction in the estimated rotation should be the same as the measured gravity. If an iterative process solves the optimization, in each step, the updates are projected to act around the gravity direction. As a result, the problem becomes the same as what is presented in Section B.

Besides varying the problem from non-constrained to hard-constrained ones, the above formulation also bridges the proposed method to the original 3DoF optimization. When $\lambda = \infty$, suppose the rotation are pre-aligned that gravity direction $\boldsymbol{g}_i$ is $[0, 1, 0]^\top$. Plugging this into the optimization process gives that $\boldsymbol{v}_x^t = \boldsymbol{v}_z^t = 0$. This means the x, z components remain 0 across the optimization process, thus the linear systems corresponding to these two rotation components become redundant. Removing them results in a smaller linear system equivalent to the proposed system in the paper. This observation also directly suggests the efficiency of our methods. While our method reduces the number of equations in the linear system by $\sim 2m$ where m is the number of pairs, formulation with soft penalty adds $\sim 2n$ additional equations where n is the number of images.

D Sensitivity Analysis

In this section, we analyze the sensitivity of the method in the presence of noise and outliers. In particular, we report the AUC score as a function of gravity noise and the outlier ratio in the pre-estimated relative poses.

To analyze the sensitivity of the method to gravity noise, we alternate the level of noise in gravity by interpolating between measured and ground-truth

gravity. More concretely, it is formulated as

$$\boldsymbol{g}_i^{\alpha} = \frac{\alpha \cdot \boldsymbol{g}_i + (1-\alpha) \cdot \boldsymbol{g}_i^{gt}}{\left\| \alpha \cdot \boldsymbol{g}_i + (1-\alpha) \cdot \boldsymbol{g}_i^{gt} \right\|}, \tag{3}$$

where $\boldsymbol{g}_i$ is the measured gravity, $\boldsymbol{g}_i^{gt}$ is the ground truth gravity, and α controls the level of noise. Such a setting allows for progressively mixing the noise in the measurements to the ground truth direction. We vary α from 0.5 to 5, and results are summarized in Figure 3a. The x-axis in the plot is the median gravity error, and $\alpha = 1$ corresponds to $\sigma_{\text{grav}} = 0.54°$, $0.20°$ and $0.31°$ respectively. From the figure, it can be observed that methods incorporating gravity, *i.e.* the proposed method and Theia with extra penalty term (Theia$^{\text{reg}}$), produce more accurate results compared to the original 3DoF optimization. This finding further motivates the use of gravity in the rotation averaging problem. Additionally, the proposed method obtains the most accurate estimations among all methods. This holds until $\sim 0.7°$ on the datasets we tested. Such levels of accuracy can be satisfied by commercial sensors as the median gravity errors on EuRoC [1] and KITTI [6] are $0.54°$ and $0.20°$, respectively. On LaMAR [10], the median error of HoloLens images is $0.34°$, and that of smartphone images is $0.26°$.

To analyze the sensitivity of the method to outliers, we conducted synthetic experiments with varying levels of outliers. There, we apply the same grid setting as in the synthetic experiments in the main paper. For a brief recap, we assume that the cameras form a 2D grid, and each camera is connected with the closest 24 grid neighbors. In the experiments, the noise level for gravity prior is fixed to be $0.25°$ which is comparable to that of gravity measurements. As for the relative pose, we add zero-mean noise by applying random rotations with a standard deviation of $1°$. To inject outliers, a uniformly randomly drawn set of relative poses is replaced by random rotations. Results are summarized in Figure 3b. The plot shows that the performance of the proposed method and baselines are similar when the relative poses are outlier-free. As the outlier ratio increases, the accuracy in estimation for all methods degrades at different rates. LAGO [2] is subject to drastic performance drop as the outlier ratios increase, validating the claim that its heuristic approach for fixing the period is not robust. CPL-Sync [5] is similarly sensitive to outliers. In contrast, Theia-based baselines and our method are robust in the presence of outliers. Among all, the proposed method undergoes the least performance drop as the outlier ratio increases. For the proposed method, the AUC score at $1°$ only decreases by 11% while those for Theia and Theia with penalty terms drop by 32% and 24%, respectively.

