

PoserNet: Refining Relative Camera Poses Exploiting Object 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. The estimation of the camera poses associated with a set of images commonly relies on feature matches between the images. In contrast, we are the first to address this challenge by using objectness regions to guide the pose estimation problem rather than explicit semantic object detections. We propose Pose Refiner Network (PoserNet) a light-weight Graph Neural Network to refine the approximate pairwise relative camera poses. PoserNet exploits associations between the objectness regions - concisely expressed as bounding boxes - across multiple views to globally refine sparsely connected view graphs. We evaluate on the 7-Scenes dataset across varied sizes of graphs and show how this process can be beneficial to optimisation-based Motion Averaging algorithms improving the median error on the rotation by 62° with respect to the initial estimates obtained based on bounding boxes. Code and data are available at github.com/IIT-PAVIS/PoserNet.

1 Introduction

A common problem in computer vision is the recovery of multiple camera poses in a common reference frame starting from a set of images, with applications in tasks such as Structure from Motion (SfM), Simultaneous Localisation And Mapping (SLAM) or visual odometry. Traditionally, most solutions rely on the extraction and matching of keypoint features from the images [31]. Those approaches are vulnerable to many factors, such as: changes in viewpoint, illumination, repeated patterns that can lead to mismatches; and the presence of featureless, transparent or reflective surfaces which can result in a scarcity of useful keypoints. Since all these factors commonly occur in real-world scenes, feature-based approaches often rely on extensive refinement steps of the matches, and on outlier rejection methods like RANSAC [8].

We propose an alternative approach, instead of focusing on keypoints, we leverage on the continuous improvements in object detectors [15, 27] to extract the location of potential objects in the images as the basis of our approach. This presents multiple advantages: *i*) detections matching is more reliable; *ii*) working

^{*} This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 870743.

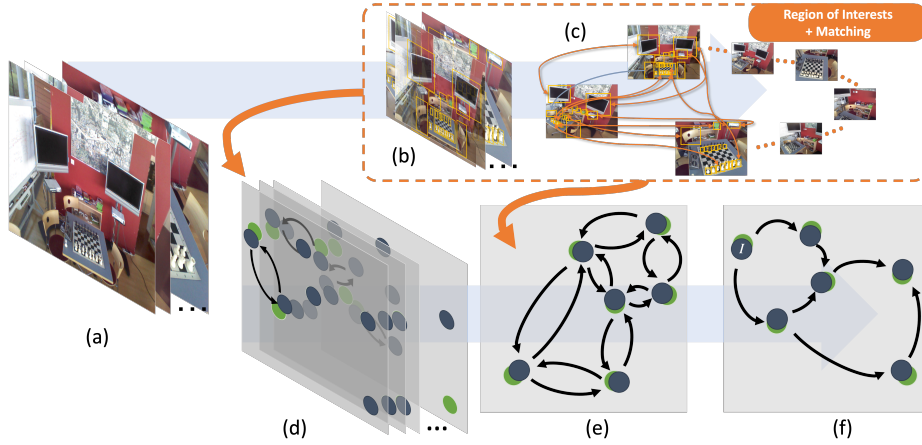


Fig. 1: (a) Our method takes a set of unordered images as input; (b) it detects ROIs; and (c) matches them between frames; (d) ROI & matches are then used to estimate an initial relative pose using 5-point algorithm; (e) the initial poses are refined exploiting the matched detections by PoserNet; and, finally, (f) the refined poses are processed by Motion Averaging to compute absolute poses.

with objects is more efficient than handling tens of thousands of keypoints; and *iii*) for many downstream applications it is convenient to reason in terms of objects present in the scene, rather than 3D points. The use of objects for camera pose estimation, *i.e.* Structure from Motion with Objects (SfMO), has been previously proposed [6, 12], however, it relies on a closed-form solution based on multiple disparate views of objects and a simpler camera model that makes its application highly constrained and challenging to apply in the wild.

To overcome the limitation of SfMO and test our object-based assumptions, we propose: *i*) to extract a set of 2D bounding boxes using an objectness detector indicating the probable location of objects of interest; *ii*) to match these detections across images; *iii*) to use them to provide an initial estimate of the relative transformation between each pair of images that contain common scene elements; and *iv*) to exploit the matched detections and rough relative poses to create a view graph, which we then refine via a novel GNN-based architecture called Pose Refinement Network (PoserNet). This allows us to obtain more accurate relative camera poses, which *v*) we then use to recover absolute poses via Motion Averaging algorithms and place camera poses in a common reference frame (as seen in Fig. 1).

To summarise, the contributions of this paper are threefold:

- A GNN-based architecture, PoserNet, that exploits matched region detections to refine rough estimates of the pair-wise camera transformations, producing accurate relative poses;

- we combine the PoserNet module with two paradigms of Motion Averaging algorithms (EIG-SE3 [2] & MultiReg [36]) to estimate absolute camera poses;
- we provide experimental evidence using object detections and keypoints, comparing our method against a state-of-the-art methods for small view-graphs on 7-Scenes.

