

PRAGO: Differentiable Multi-View Pose Optimization From Objectness Detections*

Matteo Taiana, Matteo Toso, Stuart James and Alessio Del Bue
Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia (IIT)
Genoa, Italy
{name.surname}@iit.it

Abstract

Robustly estimating camera poses from a set of images is a fundamental task which remains challenging for differentiable methods, especially in the case of small and sparse camera pose graphs. To overcome this challenge, we propose Pose-refined Rotation Averaging Graph Optimization (PRAGO). From a set of objectness detections on unordered images, our method reconstructs the rotational pose, and in turn, the absolute pose, in a differentiable manner benefiting from the optimization of a sequence of geometrical tasks. We show how our objectness pose-refinement module in PRAGO is able to refine the inherent ambiguities in pairwise relative pose estimation without removing edges and avoiding making early decisions on the viability of graph edges. PRAGO then refines the absolute rotations through iterative graph construction, reweighting the graph edges to compute the final rotational pose, which can be converted into absolute poses using translation averaging. We show that PRAGO is able to outperform non-differentiable solvers on small and sparse scenes extracted from 7-Scenes achieving a relative improvement of 21% for rotations while achieving similar translation estimates.

1. Introduction

Most current Structure from Motion (SfM) approaches [12, 32, 40, 42] require an initial robust estimation of absolute camera poses from multi-view images. This first step is critical, as it is the basis for obtaining a coarse 3D point cloud by triangulation and then feeding standard non-linear solvers [1, 46] that estimate a refined set of camera poses (rotation and translation) and 3D geometry. Given the non-linearity of the problem, an initialization farther away from the optimal solution will increase the chances of obtaining

poor-quality reconstructions or failing entirely.

As SfM pipelines evolved through time, a new trend is to inject semantics in multi-view problems. Semantic and object SfM approaches [4, 11] provide not only a raw 3D point cloud, but they are able to localize objects in 3D and to estimate their pose. Object semantic is also used in SLAM methods [7, 15, 29, 33, 49] to better localize the camera and objects in structured scenes and to build 3D scene graphs for improved scene understanding or robotic navigation [38].

On a similar trend, this paper explores how the semantic information of objects can help to estimate better absolute camera rotations and translations in an SfM scenario. In particular, we address one of the hardest cases of SfM, small-scale scenes with fewer cameras and wider baselines, as most approaches make assumptions of large and dense camera pose graphs, with relatively short baselines or expendable images. Whenever few cameras are available, obtaining a robust initialization of camera poses is harder, as the 2D correspondences are fewer and noisier, so even outliers rejection methods [47] struggle to find good inliers or eliminate too many cameras/points from the initial set.

Previous work has shown that objectness-based image detections can refine relative poses [44] by reasoning on higher-level semantic information. The method defines a pose graph where each node represents a camera and its corresponding detections, while edges relate together nodes that have matching detections in multi-view. Taiana *et al.* [44] proposed a novel graph-based computational approach embedding the images' detections on the nodes of the graph, and the aggregate of predictions of the refined relative pose on the edges. Message passing propagated information in an MLP-style alternating updates between nodes and edges. This approach successfully improved relative poses, but was trained independently from the rest of the pipeline, and relied on an optimization algorithm to perform motion averaging (*i.e.* EIG-SE3 [2]).

In contrast, we propose PRAGO, a differentiable pose estimation method, which focuses on small and sparse pose graphs where camera outlier rejection (e.g. 1DSfM [47])

*This project has received funding from the European Union's Horizon research and innovation programme under grant agreement No 870743 and under grant agreement No 101079995.

might be unpractical. PRAGO uses noisy pairwise camera poses estimated from an objectness-based approach, refines the relative rotations and then integrates rotation averaging to estimate absolute camera poses. As we show experimentally, optimizing these tasks together can provide a significant improvement. As PRAGO does not apply camera outlier rejection but instead focuses on the consistency of the matched objectness detections, we avoid making decisions early on the viability of the pose graph, which makes the method highly suitable for small and sparse graphs where the removal of even a couple of cameras would have a detrimental effect in later steps. Specifically, we propose a graph where edges encode the initial relative pose and bounding box on each respective image, while the node simply store the intrinsic parameters for each camera. Through message-passing updates between edges and nodes in the graph, the nodes are updated and a new pose is regressed as input to motion averaging. At this point, PRAGO iteratively refines the pose graph, re-weighting the contribution of rotational information mimicking an optimization process, where the RAGO [27] blocks incrementally improve the accuracy of the estimates. The rotations are combined with the refined translations to recover absolute camera poses. To understand the importance of pose refinement within PRAGO, we provide an analysis of the orientation error and identify that it frequently fixes chirality errors. In addition, we compare to state-of-the-art combinations of independent methods, both differentiable or optimization-based.

We make three major contributions: *i*) we introduce the PRAGO network for differentiable pose estimation with objectness detections where the back-propagation improves performance in contrast to non-differential pipelines; *ii*) we revise the objectness pose refinement [44] which leads to improved accuracy; and *iii*) we provide an analysis of how objectness-based pose refinement improves relative poses and its positive effect on chirality.

2. Related work

In this section, we provide an introduction to relative pose estimation (Sec. 2.1) and motion averaging (Sec. 2.2), a more detailed discussion on methods dedicated to averaging the rotation (Sec. 2.3) and translation (Sec. 2.4) components of the pose. For a full review of SfM see Bianco *et al.* [6].

2.1. Relative Pose Estimation

The first step for obtaining a set of absolute camera poses is to estimate pairwise relative poses from 2D image correspondences. Given a pair of images, their relative pose can be easily estimated via the five-point algorithm [26, 34] or the eight-point algorithm [19]. Outlier-rejection strategies such as RANSAC [14] are then used to obtain a robust estimate with a desirable low noise. Computing a rotation matrix and a translation unit vector from the Essential ma-

trix presents an inherent ambiguity, though, as two possible solutions exist [34]. One of the two solutions per view is chosen based on chirality considerations (triangulated 3D points) [20]. Picking the incorrect combinations results in outlier errors that can be discarded in the following steps.