E Effect of Iteration on Period Estimation

LAGO [2] estimates periods for edges from a coarse initialization. Such a heuristic approach often fails on different datasets. For example, on the LaMAR dataset, the AUC score of LAGO at $1°$ is only 0.94, while that of the proposed is 23.85, and on 1DSfM, the scores are 7.96 and 29.41, respectively. We conduct experiments on 1DSfM [14], EuRoC [1], KITTI [6], and LaMAR [10] to analyze the

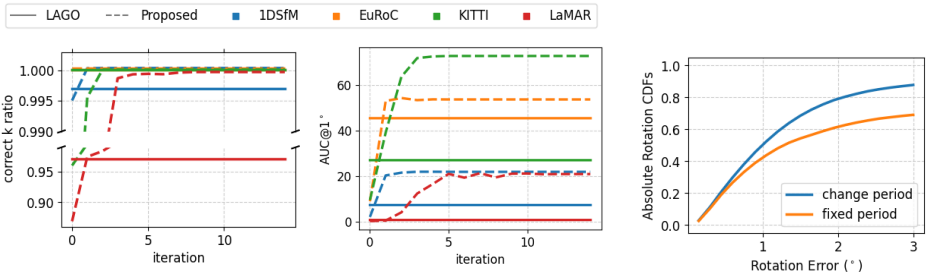

(a) accuracy of period and camera pose estimation across iterations. k is the period number

(b) effect of fixing period

Fig. 4: Period estimation can have a large impact on the accuracy of rotation averaging.

accuracy of period estimation (*correct ratio*) for LAGO [2] and the proposed iterative solution. Also, we analyze the relationship between period estimation and accuracy.

To calculate the ground truth periods, first, the ground truth rotations are globally aligned with the estimated results and are adjusted to be at most π radian to the estimation. Then, ground truth periods for edges are calculated using the same formulation as in LAGO [2] with the adjusted ground truth rotations.

Results are in Fig. 4a. For iteration 0, the rotations for our method are initialized to be identity. As LAGO is not an iterative method, we report its error as a constant number across iterations. From the plot, one can observe that the *correct ratio* is monotonically increasing for the proposed method. Also, the convergence rate is generally fast and the proposed method can achieve almost 100% *correct ratio* within 5 iterations. This indicates the effectiveness of the proposed mechanism. While for LAGO [2], though we can observe it obtain high *correct ratio* on the EuRoC [1] and KITTI [6] datasets, which are within the domain the method was originally tailored for, its performance largely drops on other datasets. As a result, LAGO fails.

The right plot of Fig. 4a shows the relationship between estimation accuracy and iteration. By reading the *correct ratio* from the left plot and the AUC score on the right plot, one can notice a strong positive correlation between these two variables. Across iterations, the accuracy of the rotation orientation increases, leading to a more accurate estimation of periods. These estimations can later contribute to a more accurate rotation estimation, and finally, the estimation converges to the desired result.

Fig. 4b shows the importance of alternating period estimation in the optimization as proposed in the main paper. As from Fig. 4a, the estimations of periods on EuRoC [1] and KITTI [6] are error-free, we only show the result on the LaMAR [10] dataset. From the plot, it can be seen that the accuracy of the rotation averaging is largely affected by the accuracy of initial period estimation,

Table 3: Results on 1DSfM dataset [14]. Columns marked with † are from the respective paper. We use ground truth with 0.5° standard deviation error as gravity.

	DMF-SYNCH† [13]	NeuRoRa† [9]	MSP† [15]	PoGO-Net† [7]	Proposed
ALM	1.2	1.2	1.07	**0.85**	0.90
ELS	0.8	0.6	0.83	**0.43**	0.52
GDM	10.5	2.9	3.69	-	**0.96**
MDR	2.3	1.1	1.09	0.96	**0.67**
MND	0.6	0.6	0.50	**0.37**	0.51
NYC	1.8	1.1	1.12	**0.88**	1.03
PDP	1.0	0.7	0.76	0.81	**0.66**
PIC	-	1.9	1.80	1.75	**0.80**
ROF	1.8	1.3	1.19	**0.69**	0.71
TOL	2.7	1.4	1.25	**0.43**	0.64
TFG	-	2.2	-	1.70	**0.86**
USQ	4.4	2.0	1.85	**1.25**	1.52
VNC	1.6	1.5	1.10	1.44	**0.74**
YKM	1.7	0.9	0.91	**0.72**	**0.72**

even if a robust estimation scheme is deployed. This indicates the necessity of deploying the circular regression.