2 Related work

Recovering camera poses from a sparse set of images is a well studied problem, especially in the context of SfM, and there are extensive reviews on the topic (see Bianco *et al.* [3]). Similarly, many works discussed in this paper, and our PoserNet model, are based on Graph Neural Networks (GNN). These have recently received a lot of attention, due to their successful applications in many domains, and we point to Wu *et al.* [32] and Zhou *et al.* [37] for an extensive review of GNNs. In this section, we will instead discuss: methods for relative camera pose estimation (Sec. 2.1), as these are the approaches PoserNet aims to improve on; Motion Averaging works (Sec. 2.2), which allow us to evaluate the effects of PoserNet pose refinements; and object-based approaches in 3D Computer Vision (Sec. 2.3), as PoserNet takes as inputs matched object detections.

2.1 Relative Pose Estimation

The most common approach to relative pose estimation is to extract a large number of keypoints, matched across the images, and to apply geometrical methods like Hartley’s normalised eight-point algorithm [14], Nistér’s five-point algorithm [25] or its simplified implementation introduced by Li and Hartley [17]. This problem is hard in real scenarios, since pixel locations and camera extrinsics can assume a wide range of values, and keypoint matches are noisy and unreliable in many practical application. For this reason, such geometrical methods are typically paired with tools like RANSAC to reduce the impact of outliers. The advancement of deep learning allowed approaching relative pose estimation as a learning problem: Li *et al.* [18] directly trained a deep network to predict the relative pose and orientation between two images, in the context of image relocalisation. Moran *et al.* [21] directly recovered global camera poses from feature matches and a sparse scene reconstruction, minimizing an unsupervised reprojection loss exploiting one encoder equivariant to feature permutations. Cai *et al.* [4] proposed an encoders-based network, with dense correlation volume representation, trained to estimate the relative pose from image pairs via cues like light source directions, vanishing points, and scene symmetries. Alternatively, Neural Radiance Fields for View Synthesis (NeRF) [20] was re-framed by Yen-Chen *et al.* [35] and used for camera pose estimation. Yan *et al.* [33], used a GNN to frame the problem of image retrieval for camera relocalisation as a node binary classification in a subgraph.

Unlike these methods, we suggest that the relative poses estimated via geometrical approaches - which are typically computationally efficient and do not

require training - can provide competitive performance, if properly refined. We exemplify this by showing how even rough estimates obtained by applying a standard geometrical method to a small number of loosely matched points (the centres of matched detections) can be refined by our PoserNet.

2.2 Motion Averaging

Given a set of relative poses, there are different approaches to reconstruct absolute poses in a common reference frame (*i.e.* the Motion Averaging or synchronisation problem). Classically, this is done via optimisation algorithms, like the ones of standard global SfM [22], or of Arrigoni *et al.* [2], who proposed a closed-form solution based on spectral decomposition, using an iteratively reweighted Least Squares scheme to reduce the impact of outliers. Recently, Arrigoni *et al.* [1] showed how cycle consistency allows establishing whether a view graph is solvable or not. Lee and Civera [16] proposed a rotation-only Bundle Adjustment (BA) framework, decoupling rotation estimation from translation and scene reconstruction, and Chen *et al.* [5] combined a global optimiser with fast view graph filtering and a local optimiser, improving on the traditional BA loss.

Also in the case of Motion Averaging, there has been a proliferation of deep models, such as NeuRoRA [26], which employed two GNNs to clean the relative camera poses before computing absolute poses. Yang *et al.* [34] introduced an end-to-end neural network for multiple rotation averaging and optimisation. Its view-graph is refined by fusing image context and features to predict outlier rotations, and initialising the rotation averaging process through a differentiable module, that allows outlier refinement at test time. Similarly, Yew and Lee [36] used a GNN to learn transformation synchronisation, but use an iterative approach in the tangent space where each step consists of a single weight-shared message passing layer that refines the absolute poses.

We test the effects of relative pose refinement via PoserNet on one optimisation-based [2] and one deep model [36], as both cases present interesting advantages. Classical methods do not require training, and therefore generalise better to different use cases, and are typically efficient, easy to deploy and well-studied; on the other hand, successfully combining PoserNet with a deep learning method would allow merging them in an end-to-end pipeline. While PoserNet shares some similarities with methods like [26], we focus on refining the relative poses before the Motion Averaging task, providing better inputs. Moreover, our method reasons about all relative poses of the graph while also exploiting bounding box information usually not used by models like [26] (*e.g.* the location of the object detections).

2.3 Object-Based Computer Vision

One of our key insight is the use of object locations instead of generic keypoints, as initial step of the camera pose reconstruction process. Other works have already suggested reasoning about the world in terms of object detections:

Rubino *et al.* [28] proposed a closed-form solution for estimating the 3D occupancy of an object, from multiview 2D object detections, and Gaudilliere *et al.* [9] used a similar representation to show how, in a scene represented as a collection of ellipsoids, a single ellipse-ellipsoid match and an approximate orientation from IMUs or vanishing point algorithms is enough to recover the camera location. They later [10] showed how, given two ellipse-ellipsoid correspondences, the camera rotation is recovered via an optimisation problem over the three Euler angles. The problem can then be further simplified [11], reducing it to an optimisation over a single angle by introducing constraints satisfied in many real world applications, *i.e.* that the images have null azimuth angle, and that the segment connecting the centres of the ellipsoids projects on the one connecting the centres of the detected ellipses.