More recent works on relative pose estimation leverage deep learning to directly predict a relative pose given two images [5, 28], or a set of feature matches and a sparse scene reconstruction [31]. Learning-based models can also exploit cues like light source directions, vanishing points, and scene symmetries [9], or use Neural Radiance Fields (NeRF) [30] to predict relative poses [50]. Finally, works like PoserNet [44] combine the two approaches, initializing the relative pose via geometrical methods, that are more efficient and do not require training, and then training a model to predict and correct the noise on the predictions.

Noisy relative poses, and especially outliers, have a big impact on motion averaging algorithms, so such algorithms are often designed to explicitly deal with these kinds of errors. Classical approaches usually address this by identifying the outliers measurements by imposing cycle consistency, *i.e.* chaining relative transformations along a closed loop [52], or by projecting the scene graph in a lower dimension and solving a minimum feedback arc set problem to identify the outliers, like 1DSfM [47]. Considering that in the case of sparse or small pose graphs rejecting relative poses can lead to catastrophic results (a graph with multiple components does not contain enough information for estimating coherent absolute poses), we design a pose refinement module to learn to refine, rather than to reject them.

2.2. Motion Averaging

Relative poses are combined to extract absolute poses through optimization algorithms that simultaneously optimize over the orientation and translation components of the pose [2, 3, 10, 32, 48, 51]. Other optimization-based approaches instead decouple rotation estimation from translation and scene reconstruction, and focus on either rotation [24, 37] or translation [54] averaging. Finally, there are methods that focus on object detections, performing diffusion and clustering to compute the candidate poses of the input objects, optimizing the best pose for each object and then refining the camera poses [43].

2.3. Rotation Averaging

To recover the absolute camera orientations, some methods decouple in Bundle Adjustment (BA) rotation estimation from translation and scene reconstruction [24]; or improve on it by combining a global optimizer with fast view graph filtering and a local optimizer [10]. Other methods propose incremental approaches, prioritizing the averaging over camera pairs whose relative pose is supported by more triplets in the view graph, to increase robustness

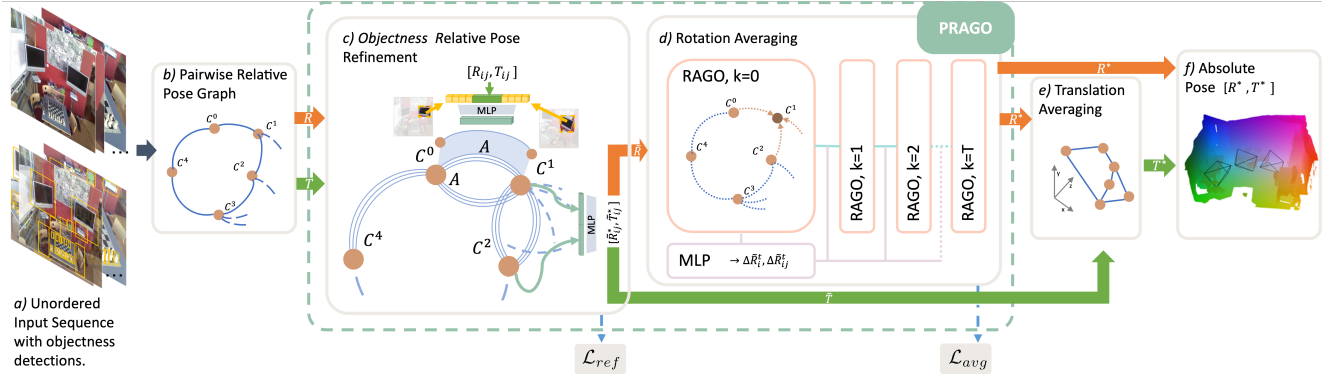


Figure 1. From a set of unordered images *a)*, we compute objectness detections and construct an initial relative pose graph from *b)* the 5-Point pose estimates. *PRAGO* refines the relative pose using the objectness detections resolving, among others, the chirality issues before inferring the rotational poses using RAGO blocks which iteratively infer the rotational pose by reweighting the graph edges. The absolute rotations and refined translations are then combined in *e)* translation averaging using BATA to construct the *f)* absolute camera poses.

to outliers [25]. Regarding GNN-based approaches, NeuRoRA [37] combines a GNN that generates outlier scores and an estimated rotation error for each input relative pose (CleanNet), with a graph network that performs rotation averaging only on the inliers. Subsequent works on multiple rotation averaging include GNNs that perform both initialization and optimization through a differentiable multi-source propagation module [48], or iterative approaches like RAGO [27], which learns to incrementally optimize relative pose estimates using only supervision on the relative poses. Drawing from the success of NeRFs in computer vision, other methods perform multiple motion averaging in the neural volume rendering framework [21]. More recent works investigate the possibility of unsupervised models for rotation averaging, based on deep matrix factorization [45] and directly propagating the uncertainty from the keypoint correspondences used in relative pose estimation, to more accurately model the underlying noise distributions [53].

2.4. Translation Averaging

Translation averaging uses as input pairwise relative translations and the (possibly) noisy absolute rotations from the previous rotation averaging step. Classical approaches [18] predict absolute translation by optimizing the l_2 norm of the vector product between each input relative translation direction and the distance between the corresponding predicted absolute translations. Subsequent works improved results by relaxing the loss to l_∞ [32], by introducing a least unsquared deviations (LUD) form [35], or by redefining the problems as the minimization of the projection of the absolute translations difference onto the relative translation direction [17]. These methods manage to frame the problem as the minimization of a convex cost function, but the scale ambiguity on each translation still requires extensive preprocessing strategies. Other approaches consider

instead angular residuals θ_{ij} between the predicted and observed translation directions, framing the task as a quasi-convex problem on $\tan \theta_{ij}$, solved through a sequence of Second Order Cone Programming (SOCP) problems [41]. To reduce the impact of outliers, other works directly minimize a non-linear cost function of $\sin \theta_{ij}$ via Levenberg-Marquard [47] or iteratively reweighted Least Squares [54].