F Comparison with Learning-based Methods

F.1 Unordered Image Collections

In this section, we compare with the recent learning-based DMF-SYNCH [13], NeuRoRA [9], MSP [15] and PoGO-Net [7]. Results on the 1DSfM dataset are reported in Table 3. The results in the table are directly taken from the respective papers. Given that we do not have a measured gravity direction, similarly as in the main paper, we take the gravity from the ground truth rotations and add zero-mean Gaussian noise with 0.5° standard deviation. The most accurate learning-based method is PoGO-Net. Over the scenes where results are provided (all, except for Gendarmenmarkt), the average rotation error of PoGO-Net is 0.96°, while that of the proposed method is 0.79°. On Gendarmenmarkt, PoGO-Net results are missing due to its high memory requirement, while the proposed method leads to significantly more accurate results than others. Given that the gravity error is lower than 0.5° on most of the tested datasets, we expect even larger differences in practice.

F.2 Sequential Image Collections

On the KITTI [6] and EuRoC [1] datasets, we failed to achieve reasonable results for DMF-SYNCH [13], NeuRoRA [9] and MSP [15] even when using the code provided by the authors and retraining their models. As DMF-SYNCH is a method based on matrix completion, it requires not too sparsely connected graphs, which sequential datasets generally do not satisfy. For the shortest sequence 04 in KITTI, which only contains 277 images, DMF-SYNCH can achieve

a median error at the level of 10°. But for sequences with more than 1000 images, we failed to achieve a median error below 30°. Also, the parameters to learn grow quadratically with the number of images, making the method impractical for sequences with more than a few thousand images.

The sparsity also poses challenges to NeuRoRA [9] and MSP [15], which are graph-neural-network-based methods. They contain two parts: a view-graph cleaning component and a rotation refinement component. The first component in both works rejects outliers and refines relative rotations. Then, NeuRoRA and MSP initialize global rotation estimation via a minimal spanning tree on the input view graph. After coarse initialization, they both proceed with a network called FineNet, which refines the rotations.

We attempted to run NeuRoRA [9] in different ways. We tried the weights provided by the authors, finetuning the network on 1DSfM, finetuning on KITTI, and finetuning on EuRoC. All such tests led to inaccurate results. To eliminate the effect of the potentially failing first step (*i.e.*, view-graph cleaning), we tried providing the ground truth outlier information as initialization for the refinement module. Since the remaining relative poses are generally accurate, the obtained initialization is reasonable, with median errors generally below 3° for EuRoC and KITTI sequences. Despite this accurate initialization, FineNet consistently (with all training strategies) reduced the accuracy instead of improving it. We conclude that these two methods are not suitable for the rotation averaging on KITTI, EuRoC, and LaMAR datasets featuring long sequential image streams.

Table 4: Mean rotation error (°) on sequence 02 and 08 in KITTI [6]. Results with † are taken from the original paper.

	PoGO-Net [7]†	Proposed
KITTI-02	1.08	**0.46**
KITTI-08	2.17	**0.53**

As PoGO-Net [7] has no official code available, we report the errors on the KITTI dataset in Table 4 as provided in the original paper. The proposed method substantially improves upon PoGO-Net on both sequences.

G Results with Gravity-Aligned Essential Matrix

The results when using 3-DoF and 5-DoF relative poses as input are reported in Table 5. The 3-DoF essential matrix solver is incorporated within the same LO-RANSAC framework used in other experiments to estimate gravity-aligned essential matrix. We allow for full 5 DoF optimization in the essential matrix estimation. We can see that the gravity-aligned essential matrix improves almost all results, except for Theia's AVG error on KITTI. The proposed method is still significantly more accurate than Theia [12].

Table 5: Results with relative poses from both 5-DoF solvers and 3-DoF solvers for EuRoC and KITTI dataset.