In the context of object-level SLAM, notable contributions include Fusion++ [19], which combines Microsoft Fusion’s [23] TSDF and instance segmentation for indoor scene understanding; or SLAM++ [29], which implements real-time object detection, tracking and mapping in a dense RGB-D framework, but requires a database of well defined 3D models of the objects; or QuadricSLAM [24], which uses noisy odometry and object detections (as bounding boxes) to simultaneously estimate camera poses and objects locations, expressed as ellipsoids.

Unlike the aforementioned methods, we do not use a specific parameterisation of the objects shape (*e.g.* ellipsoid), and only rely on the approximate location of generic ROIs. Moreover, our method can be seen as complementary to theirs, as in principle we could take the bounding boxes associated with our ROIs, and the camera poses produced by combining PoserNet with a Motion Averaging approach, and use one of the approaches to localise the scene elements in 3D. Regarding SLAM methods, a direct comparison is not possible, as our approach is aimed at scenes with large camera baseline displacement and sparsely connected view-graphs, with minimal overlapping views.

3 Methodology

Our approach takes as input an unordered set of images $\{I_i\}$, $i \in [1, \dots, N]$ and their respective camera matrix intrinsics (K_i). We then aim to reconstruct the extrinsics $M_i = [R_i | t_i]$, with rotation $R_i \in SO(3)$ *i.e.* Special Orthogonal Group, and translation $t_i \in \mathbb{R}^3$ that map from the world reference frame to the camera reference frame (as shown in Figure 1). Firstly we detect ROIs within I_i , then we perform matching between ROIs and in turn solve for an initial estimate of the pairwise problem (Sec. 3.1). The ROIs represented as bounding boxes (BB) and the initial estimates are then passed to our PoserNet module to refine the relative poses (Sec. 3.2). Finally, the refined relative poses are processed by a rotation averaging method, which computes the absolute poses estimates of the cameras (Sec. 3.3).

3.1 Objectness, matching and initial poses

For each image, we compute a set of bounding boxes $bb_{i,a}$, $a \in [1, \dots, 50]$ which represent the top most confident ROIs candidates for objects. We opt for ROIs instead of the explicit object detections, as we make no assumption on semantic classes and they provide bounding boxes even in scenes with few objects. We use the *objectness*-based detector, Object Localisation Network (OLN) [15]. OLN is a generalisation of common detectors, substituting the classification layer (object semantic class) with a localisation confidence and trained in a weakly supervised manner relying on intersection over union and centreness.

Once we have a set of bounding boxes for each image, we then match them to obtain tracks of detections across the multi-view images. The matching of bounding boxes, or their image patches, between images is in a non-trivial problem, therefore we treat it as a black-box task. For simplicity, we solve the matching problem by applying the pre-trained SuperGlue model [30], which can be replaced by any matching approach. From the matches we construct a view graph, two images are considered connected if they share at least five matched ROIs. In addition, only images that are connected to at least another image are used in the graph, and the graph is built so that its connections form a chain that passes through all N images. This is necessary to ensure that the set of images and their matches constitute a solvable problem, in which no image view is completely disconnected from the others.

For each connected pair of images I_i, I_j , we estimate the relative transformation mapping from one image to the other, *i.e.* $\hat{R}_{i,j}$ and translation $\hat{t}_{i,j}$ that satisfy: $\hat{R}_{i,j} = R_j R_i^T$ and $\hat{t}_{i,j} = t_j - R_j^T t_i$. We compute the relative transformation using the 5-point algorithm and RANSAC. We experiment using both BB centres, which are sensitive to the quality of the BB detection, and SuperGlue keypoint matches within the region, which are sensitive to mismatches.

3.2 PoserNet: Graph-Based Relative Camera Pose Refinement

We define the PoserNet Graph Neural Network (GNN) as a special case of a multi-layer GNN, with shared edges representing the relative transformations and a partly shared node encoding. Therefore, we first review a general formulation of GNNs and then we describe our directed GNN with matched detections.

General Graph Neural Networks: The standard GNN formulation employs a graph $\mathcal{G} = \langle N, E \rangle$, representing a structure of nodes $N = \{N_1 \dots N_p\}$ and edges $E = \{E_1 \dots E_q\}$. Each N_i or E_i has an embedding (h_i^N, h_i^E) representing its latent information, which can be initialised based on *a priori* information (*e.g.* relative poses). The representations are then updated independently via a combination of message passing, nonlinear update functions ($\Psi^N(\cdot)$ and $\Psi^E(\cdot)$) and aggregation functions. Each update constitutes a step (k) of the graph. The message m^N between nodes from j to i , for step (k), is defined as:

$$m_{i,j,k}^N = \Psi^N(h_{i,k}^N, h_{j,k}^N, h_{i,j,k}^E). \quad (1)$$

To update N_i , a message is computed for all its neighbouring nodes $j \in \mathcal{N}_i$. Finally, the aggregation function averages the incoming messages to produce the latent information for node i for the $(k+1)^{th}$ iteration:

$$h_{i,k+1}^N \leftarrow \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} m_{i,j,k}^N. \quad (2)$$

Edges are updated in a similar fashion by applying a nonlinear transformation ($\Psi^E(\cdot)$) to information from the neighbour nodes and from their current latent space, *i.e.* $h_{i,j,k+1}^E \leftarrow m_{i,j,k}^E = \Psi^E(h_{i,k}^N, h_{j,k}^N, h_{i,j,k}^E)$. The formulation can be generalised to multi-layer graphs simply by connecting nodes across layers and having a custom (or repeating) \mathcal{G} topography for each layer.