3. PRAGO for Objectness Refined Rotation Averaging

We propose PRAGO, an end-to-end network for relative pose refinement and rotation averaging, aimed towards the estimation of absolute camera poses $[R^*, T^*]$. Our model takes as input a set of images with matched objectness detections (Fig. 1a) and, as proposed in PoserNet [44], we exploit information on matched detections to define co-visibility, and employ the 5-point algorithm to generate relative pose estimation (Fig. 1b). This results in a sparsely connected view graph, with noisy relative pose estimates which may contain outliers. First, we use objectness-based relative pose refinement (Fig. 1c) to correct or refine outliers (Sec. 3.2). Then the refined poses are passed to a rotation averaging module (Fig. 1d) to produce a set of absolute camera orientations (Sec 3.3). Both refined relative poses and absolute rotations are input to a translation averaging algorithm (Fig. 1e) to predict the absolute camera poses (Fig. 1, Sec. 3.4). To frame the problem, we first establish the preliminaries and notation in Sec 3.1, before defining the PRAGO components.

3.1. Preliminary Problem Formulation

Motion averaging is defined as the problem of recovering the absolute orientation \hat{R} and translation \hat{T} for a set of N cameras, given a set of noisy observed relative orientations R_{ij} and translations unit vectors T_{ij} . For a pair of cameras

(i, j) , given perfect relative transformations \hat{R}_{ij} and \hat{T}_{ij} , their relations with absolute poses are defined as:

$$\hat{R}_{ij} = \hat{R}_i \hat{R}_j^T, \quad \hat{R}_i \hat{T}_{ij} = \frac{\hat{T}_j - \hat{T}_i}{\|\hat{T}_j - \hat{T}_i\|_2}. \quad (1)$$

The observed relative poses are, however, affected by noise and can include outliers.

A standard solution to motion averaging is to split the problem into two optimization steps, one for the rotations and one for the translations, computing the absolute poses more compatible with the observations. The absolute orientations R^* are estimated by minimizing a robust cost function $\rho(\cdot)$ of the distance $d^R(\cdot)$ between R_{ij} and the predicted $\tilde{R}_{ij} = \tilde{R}_i \tilde{R}_j^T$ as:

$$R^* = \arg \min_{R=\{R_i\}} \sum_{(i,j) \in \mathcal{E}} \rho(d^R(R_{ij}, R_i R_j^T)). \quad (2)$$

The solution to this problem is, however, not unique. Applying the same rigid transformation to all absolute orientations \bar{R} does not change the relative orientations: $R_i R_j^T = R_i \bar{R} \bar{R}^T R_j^T = (R_i \bar{R})(R_j \bar{R})^T$. To make the solution unambiguous, a fixed *gauge transformation* is needed.

Given the predicted R^* , a cost function σ and a vector distance $d^T(\cdot)$, the absolute translations are obtained as:

$$T^* = \arg \min_{T=\{T_i\}} \sum_{(i,j) \in \mathcal{E}} \sigma(d^T(T_{ij}, R_i^{*T}(T_j - T_i))). \quad (3)$$

In Eq. 2 and Eq. 3, the cost functions $\rho(\cdot)$ and $\sigma(\cdot)$ model the noise on the relative transformations; the same problems can be approached with a learning paradigm, training a network to implicitly model the noise. In this case, graph networks are especially well suited, as the problem can be naturally described as a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where nodes $N_i \in \mathcal{N}$ represent the cameras, and edges $e_{ij} \in \mathcal{E}$ represent the observed relative poses between nodes i and j .

3.2. Relative Pose Refinement

To learn how to correct the relative pose outliers, we opt for a graph-based approach. Each camera C^i is associated to a node, whose initial embedding is defined as $h = [f, H, W]$, where f is the camera focal length and H and W are the height and width of the associated image. We then draw for each image pair (i, j) a set of directed edges, for each pair of matched objectness detections (l, m) with $l \in i, m \in j$. This is, in contrast, with [44], which draws a single edge between a pair of nodes, and handles the multiple detections on one node in an *ad hoc* way. We define the edges embeddings by concatenating the relative pose between the cameras, $[R_{ij}, T_{ij}]$, and the bounding boxes BB_l^i and BB_m^j associated to the two detections. The former are represented as a quaternion for the rotation and as a unit vector for the

translation. The latter are defined as $BB = [x, y, bb_w, bb_h]$, where x and y are the pixel location of the upper left bounding box corner and bb_w and bb_h respectively represent the bounding box's width and height in pixels. This graph is passed through two GATv2 [8] layers, modified to include the edge embeddings in the node's updates (drawing from Edge-GAT [22]) and to apply attention. During this process, only the node embeddings are updated, while the edge embedding preserve information about the bounding boxes locations and the input relative poses. In contrast, [44] applies an *ad hoc* GNN which updates both edge and node information. To update the relative pose between camera pair (i, j) , we first use an MLP to merge together the concatenation of the learned embeddings of the nodes, $h_{ij} = \text{MLP}([h_i, h_j])$. Then, a second MLP takes as input the noisy relative pose and this joint embedding, to generate a refined pose as

$$\tilde{R}_{ij}, \tilde{T}_{ij} = \text{MLP}([R_{ij}, T_{ij}, h_{ij}]). \quad (4)$$

Both the graph network and the two MLPs are trained to minimize the loss

$$\mathcal{L}_{ref} = \mathcal{L}_{orient} + \mathcal{L}_{tr_dir} + \gamma(\mathcal{L}_{q||} + \mathcal{L}_{tr||}), \quad (5)$$

where \mathcal{L}_{orient} and \mathcal{L}_{tr_dir} are angular losses, based respectively on the quaternion combination $\tilde{R}_{ij} \circ \hat{R}_{ij}$ and on $\cos^{-1}(\hat{T}_{ij} \cdot \tilde{T}_{ij} / \|\hat{T}_{ij}\| \|\tilde{T}_{ij}\|^{-1})$; $\mathcal{L}_{q||}$ and $\mathcal{L}_{tr||}$ are respectively normalization losses on the quaternions and unit vectors; and γ is a coefficient used to tune the contribution of the different components of the loss empirically set to 0.5.