			AVG (°)	MED (°)	AUC@0.5°	@1°	@2°
EuRoC	Proposed	5DoF	0.65	0.35	32.68	53.89	72.11
		3DoF	**0.62**	**0.34**	**33.94**	**56.64**	**74.12**
	Theia	5DoF	0.88	0.54	22.01	43.30	64.92
		3DoF	**0.77**	**0.48**	**23.56**	**47.05**	**68.80**
KITTI	Proposed	5DoF	**0.39**	0.29	48.71	69.79	83.86
		3DoF	**0.39**	**0.27**	**50.69**	**70.51**	**83.98**
	Theia	5DoF	**0.68**	0.64	25.96	47.48	71.77
		3DoF	2.52	**0.54**	**28.79**	**52.58**	**72.92**

H Implementation Details for Theia Baselines

For the ease of implementation of Theia with regularization terms, we pre-align the camera rotations with the gravity prior. Thus, the estimated rotations are around $(0, 1, 0)^\top$. In this case, the first and the last term for camera rotations should be 0. We can directly penalize these two terms as the deviation from the prior. After appending these linear systems to the system, the problem can be solved as in Theia [12].

I Detailed Experiment Results

Per sequence result for EuRoC [1] and KITTI [6] for the proposed method are summarized in Table 6, 7. From the tables, one can observed that the proposed method achieves accurate rotation estimation consistently across the datasets.

Table 6: Full results of the proposed method for EuRoC [1]

	MH_01	MH_02	MH_03	MH_04	MH_05	V1_01	V1_02	V1_03	V2_01	V2_02	V2_03
mean	0.78	0.21	0.20	0.46	0.18	1.79	0.32	0.58	0.48	0.69	0.63
median	0.26	0.21	0.15	0.35	0.16	1.48	0.32	0.56	0.48	0.58	0.57

Table 7: Full results of the proposed method for KITTI [6]

	01	02	04	05	06	07	08	09	10
mean	0.35	0.46	0.07	0.39	0.20	0.38	0.53	0.23	0.22
median	0.32	0.40	0.06	0.35	0.15	0.24	0.30	0.19	0.20

References

1. Burri, M., Nikolic, J., Gohl, P., Schneider, T., Rehder, J., Omari, S., Achtelik, M.W., Siegwart, R.: The euroc micro aerial vehicle datasets. The International Journal of Robotics Research **35**(10), 1157–1163 (2016)
2. Carlone, L., Aragues, R., Castellanos, J.A., Bona, B.: A fast and accurate approximation for planar pose graph optimization. The International Journal of Robotics Research **33**(7), 965–987 (2014)
3. Chatterjee, A., Govindu, V.M.: Efficient and robust large-scale rotation averaging. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 521–528 (2013)
4. Crandall, D.J., Owens, A., Snavely, N., Huttenlocher, D.P.: Sfm with mrfs: Discrete-continuous optimization for large-scale structure from motion. IEEE transactions on pattern analysis and machine intelligence **35**(12), 2841–2853 (2012)
5. Fan, T., Wang, H., Rubenstein, M., Murphey, T.: Efficient and guaranteed planar pose graph optimization using the complex number representation. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1904–1911. IEEE (2019)
6. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. The International Journal of Robotics Research **32**(11), 1231–1237 (2013)
7. Li, X., Ling, H.: Pogo-net: pose graph optimization with graph neural networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5895–5905 (2021)
8. Moré, J.J.: The levenberg-marquardt algorithm: implementation and theory. In: Numerical Analysis: Proceedings of the Biennial Conference Held at Dundee, June 28–July 1, 1977. pp. 105–116. Springer (2006)
9. Purkait, P., Chin, T.J., Reid, I.: Neurora: Neural robust rotation averaging. In: European Conference on Computer Vision. pp. 137–154. Springer (2020)
10. Sarlin, P.E., Dusmanu, M., Schönberger, J.L., Speciale, P., Gruber, L., Larsson, V., Miksik, O., Pollefeys, M.: Lamar: Benchmarking localization and mapping for augmented reality. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII. pp. 686–704. Springer (2022)
11. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
12. Sweeney, C.: Theia multiview geometry library: Tutorial & reference. `http://theia-sfm.org`
13. Tejus, G., Zara, G., Rota, P., Fusiello, A., Ricci, E., Arrigoni, F.: Rotation synchronization via deep matrix factorization. arXiv preprint arXiv:2305.05268 (2023)
14. Wilson, K., Snavely, N.: Robust global translations with 1dsfm. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III 13. pp. 61–75. Springer (2014)
15. Yang, L., Li, H., Rahim, J.A., Cui, Z., Tan, P.: End-to-end rotation averaging with multi-source propagation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11774–11783 (2021)