PoserNet: The Pose Refinement network (PoserNet) takes as input a graph which contains the noisy relative camera poses obtained via pairwise geometry estimation (Sec. 3.1) to initialise its edges, and it outputs a graph with the same topology, but having refined estimates for the same relative poses. Therefore, PoserNet is a special case of the multi-layer GNN paradigm. At its simplest, it works on a graph representing an enriched version of the view graph of a scene, however, across layers of matched detections $l = \{l_1 \dots l_r\}$, with a shared node encoding (h^N), and with a shared edge representation (h^E). See Fig. 2 for a graphical representation of the structure of such a graph.

Specifically, PoserNet works on a graph in which nodes (N_i) represent camera poses, while directed edges ($E_{i,j}$) represent pairwise relative camera poses. Each node is associated with one of the input images, I_i . It contains the corresponding intrinsic camera parameters and it is associated with the list of bounding boxes detected on that image, arranged as layers. The embedding for a detection consists of the normalised information on the BB location, height, and width: $bb_{i,l} \in \mathbb{R}^4$. Detections from different nodes are connected pairwise when matched (*i.e.* they represent the same object in multi-view), where matched detections are represented as logical layers in the graph. Edge $E_{i,j}$ exists in the graph only if nodes N_i and N_j have at least one detection in common. For the message passing, the Ψ functions are implemented as a multi-layer perceptron (MLP).

Edge representation & update: The relative camera transformations are encoded in the edges as a translation vector plus a quaternion: $h_{i,j}^E \in \mathbb{R}^7$. We encode the transformation induced by traversing the edge in one, specified, direction, and we enforce that the transformation observed by traversing the edge in the opposite direction is the inverse of that transformation. Updating the edge representation requires creating one message per active layer ($l \in \mathcal{N}_l$), based on a pair of matched detections at its sides. Thus, for PoserNet the message is enriched in contrast to traditional GNNs, as it takes as additional input the representations of a pair of detections. A message is defined as:

$$m_{i,j,k,l}^E = \Psi^E(h_{i,k}^N, h_{j,k}^N, h_{i,j,k}^E, bb_{i,l}, bb_{j,l}), \quad (3)$$

and the updated representation for the edge is computed as the average of the messages for all layers: $h_{i,j,k+1}^E \leftarrow \frac{1}{|\mathcal{N}_l|} \sum_{l \in \mathcal{N}_l} m_{i,j,k,l}^E$.




Fig. 2: (a) The structure of graphs processed by PoserNet. Each image I_i is associated with a node N_i where the corresponding edges are relative transformations. Each set of corresponding detections (orange lines) between nodes refers to a layer of the graph for message passing (Detection track). (b) Nodes and (c) edges are updated by passing through $\Psi(\cdot)$ the node and edge information and respective bounding boxes (bb) for a given layer.

Node representation & update: The information contained in each node consists of image height and width, the camera focal length and the first coefficient of radial distortion, all normalised: $h_i^n \in \mathbb{R}^4$. For updating the latent representation of a node in PoserNet, each of its neighbours can send not just one, but a set of messages, each of which corresponds to a layer with a matched pair of detections. A message is defined as:

$$m_{i,j,k,l}^N = \Psi^N(h_{i,k}^N, h_{j,k}^N, h_{i,j,k}^E, bb_{i,l}, bb_{j,l}), \quad (4)$$

and the updated representation for the node is computed as the average of the messages from all of its neighbours ($j \in \mathcal{N}_i$), for all layers ($l \in \mathcal{N}_l$):

$$h_{i,j,k+1}^N \leftarrow \frac{1}{|\mathcal{N}_j|} \sum_{j \in \mathcal{N}_j} \frac{1}{|\mathcal{N}_l|} \sum_{l \in \mathcal{N}_l} m_{i,j,k,l}^N. \quad (5)$$

This message-passing and data processing scheme is repeated a fixed number of times (one of the hyperparameters of GNNs), and it finally produces the output edge embeddings which represent the refined estimates for inter-camera pose transformations. During this process, the embedding on the nodes and edges are updated, but the information relative to bounding boxes is kept constant.