3.3. Rotation Averaging

Given the refined relative poses, we combine them to predict the best fitting absolute orientations \bar{R} , drawing from the RAGO [27] architecture. This module uses two Gated Recurrent Units (GRU) to iteratively update the global orientation of node i and the relative orientations of its nearest neighbors \mathcal{N}_i . First, the absolute poses R_i are randomly initialized; then, R_i, \tilde{R}_{ij} and \mathcal{G} are passed to four MPNN to generate features g for each node and edge, and initialization embeddings for the hidden states of the node and edges in the GRUs. At each optimization step $k = 1, \dots, K$, the cost of the current absolute and relative orientations for node i are respectively d_i^k and d_{ij}^k :

$$d_i^k = \frac{1}{\mathcal{N}_i} \sum_{j \in \mathcal{N}_i} \|R_i^k - \tilde{R}_{ij} R_j^k\|. \quad (6)$$

$$d_{ij}^k = \|R_{ij}^k - R_i^k R_j^k{}^T\| + \|R_{ij}^k - \tilde{R}_{ij}\|. \quad (7)$$

Here, R_{ij}^t is an estimated relative pose, obtained by rectifying $R_i^k R_j^k{}^T$ to align it to \tilde{R}_{ij} ; the cost of Eq. 7 will then act as an outlier score, as it will tend to 0 for inliers and will be large for outliers. Another MPNN extracts features \mathcal{C}^k

from the cost function, that are concatenated with the current orientation R^k and the features into $I^k = [C^k, R^k, g]$. The current orientations can then be updated as:

$$h^k = \text{GRU}(h^{k-1}, I^k) \quad \Delta R^k = \Phi(h^k), \quad (8)$$

where h^k is the hidden state of the GRU nodes as time k , Φ is an MLP. ΔR^k is a correction to the orientation estimate, allowing to update the current predictions as

$$R_i^{k+1} = R_i^k \Delta R_i^k \quad R_{ij}^{k+1} = R_{ij}^k \Delta R_{ij}^k. \quad (9)$$

To speed up convergence, during training nodes and edges are optimized in alternated fashion; for a fixed number of node (K_n) and edge (K_e) iterations, the loss is given by:

$$\begin{aligned} \mathcal{L}_{avg} = \frac{1}{|\mathcal{E}|} \sum_{k=1}^{K_n} \sum_{(i,j) \in \mathcal{E}} \gamma^{K_n-k} \|R_i^T R_j^k T - \hat{R}_{ij}\|_1 + \\ \frac{1}{|\mathcal{E}|} \sum_{k=1}^{K_e} \sum_{(i,j) \in \mathcal{E}} \gamma^{K_e-k} \|R_{ij}^T - \hat{R}_{ij}\|_1. \end{aligned} \quad (10)$$

Note that this loss enforces supervision only on the ground-truth relative pose \hat{R}_{ij} , and not on the absolute poses. This avoids the gauge ambiguity issue. At convergence, the absolute rotation poses prediction is given by $R^* = \{R_i^K\}$.

The PRAGO components for pose refinement and rotation averaging modules are fully differentiable, and can be trained end-to-end by minimizing the final loss:

$$\mathcal{L} = \alpha \cdot \mathcal{L}_{ref} + \beta \cdot \mathcal{L}_{avg}, \quad (11)$$

where α and β are scaling coefficients. In practice, the residual on \mathcal{L}_{ref} is approximately two orders of magnitude larger than \mathcal{L}_{avg} , and we set $\alpha = 1, \beta = 0.01$.

3.4. Translation Averaging

The relative poses from Pose Refinement and absolute orientations generated by Rotational Averaging are combined to predict absolute translations. For this, we apply BATA [54], which frames translation averaging as an optimization problem of a bilinear objective function. BATA introduces a variable that performs normalization on the translations to base the loss on the angular error on the translation unit vectors, avoiding numerical instability due to scale ambiguity (Sec. 2.4). This prevents translation vectors with a larger module from dominating the loss; in fact, at parity of noise on the orientation and on the module, the difference between ground truth and estimated translation vector will be proportional to the vector module. In contrast, a loss based on an angular distance will treat all relative translations similarly, regardless of the actual pairwise distance. The absolute translations of Eq. 3 can be found by minimizing:

$$T^* = \arg \min_{T,d} \sum_{(ij) \in \mathcal{E}} \mu(\|(T_j - T_i)d_{ij} - \tilde{T}_{ij}\|_2) \quad (12)$$

where $\mu(\cdot)$ is a robust M-estimator and d is a non-negative variable, used to scale the module of $T_j - T_i$. This can be solved by including two constraints:

$$\sum_{i \in \mathcal{V}} T_i = 0 \quad \sum_{(ij) \in \mathcal{E}} \langle T_j - T_i, \tilde{T}_{ij} \rangle = 1, \quad (13)$$

which respectively impose a reference system and a scale. Note that Eq. 12 essentially is a Least Squares objective embedded in an M-estimator to account for the presence of outliers in the input data, and can be efficiently resolved using an iteratively reweighted Least Squares scheme. As rotation averaging is typically considered more reliable than translation averaging, BATA uses the following residuals:

$$\epsilon = \sqrt{\|(T_j - T_i)d_{ij} - \tilde{T}_{ij}\|_2^2 + \|\tilde{R}_{ij} - R_i^* R_j^{*T}\|_2^2}, \quad (14)$$

to update the weights. Combined with the absolute orientations, this results in full camera poses $[R^*, T^*]$, which we can compare against the ground truth poses $[\hat{R}, \hat{T}]$ to evaluate the full pipeline.

4. Evaluation

We first present results highlighting the effectiveness of PRAGO on real-world data, comparing our pipeline against baseline motion averaging approaches (Sec. 4.1); then, we report on ablation tests that explore how to best combine the components of the loss of PRAGO and how our relative pose refinement compares against PoserNet [44] (Sec. 4.2). Finally, in Sec. 4.3, we investigate how the relative pose accuracy changes after pose refinement, and how these changes affect the accuracy of absolute pose estimates.