Loss function: PoserNet is trained using graphs for which the relative transformations between camera poses are known. The loss function comprises four

components: the first and second component drive the network to produce accurate estimates for the relative poses:

$$\mathcal{L}_{orient} = \angle(q_{GT}^* \circ q_{est}), \quad (6)$$

$$\mathcal{L}_{tr_dir} = \angle(tr_{GT}, tr_{est}). \quad (7)$$

The orientation loss, \mathcal{L}_{orient} , encodes the angle between (\angle) the ground-truth and the estimated quaternion, where \circ represents quaternion composition, and q_{GT}^* is the conjugate of the ground-truth quaternion. The translation direction loss, \mathcal{L}_{tr_dir} , encodes the angle between ground-truth (tr_{GT}) and estimated translation vectors. The remaining two components are the quaternion normalisation loss, $\mathcal{L}_{q||}$, and the translation normalisation loss, $\mathcal{L}_{tr||}$. They push the network towards producing unit norm outputs:

$$\mathcal{L}_{q||} = ||q_{est}|| - 1, \quad (8)$$

$$\mathcal{L}_{tr||} = ||tr_{est}|| - 1. \quad (9)$$

The total loss function of PoserNet is defined as:

$$\mathcal{L}_{PoserNet} = \mathcal{L}_{orient} + \mathcal{L}_{tr_dir} + \alpha(\mathcal{L}_{q||} + \mathcal{L}_{tr||}), \quad (10)$$

where α is a coefficient used to tune the strength of the contribution of the different components of the loss.

3.3 Absolute Pose Estimation

The relative poses refined by PoserNet can be used to compute absolute camera poses; for this, we look at State-of-the-Art methods for two common approaches: optimisation and deep learning. With the former, we want to show how modern convolution techniques can improve “classical” approaches, while combining PoserNet with a deep-learning model could lead to an end-to-end pipeline.

EIG-SE3 [2]: This method solves the Motion Averaging problem in closed-form using spectral decomposition. Given a fully connected graph with n nodes, and accurate estimates of the relative and absolute transformations between the cameras expressed as:

$$M_{i,j} = \begin{pmatrix} R_{i,j} & t_{i,j} \\ 0 & 1 \end{pmatrix}, M_i = \begin{pmatrix} R_i & t_i \\ 0 & 1 \end{pmatrix}, \quad (11)$$

with $M_{i,j} \in SE(3)$, $SE(3)$ the Special Euclidean Group (3), the absolute camera poses $M_i \in SE(3)$, with $M_{i,j} = M_i M_j^{-1}$, can be recovered solving the problem $\mathbf{X}\mathbf{M} = n\mathbf{M}$, where \mathbf{M} is the $4n \times 4$ matrix obtained concatenating all matrices M_i and \mathbf{X} is a block matrix with block $i, j = M_{i,j}$.

If the relative transformations are exact, this problem can be solved in closed form in two steps: *i*) the columns of M are four eigenvectors of \mathbf{X} associated

with the eigenvalue n , found as the 4-dimensional basis of the null-space of $L = (n\mathbf{I}_{4n} - \mathbf{X})$; *ii*) the basis found at the previous point is not guaranteed to be composed of Euclidean motions, *i.e.* it is not guaranteed that each 4×4 block belongs to $\text{SE}(3)$. This is addressed by extracting every fourth row of M and finding the change of basis that maps them to $[0, 0, 0, 1]$.

In the original work, the authors show how this approach can be generalised to problems with noisy relative transformation, by solving the Least Squares problem $\min_{M \in \text{SE}(3)} \|LM\|_F^2$. They also extend the approach to work on graphs that are not fully connected - *i.e.* some of the off-diagonal blocks of \mathbf{X} are zero - and show how the optimisation problem can be included in an iteratively reweighted Least Squares scheme to reduce the impact of outliers, which are a common occurrence in most camera pose estimation applications. See [2] for a full description of the optimisation algorithm.

MultiReg: computes incremental pose updates within tangent space to ensure that poses remain on the $\text{SO}(3)$ or $\text{SE}(3)$ manifold. Like PoserNet, a GNN is used where, however, nodes represent camera poses and the edges represent transformations, $\mathcal{G}_a = \langle N_a, E_a \rangle$. In addition, only the absolute pose and a latent representation is updated on a node, with no edge update. To achieve this each increment ϵ_i is computed from the pose residuals $\Gamma_{i,j} = T_i T_j^{-1} \hat{T}_{i,j}^{-1}$ based on the current absolute pose and the latent encoding (h^N) stored in the node. Therefore each update is equivalent to:

$$(\epsilon_i, \Delta h_i^N) \leftarrow \Psi^N(h_i^N, u, \Psi^E(h_i^N, h_j^N \Gamma_{i,j})), \quad (12)$$

where Ψ^N and Ψ^E are MLPs and u is a global encoding from average pooling of f over the graph. The transform update is therefore $T_i \leftarrow \exp(\hat{\epsilon}_i) T_i$ which corresponds to addition in Euclidean space for Lie Groups. Finally, similar to pair-wise we compute the loss based on the decomposed rotation and translation.

$$\mathcal{L}_{abs} = \frac{1}{|\epsilon^c|} \sum_{(i,j) \in \epsilon^c} (|R_{i,j} - R_{i,j}^{GT}| + \lambda_t |t_{i,j} - t_{i,j}^{GT}|). \quad (13)$$

4 Experiments

In this section, we provide an evaluation of the proposed PoserNet model. We validate our choice of PoserNet structure (*i.e.* depth, choice of training data), and show how PoserNet can improve the accuracy of relative pose estimation. We then investigate the effects of this refinement on Motion Averaging. As an implementation detail, PoserNet partially relies on PyTorch Geometric (PyG) [7].

For each experiment we report separately the error on the camera rotation and the translation error; the former is expressed in degrees, while the translation is reported in meters for absolute poses and in degrees for relative poses. This is due to the fact that relative poses are defined up to a scale, and we therefore report the angular distance between the ground-truth and the estimated translation unit vector. These results are expressed in form of the median

Dataset part	Graphs	Nodes	Edges	Nodes/graph	Edges/graph	Edges/node
Small graphs, train	14000	112000	215790	8	15.4	1.9
Small graphs, test	7000	56000	110367	8	15.8	2.0
Large graphs, train	5	625	22674	125	4534.8	36.3
Large graphs, test	5	625	22606	125	4521.2	36.2

Table 1: Composition of the small and large graph datasets used to validate the proposed method. The different graph-building strategies result in large graphs which have a higher (18-fold higher) connectivity than the small graphs.

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 metrics, and 0.05, 0.1, 0.25, 0.5 and 0.75 for errors expressed in meters. For all experiments that required training network, we trained three separate models and report their average performance, with their standard deviation.

4.1 Dataset

For our experiments, we use graphs built from the 7-Scenes [13]: this is a collection of 41k images acquired in seven small-sized indoor environments with a Microsoft Kinect. The RGB-D information has been used to compute accurate ground-truth camera poses. For each scene, several sequences were recorded by different users, and they are split into distinct training and testing sequence sets. Starting from 7-Scenes data, we created two datasets comprising graphs of different sizes: a small graphs dataset and a larger graph dataset. In each case, we followed the split between training and test data or the original work.

Small graphs dataset: It contains 8-node graphs. For each scene we created 2k graphs from the training sequences, and 1k from the test ones. Each graph was created by randomly selecting eight images that satisfy the connectivity constraints outlined in Section 3.1. We thus ensure that all graph nodes are connected, while not requiring the graph to be fully connected.

Large graphs dataset: The graphs are generated sampling, for each scene, 125 frames from the aggregated training or testing sequences. In this case, we do not explicitly enforce a connectivity constraint, because the limited size of the scenes and the large number of nodes guarantee a large number of connected edges. We extracted only one training and one testing graphs from the first five scenes, due to the difficulty in identifying large graphs.

4.2 PoserNet Refining Relative Poses

We present an analysis designed to assess the effectiveness of PoserNet in refining relative camera poses. We first show how PoserNet successfully improves relative poses provided by different initialisation algorithms, and secondly how PoserNet generalises well to graphs of different sizes. In Supplementary Material we discuss

the impact of choosing different depths for PoserNet’s GNN, we identify a depth of 2 as the best and we use it throughout the remaining experiments.

Refining Relative Poses from Different Initialisation Modes: We assessed the effectiveness of PoserNet in refining the relative pose estimates initialised from bounding box centres (BB) and keypoints (KP), as discussed in Section 3.1. We evaluate the PoserNet model on the small graphs dataset as shown in Table 2. In both cases we see a significant performance gain, with median rotation error reduction of over 76° and 28° for the BB and KP cases respectively, and an improvement of over 42° and 72° for the translation median angular error. This confirms PoserNet’s effectiveness at refining relative poses generated by different algorithms, with different levels of noise.

Dealing with graphs of different sizes: We explore the impact of using the large and small graphs to train and test PoserNet. The results are summarised in Table 3. For each test set, we provide the baseline errors of the unrefined graphs, and the performances of models trained on either the large graph or small graph dataset. From the results, we make three observations. First, we verify that PoserNet is effective at refining relative poses: models trained on either small or large graphs and tested on the corresponding test set achieve a significant accuracy improvement over the baseline. Second, we see that PoserNet has good generalisation capabilities: the pose refinement performance generalises quite well when training on one dataset, *e.g.* small graphs, and testing on large graphs. Finally, we observe that models trained on small graphs perform better, while testing on large graphs leads to more accurate estimates.

Generalisation to unseen graphs: To test the ability of our method to generalise to novel scenes, we evaluate PoserNet’s performance on relative pose refinement for 8-node graphs using a leave-one-out scheme. This is done evaluating on each of the seven scenes of the dataset a model trained only on the other six scenes. The results of this test, reported in the Supplementary Material, show median orientation and translation orientation errors of 8.04° and 19.75° respectively, comparable with the performance on the full dataset (7.31° and 14.54°).

4.3 Absolute Pose Estimates

The last experiments show the impact of PoserNet on the motion averaging task (*i.e.* extract absolute camera poses from the relative ones). We consider two possible approaches: a modern, graph-network based model (MultiReg); and an optimisation-based algorithm (EIG-SE3). Results for both Motion Averaging approaches are reported in Table 4, using the same notation of previous experiments. However, since the Motion Averaging process reintroduces the scale of the translation vector, the error on the translations is reported as the Euclidean distance between the ground-truth and estimated vectors, expressed in meters.