Dataset – For the evaluation setting, we adopt the 7-scenes dataset [16] following the methodology of [44]. In particular, we use the view graphs from the “small graphs dataset”, which consists of 14000 training graphs (2000 per scene) and 7000 test graphs (1000 graphs per scene). Each graph’s image contains objectness detections generated through a pre-trained Object Localization Network (OLN) [23], and matches defined by an oracle, using the intersection over union of the projection of one bounding box onto the other, for each possible pair of bounding boxes. Each graph contains eight nodes, with a connectivity dictated by having at least three matched detections, and relative camera poses between the two images estimated using SuperGlue [39], the five-point algorithm and RANSAC. This results in an average of 15.5 edges per graph out of the possible 28 edges (55.36% connectivity). We note that 3.8% of the test graphs contain multiple components, which is not an issue for relative pose (as the focus of [44]) but a problem for absolute pose estimation. To overcome this issue, we present results on the 6731 test graphs, which contain one single component, and re-evaluate the PoserNet baseline in ablation. For a full discussion of the object detections, matching and graph generation process see [44].

| Pipeline | Orientation | | | | | | Translation | | | | | | |
|--------------------------|---------------------------|---------------|---------------|---------------|---------------|---------------|--------------|--------------|---------------|---------------|---------------|---------------|---------------|
| | 3° | 5° | 10° | 30° | 45° | η | 0.05 m | 0.1 m | 0.25 m | 0.5 m | 0.75 m | η | |
| Non-fully differentiable | PoserNet + EIG-SE3 | 28.15% | 60.26% | 85.72% | 94.43% | 96.55% | 4.21° | 0.04% | 2.69% | 34.88% | 76.21% | 94.19% | <u>0.33 m</u> |
| | PoserNet + NeuRoRA + BATA | 22.60% | 55.91% | 84.89% | 95.97% | 97.37% | 4.59° | <u>1.20%</u> | <u>10.07%</u> | <u>36.89%</u> | 66.68% | 85.77% | 0.35 m |
| | PoserNet + RAGO + BATA | <u>28.84%</u> | 59.68% | <u>87.28%</u> | <u>96.17%</u> | <u>97.53%</u> | 4.26° | 1.50% | 11.59% | 40.97% | <u>71.58%</u> | <u>90.09%</u> | 0.31 m |
| Differentiable | PRAGO (Ours) + BATA | 37.88% | 65.59% | 89.29% | 96.32% | 97.56% | 3.77° | 0.34% | 4.47% | 33.23% | 67.84% | 89.41% | 0.36 m |

Table 1. The Table presents rotation and translation median errors η and the percentage of test graphs that falls under a given error threshold for rotation (3°, 5°, 10°, 30°, 45°) and translation (0.05 m, 0.1 m, 0.25 m, 0.5 m, 0.75 m). We evaluate three baseline approaches that use non-fully-differentiable pose refinement against PRAGO. Best and second best values for each column are highlighted in bold and underlined, respectively.




Figure 2. Histograms with natural (left) and logarithmic frequency axis (right) of absolute orientation errors produced by PRAGO (green) and by applying the pose refinement and the rotation averaging modules independently (pink). Histograms are overlaid with partial transparency.

4.1. PRAGO Evaluation

We compare the performance achieved by the fully-differentiable PRAGO with that achieved by applying the pose refinement module and the rotation averaging module trained independently. For reference, we also include baseline results obtained using the EIG-SE3 [2] motion averaging algorithm, or using the NeuRoRA [37] rotation averaging method instead of RAGO. Results are reported in Tab. 1, expressed in the form of the median error η (the smaller the better) and as the percentage of testing graphs with and average error smaller than predefined thresholds (the higher the better). Such thresholds are 3, 5, 10, 30 and 45 degrees for the angular metric, and 0.05, 0.1, 0.25, 0.5 and 0.75 for errors expressed in meters. It can be seen that applying PRAGO achieves better orientation accuracy than the non-fully-differentiable baselines, both in terms of median error (3.77°, against the second best of 4.21° achieved by PoserNet followed by EIG-SE3) and in terms of percentage of graphs with errors below the specified thresholds. This improvement comes at the cost of a slightly lower accuracy for the estimated translations, with a median of 0.36 m compared with the 0.31 m achieved by running PoserNet, RAGO and BATA in a sequence. The histogram of orientation errors obtained by applying PRAGO versus applying separately trained pose orientation and rotation averaging modules are shown in Fig. 2. It can be seen that PRAGO

produces an overall better output distribution, with more errors concentrated bin corresponding to lowest error values.

4.2. Ablations

To provide further insight on the design choices of the proposed pipeline, first we ablate across the hyper-parameters of PRAGO (Eq. 11); then, we compare our revised version of objectness pose refinement against PoserNet.

Combining partial loss functions – When training the fully differentiable PRAGO, we have the problem on how best to combine the losses as in Eq. 11. We devised three strategies: driving the learning using exclusively the rotation averaging loss ($\alpha = 0, \beta = 1$), driving it by naively adding the two loss components ($\alpha = 1, \beta = 1$) and, finally, driving the learning by weighting the loss components so that they would have comparable values at the end of the training of the single components ($\alpha = 1, \beta = 0.01$). As can be seen in Tab. 2, relying exclusively on the rotation averaging loss leads to the overall best result in terms of rotation averaging, but this comes with a large penalty for translation averaging, with a median error almost twice as high as that achieved by the other loss combination schemes (0.61 m versus 0.36 m and 0.32 m).

| \mathcal{L}_{ref} | \mathcal{L}_{abs} | Orientation error (deg) | Translation error (meters) |
|---------------------|---------------------|-------------------------|----------------------------|
| 0.00 | 1.00 | 3.75 | 0.61 |
| 1.00 | 1.00 | 3.77 | 0.36 |
| 1.00 | 0.01 | 3.87 | 0.32 |

Table 2. Absolute orientation and translation errors achieved by using different weights for combining PRAGO’s losses. Weights for each component are listed in the \mathcal{L}_{ref} and \mathcal{L}_{abs} columns.

| Pipeline | Orient. error (deg) | Transl. error (meters) |
|------------------------------|---------------------|------------------------|
| PoserNet + EIG-SE3 | 4.97 | 0.32 |
| Enhanced pose ref. + EIG-SE3 | 4.21 | 0.33 |

Table 3. Absolute orientation and translation errors achieved by using either PoserNet or the proposed relative pose refinement architecture prior to applying EIG-SE3.