Input		Orientation error (deg)						Translation direction error (deg)					
		3	5	10	30	45	η	3	5	10	30	45	η
BB	Initial	0.00	0.00	0.00	0.16	1.44	96.48	0.00	0.00	0.00	0.00	0.00	89.30
	PoserNet	0.00	0.00	0.74	88.18	99.02	20.39	0.00	0.00	0.00	7.92	45.63	46.60
		± 0.00	± 0.00	± 0.14	± 0.28	± 0.02	± 0.20	± 0.00	± 0.00	± 0.00	± 0.25	± 0.43	± 0.14
KP	Initial	0.74	4.01	11.16	41.46	62.53	36.26	0.00	0.03	0.10	1.37	5.49	87.23
	PoserNet	1.13	17.24	78.19	98.78	99.86	7.31	0.01	1.42	27.67	80.07	90.01	14.54
		± 0.11	± 0.86	± 1.36	± 0.02	± 0.01	± 0.11	± 0.01	± 0.28	± 0.70	± 0.31	± 0.68	± 0.14

Table 2: PoserNet performances in refining relative pose computed from the centres of matched bounding boxes (BB) or from a standard keypoints approach (KP). For both cases, the initial orientation and translation direction error are improved by PoserNet processing.

Train	Test	Orientation error (deg)						Translation direction error (deg)					
		3	5	10	30	45	η	3	5	10	30	45	η
Raw	Small	0.74	4.01	11.16	41.46	62.53	36.26	0.00	0.03	0.10	1.37	5.49	87.23
Small	Small	1.13	17.24	78.19	98.78	99.86	7.31	0.01	1.42	27.67	80.07	90.01	14.54
		± 0.11	± 0.86	± 1.36	± 0.02	± 0.01	± 0.11	± 0.01	± 0.28	± 0.70	± 0.31	± 0.68	± 0.14
Large	Small	0.00	0.04	13.26	96.98	99.61	13.97	0.00	0.00	0.11	57.19	81.26	27.56
		± 0.00	± 0.03	± 5.55	± 0.37	± 0.09	± 0.94	± 0.00	± 0.00	± 0.10	± 5.72	± 1.16	± 2.02
Raw	Large	0.00	0.00	0.00	20.00	60.00	37.70	0.00	0.00	0.00	0.00	0.00	79.60
Small	Large	0.00	0.00	80.0	100.00	100.00	6.85	0.00	0.00	20.00	100.00	100.00	12.62
		± 0.00	± 0.00	± 0.00	± 0.00	± 0.00	± 0.11	± 0.00	± 0.00	± 0.00	± 0.00	± 0.00	± 0.48
Large	Large	0.00	0.00	13.33	100.00	100.00	12.11	0.00	0.00	0.00	100.00	100.00	21.58
		± 0.00	± 0.00	± 9.43	± 0.00	± 0.00	± 1.10	± 0.00	± 0.00	± 0.00	± 0.00	± 0.00	± 1.36

Table 3: PoserNet performances when trained and evaluated on small (Small) or large (Large) graphs. While versions of PoserNet trained on either largely outperform the baseline error (Raw), models trained on small give the best performances on both testing datasets. (Highlighted in bold).

Model	Init	Orientation error (deg)						Translation error (meters)					
		3	5	10	30	45	η	0.05	0.1	0.25	0.5	0.75	η
MultiReg	BB	0.60	2.10	8.10	47.60	69.13	31.37	0.50	1.47	6.93	23.03	40.80	0.89
		± 0.00	± 0.00	± 0.10	± 0.53	± 0.47	± 0.30	± 0.00	± 0.06	± 0.12	± 0.70	± 0.85	± 0.01
	KP	10.23	24.23	39.7	67.97	79.90	15.69	1.00	3.77	17.77	40.30	57.23	0.64
		± 0.12	± 0.25	± 0.40	± 0.06	± 0.00	± 0.08	± 0.00	± 0.06	± 0.12	± 0.17	± 0.06	± 0.01
MultiReg + PoserNet	BB	0.50	1.60	6.10	41.03	62.50	35.60	0.60	1.50	6.53	21.60	39.33	0.91
		± 0.00	± 0.00	± 0.10	± 0.42	± 0.14	± 0.33	± 0.00	± 0.00	± 0.06	± 0.00	± 0.06	± 0.00
	KP	0.50	1.63	6.20	40.90	62.60	35.73	0.60	1.50	6.40	21.43	39.23	0.91
		± 0.00	± 0.06	± 0.10	± 0.44	± 0.53	± 0.37	± 0.00	± 0.00	± 0.10	± 0.12	± 0.12	± 0.00
EIG-SE3	BB	0.00	0.00	0.09	2.46	7.84	78.46	0.00	0.00	0.00	0.57	13.47	0.65
	KP	11.99	23.87	34.49	57.59	62.09	25.37	0.00	0.06	2.63	31.39	74.40	0.61
EIG-SE3 + PoserNet	BB	0.00	0.09	10.36	88.00	94.06	15.80	0.00	0.00	7.54	46.80	83.23	0.52
	KP	15.56	48.36	82.14	93.14	5.69	5.10	0.04	1.94	34.51	76.43	94.49	0.32

Table 4: Motion averaging performances on the 7-Scenes dataset using the MultiReg and EIG-SE3 algorithms. We report results for different initialisation methods: bounding box centres (BB) and keypoints (KP).