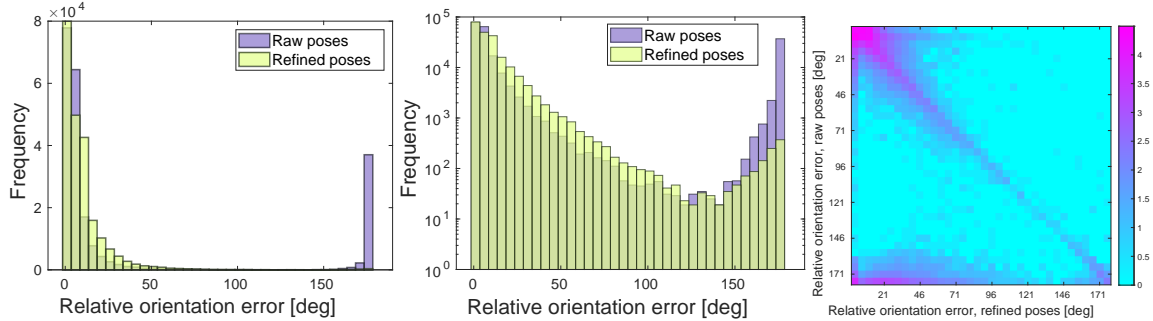


Figure 3. *Effect of objectness-based pose refinement on the input relative poses.* Left and center, natural and logarithmic histograms of edge-wise orientation error, respectively, for the test set with raw initial poses (violet) and refined poses (green). Right, shows a matrix representing the refinement effect based on relative orientation error. The color represents \log_{10} of the number of graphs that belong to one cell of the matrix.

Enhanced objectness-based pose refinement – We compare the performance obtained by EIG-SE3 when applied respectively on relative poses refined either by PoserNet or by the pose refinement architecture we propose. As it can be seen in Tab. 3, applying the proposed refinement architecture rather than PoserNet shows a drop in median orientation error from 4.97° to 4.21° , while the corresponding impact on median translation error is negligible.

4.3. Analysis of the effect of pose refinement

In the initial formulation of objectness relative pose refinement in PoserNet [44] showed the benefits of refining relative poses and, thus, leading to more accurate absolute poses. We explore this further to understand and explain a possible cause for this improvement firstly in its relative setting then in absolute when applying motion averaging.

Refining relative poses with objectness – We characterize the effect of the proposed enhanced pose refinement architecture has on estimates affected by different levels of noise. From the edge-wise orientation error histograms in Fig. 3 it can be seen that pose refinement greatly reduces the number of estimates with orientation error larger than 160° , the ones which likely stem from an incorrect chirality choice. The transfer matrix on the right side of Fig. 3 visualizes the effect pose refinement has on estimates affected by different levels of noise. Elements below and above the main diagonal correspond to estimates whose accuracy is increased or decreased, respectively. On the bottom-left corner of the matrix, it can be seen that many estimates with the wrong chirality are corrected to very accurate values (many of the input estimates that fall in the bin centered at 176° are corrected to estimates that fall in the bin centered at 1°). This chirality-correction property is not perfect, though, as other such estimates are corrected to estimates with lower, but non-negligible levels of error (e.g., estimates that are moved from the bin centered on 176° to that centered on 46°). Other effects are visible on the matrix, though perhaps not

| Motion Averaging Pipeline | Type of input poses | Orientation error (deg) | Translation error (meters) |
|---------------------------|---------------------|-------------------------|----------------------------|
| EIG-SE3 | Raw | 25.22 | 0.61 |
| | Refined | 4.21 | 0.33 |
| NeuRoRA, BATA | Raw | 5.84 | 0.61 |
| | Refined | 4.59 | 0.35 |
| RAGO, BATA | Raw | 5.09 | 0.60 |
| | Refined | 4.26 | 0.31 |
| PRAGO (Fully diff.) | Raw | 3.77 | 0.36 |

Table 4. Impact of relative pose refinement on absolute poses. Orientation and translation error are shown for different motion averaging pipelines fed with either raw or refined relative poses. The performance of PRAGO is shown for reference.

as clearly or not with a straightforward interpretation. On the top-left corner of the matrix, it can be seen that estimates with lower input errors can be affected both positively and negatively by the refinement. Finally, the high values present on the first column of the matrix indicate that many estimates are corrected to very low error values.

This analysis highlights that objectness-based relative pose refinement has a strong positive impact on estimates affected by incorrect chirality, while also improving other estimates, suggesting that it could be especially important when coupled with motion averaging systems which do not implement an outlier rejection strategy.

Impact of relative pose refinement on Motion Averaging

– To understand the cumulative effect of applying objectness relative pose refinement, we compare the performance achieved by pairing the proposed pose-refinement architecture with three different Motion Averaging systems: EIG-SE3, NeuRoRA followed by BATA, and RAGO followed by BATA. It is important to notice that EIG-SE3 is not trainable, so it is applied without any customization to the problem. For NeuRoRA, we fine-tuned its two pre-trained net-

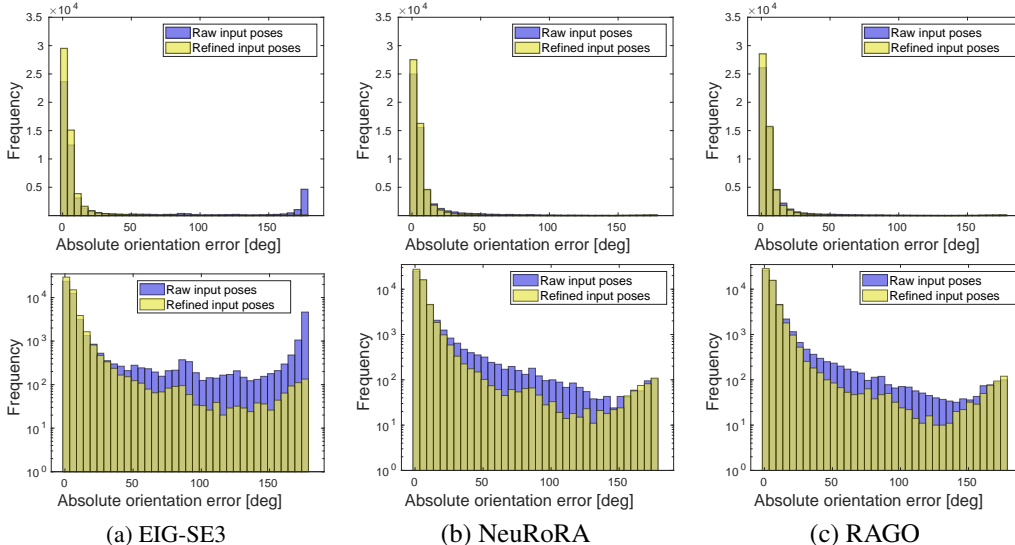


Figure 4. Natural (top) and logarithmic (bottom) histograms of the absolute orientation error achieved by different rotation averaging methods (EIG-SE3, NeuRoRA and RAGO) applied either on raw relative poses (blue) or on refined relative poses (yellow).

works (CleanNet and FineNet) on the training data from 7-Scenes as provided in [44]. We trained RAGO from scratch on the same training data and applied it on the test data directly. The authors of RAGO mention that they preprocess real data with NeuRoRA’s CleanNet, but for our experiments, we opted to apply it directly on raw relative poses or to the refined ones.

We report the orientation and translation errors on raw and refined relative poses in Tab. 4, and show the corresponding histograms of orientation errors in Fig. 4. The results show that applying objectness-based pose refinement is beneficial for all motion averaging algorithms, with the median translation error dropping from roughly $0.6 m$ to roughly $0.3 m$ for all of them. In terms of absolute orientation errors, EIG-SE3 is the algorithm which benefits the most from the relative pose refinement. When applied on raw relative transformations, EIG-SE3 produces the estimates with the highest median rotational error (25.22°), highlighting a sensitivity to highly noisy input data. As can be seen in Fig. 4a, EIG-SE3 is not able to solve all the chirality errors present in the raw input data (a considerable number of its output estimates have errors close to 180°). When applied on refined relative poses (where most very high errors have been reduced), instead, EIG-SE3 produces the lowest median orientation error of all methods (4.21°), narrowly beating RAGO.

4.4. Implementation details

As the input relative poses, we used the five-point poses from [44]. We implemented the architecture for objectness-based pose refinement using the PyTorch framework [36] and PyTorch Geometric [13], and trained it using the Adam

optimizer with learning rate 0.0001 and batch size 1. We used a schedule to reduce the learning rate by a factor of 0.316 with a patience of 3 epochs. We used the author implementation for each of the other methods, EIG-SE3 [2], RAGO [27], NeuRoRA [37] and BATA [54]. For training or fine-tuning those methods, we used the training set defined in [44], and the hyperparameters present in their code or defined in their papers. For the training PRAGO, we started from versions of the pose refinement and rotation averaging models trained in isolation, and fine-tuned the joint system using the same hyperparameters described earlier.

5. Conclusion

We have presented PRAGO, a novel differentiable camera pose estimation approach exploiting objectness detections. Our method takes as input a set of unordered images with their corresponding detections to estimate initial pair-wise relative poses. The pairwise rotations are then refined with a novel approach that can work efficiently with rotational averaging iterative optimization blocks. By differentiating through the method, we have shown that the PRAGO method can achieve superior performance in contrast to concatenating its pre-trained components. In addition, through analysis, we have shown that a key cause of the performance increase is improving the pair-wise orientation errors, especially in the case of chirality, during pose refinement, a step largely unconsidered by large-scale techniques with abundant images. PRAGO has demonstrated the benefits of learning in an end-to-end fashion, as opposed to focusing on the optimization of individual tasks alone.

References

- [1] Sameer Agarwal and Keir Mierle. Ceres solver: Tutorial & reference. *Google Inc*, 2(72):8, 2012. **1**
- [2] Federica Arrigoni, Beatrice Rossi, and Andrea Fusiello. Spectral synchronization of multiple views in se (3). *SIAM Journal on Imaging Sciences*, 9(4):1963–1990, 2016. **1, 2, 6, 8**
- [3] Federica Arrigoni, Andrea Fusiello, Elisa Ricci, and Tomas Pajdla. Viewing graph solvability via cycle consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5540–5549, 2021. **2**
- [4] Sid Yingze Bao and Silvio Savarese. Semantic structure from motion. In *CVPR 2011*, pages 2025–2032. IEEE, 2011. **1**
- [5] Axel Barroso-Laguna, Eric Brachmann, Victor Adrian Prisacariu, Gabriel J Brostow, and Daniyar Turmukhambetov. Two-view geometry scoring without correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8979–8989, 2023. **2**
- [6] Simone Bianco, Gianluigi Ciocca, and Davide Marelli. Evaluating the performance of structure from motion pipelines. *Journal of Imaging*, 4(8), 2018. **2**
- [7] Sean L Bowman, Nikolay Atanasov, Kostas Daniilidis, and George J Pappas. Probabilistic data association for semantic slam. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 1722–1729. IEEE, 2017. **1**
- [8] Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? In *International Conference on Learning Representations*, 2022. **4**
- [9] Ruojin Cai, Bharath Hariharan, Noah Snavely, and Hadar Averbuch-Elor. Extreme rotation estimation using dense correlation volumes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. **2**
- [10] Yu Chen, Ji Zhao, and Laurent Kneip. Hybrid rotation averaging: A fast and robust rotation averaging approach. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10353–10362, 2021. **2**
- [11] Marco Crocco, Cosimo Rubino, and Alessio Del Bue. Structure from motion with objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4141–4149, 2016. **1**
- [12] Michela Farenzena, Andrea Fusiello, and Riccardo Gherardi. Structure-and-motion pipeline on a hierarchical cluster tree. In *2009 IEEE 12th international conference on computer vision workshops, ICCV workshops*, pages 1489–1496. IEEE, 2009. **1**
- [13] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019. **8**
- [14] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. **2**
- [15] Dorian Gálvez-López, Marta Salas, Juan D Tardós, and José MM Montiel. Real-time monocular object slam. *Robotics and Autonomous Systems*, 75:435–449, 2016. **1**
- [16] Ben Glocker, Shahram Izadi, Jamie Shotton, and Antonio Criminisi. Real-time rgb-d camera relocalization. In *International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2013. **5**
- [17] Thomas A. Goldstein, Paul Hand, Choongbum Lee, Vladislav Voroninski, and Stefano Soatto. Shapefit and shapekick for robust, scalable structure from motion. *ArXiv*, abs/1608.02165, 2016. **3**
- [18] V.M. Govindu. Combining two-view constraints for motion estimation. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, pages II–II, 2001. **3**
- [19] R.I. Hartley. In defense of the eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(6):580–593, 1997. **2**
- [20] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. **2**
- [21] Nishant Jain, Suryansh Kumar, and Luc Van Gool. Robustifying the multi-scale representation of neural radiance fields. *arXiv preprint arXiv:2210.04233*, 2022. **3**
- [22] Kamil Kamiński, Jan Ludwiczak, Maciej Jasiński, Adriana Bukala, Rafal Madaj, Krzysztof Szczepaniak, and Stanisław Dunin-Horkawicz. Rossmann-toolbox: a deep learning-based protocol for the prediction and design of cofactor specificity in Rossmann fold proteins. *Briefings in Bioinformatics*, 23(1), 2021. **4**
- [23] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo. Learning open-world object proposals without learning to classify. *IEEE Robotics and Automation Letters*, 7(2):5453–5460, 2022. **5**
- [24] Seong Hun Lee and Javier Civera. Rotation-only bundle adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. **2**
- [25] Seong Hun Lee and Javier Civera. Hara: A hierarchical approach for robust rotation averaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15777–15786, 2022. **3**
- [26] Hongdong Li and Richard Hartley. Five-point motion estimation made easy. In *Proceedings of the 18th International Conference on Pattern Recognition - Volume 01*, page 630–633, USA, 2006. IEEE Computer Society. **2**
- [27] Heng Li, Zhaopeng Cui, Shuaicheng Liu, and Ping Tan. Rago: Recurrent graph optimizer for multiple rotation averaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15787–15796, 2022. **2, 3, 4, 8**
- [28] Qing Li, Jiasong Zhu, Rui Cao, Ke Sun, Jonathan M. Garibaldi, Qingquan Li, Bozhi Liu, and Guoping Qiu. Relative geometry-aware siamese neural network for 6dof camera relocalization. *Neurocomputing*, 426:134–146, 2021. **2**
- [29] John McCormac, Ronald Clark, Michael Bloesch, Andrew Davison, and Stefan Leutenegger. Fusion++: Volumetric object-level slam. In *2018 international conference on 3D vision (3DV)*, pages 32–41. IEEE, 2018. **1**
- [30] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. **2**

- [31] Dror Moran, Hodaya Koslowsky, Yoni Kasten, Haggai Maron, Meirav Galun, and Ronen Basri. Deep permutation equivariant structure from motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5976–5986, 2021. [2](#)
- [32] Pierre Moulon, Pascal Monasse, Romuald Perrot, and Renaud Marlet. OpenMVG: Open multiple view geometry. In *International Workshop on Reproducible Research in Pattern Recognition*, pages 60–74. Springer, 2016. [1](#), [2](#), [3](#)
- [33] Lachlan Nicholson, Michael Milford, and Niko Sünderhauf. Quadriclam: Dual quadrics from object detections as landmarks in object-oriented slam. *IEEE Robotics and Automation Letters*, 4(1):1–8, 2018. [1](#)
- [34] David. Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–770, 2004. [2](#)
- [35] Onur Özyesil and Amit Singer. Robust camera location estimation by convex programming. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2674–2683, 2014. [3](#)
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [8](#)
- [37] Pulak Purkait, Tat-Jun Chin, and Ian Reid. Neurora: Neural robust rotation averaging. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV 16*, pages 137–154. Springer, 2020. [2](#), [3](#), [6](#), [8](#)
- [38] Antoni Rosinol, Andrew Violette, Marcus Abate, Nathan Hughes, Yun Chang, Jingnan Shi, Arjun Gupta, and Luca Carlone. Kimera: From slam to spatial perception with 3d dynamic scene graphs. *The International Journal of Robotics Research*, 40(12-14):1510–1546, 2021. [1](#)
- [39] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. [5](#)
- [40] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. [1](#)
- [41] Kristy Sim and Richard Hartley. Recovering camera motion using ℓ_{∞} minimization. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, pages 1230–1237. IEEE, 2006. [3](#)
- [42] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. In *ACM siggraph 2006 papers*, pages 835–846. 2006. [1](#)
- [43] Yifan Sun and Qixing Huang. Pose synchronization under multiple pair-wise relative poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13072–13081, 2023. [2](#)
- [44] Matteo Taiana, Matteo Toso, Stuart James, and Alessio Del Bue. Posernet: Refining relative camera poses exploiting object detections. In *European Conference on Computer Vision*, pages 247–263. Springer, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [8](#)
- [45] Gk Tejus, Giacomo Zara, Paolo Rota, Andrea Fusiello, Elisa Ricci, and Federica Arrigoni. Rotation synchronization via deep matrix factorization. *ICRA*, 2023. [3](#)
- [46] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment—a modern synthesis. In *Vision Algorithms: Theory and Practice: International Workshop on Vision Algorithms Corfu, Greece, September 21–22, 1999 Proceedings*, pages 298–372. Springer, 2000. [1](#)
- [47] Kyle Wilson and Noah Snavely. Robust global translations with 1DSfM. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. [1](#), [2](#), [3](#)
- [48] Luwei Yang, Heng Li, Jamal Ahmed Rahim, Zhaopeng Cui, and Ping Tan. End-to-end rotation averaging with multi-source propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11774–11783, 2021. [2](#), [3](#)
- [49] Shichao Yang and Sebastian Scherer. Cubeslam: Monocular 3-d object slam. *IEEE Transactions on Robotics*, 35(4):925–938, 2019. [1](#)
- [50] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. iNeRF: Inverting neural radiance fields for pose estimation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. [2](#)
- [51] Zi Jian Yew and Gim Hee Lee. Learning iterative robust transformation synchronization. In *2021 International Conference on 3D Vision (3DV)*, pages 1206–1215. IEEE, 2021. [2](#)
- [52] Christopher Zach, Manfred Klopschitz, and Marc Pollefeys. Disambiguating visual relations using loop constraints. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1426–1433, 2010. [2](#)
- [53] Ganlin Zhang, Viktor Larsson, and Daniel Barath. Revisiting rotation averaging: Uncertainties and robust losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17215–17224, 2023. [3](#)
- [54] Bingbing Zhuang, Loong-Fah Cheong, and Gim Hee Lee. Baseline desensitizing in translation averaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4539–4547, 2018. [2](#), [3](#), [5](#), [8](#)