Motion averaging via MultiReg: In the case of MultiReg, the refinement of the relative poses via PoserNet does not improve the performance, leading instead to larger errors. Specifically, refining the poses makes the orientation and translation errors respectively 4.2° and 0.02 m worse in the case of BB initialisation, and 20° and 0.27 m worse in the case of KP initialisation.

While this performance is disappointing, it is not completely unexpected. Firstly, the original MultiReg model was tested on large, very connected graphs, which means its structure and parameters might not be well-suited for our small graph dataset. On the other hand, the MultiReg model is too complex to be trained only on our limited large graph dataset. Moreover, [36] discarded as outliers all edges initialised with a rotation error larger than 15° and a translation error of more than 15 cm, which makes our initial error condition very challenging in comparison.

Motion averaging via EIG-SE3: Finally, we show the performance gain obtained by combining PoserNet with a closed-form approach to solve Motion Averaging, the EIG-SE3. Unlike MultiReg, this approach does not require training, making it easier to deploy and more general. Given the deterministic nature of this model, we report results of only one evaluation run.

Results of this experiment are reported in the bottom two rows of Table 4; in this case, PoserNet results in improvement over the baseline error for both the *BB* and *KP* initialisation modes. In the *BB* case, the median error on the rotation improves by 62° , and in the *KP* case the orientation error is reduced by 20° . Performance on the translation are also clearly noticeable, with an improvement of 0.13 m and 0.29 m for the BB and KP initialisation respectively.

5 Conclusions

We have proposed the novel PoserNet module for relative pose refinement driven by matched detections, based on ROIs generated from objectness-based detectors. We have shown that such ROIs are sufficient to recover accurate camera poses, with a large improvement with respect to the initial pair-wise estimates. In addition, we have shown how PoserNet can improve the outcome of the optimisation-based Motion Averaging EIG-SE3. The proposed model is a relatively simple and fast network relying on sparse object detections (in contrast to keypoints) and relatively lightweight bounding box representation. Therefore, the combination with a ‘classical’ geometric approach, such as EIG-SE3, yields increased performance compared to the more complex state-of-the-art networks like MultiReg. While the lack of improvement in combination with MultiReg highlights a challenge of deep learning methods being able to generalise to other scenarios outside of their initial design. We point to this limitation as a direction for future work, which could result in a generalised Motion Averaging method as flexible as optimisation methods and able to incorporate specific elements such as PoserNet for task or domain-specific improvement.

References

1. 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, October 2021.
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.
3. Simone Bianco, Gianluigi Ciocca, and Davide Marelli. Evaluating the performance of structure from motion pipelines. *Journal of Imaging*, 4(8), 2018.
4. 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.
5. 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.
6. 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.
7. 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. 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.
9. V. Gaudilliere, G. Simon, and M. Berger. Camera relocalization with ellipsoidal abstraction of objects. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 8–18, Los Alamitos, CA, USA, oct 2019. IEEE Computer Society.
10. Vincent Gaudillière, Gilles Simon, and Marie-Odile Berger. Camera pose estimation with semantic 3d model. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4569–4576, 2019.
11. Vincent Gaudillière, Gilles Simon, and Marie-Odile Berger. Perspective-2-ellipsoid: Bridging the gap between object detections and 6-dof camera pose. *IEEE Robotics and Automation Letters*, 5(4):5189–5196, 2020.
12. Paul Gay, Cosimo Rubino, Vaibhav Bansal, and Alessio Del Bue. Probabilistic structure from motion with objects (psfmo). In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3075–3084, 2017.
13. 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, October 2013.
14. R.I. Hartley. In defense of the eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(6):580–593, 1997.
15. 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 (RA-L)*, 2022.
16. Seong Hun Lee and Javier Civera. Rotation-only bundle adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
17. Hongdong Li and Richard Hartley. Five-point motion estimation made easy. In *Proceedings of the 18th International Conference on Pattern Recognition - Volume 01*, ICPR ’06, page 630–633, USA, 2006. IEEE Computer Society.

18. 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.
19. John McCormac, Ronald Clark, Michael Bloesch, Andrew J. Davison, and Stefan Leutenegger. Fusion++: Volumetric object-level slam. 2018.
20. 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.
21. Dror Moran, Hodaya Koslowsky, Yoni Kasten, Haggai Maron, Meirav Galun, and Ronen Basri. Deep permutation equivariant structure from motion, 2021.
22. 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.
23. Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011.
24. Lachlan Nicholson, Michael Milford, and Niko Sunderhauf. Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam. *IEEE Robotics and Automation Letters*, PP:1–1, 08 2018.
25. D. Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–770, 2004.
26. Pulak Purkait, Tat-Jun Chin, and Ian Reid. Neurora: Neural robust rotation averaging. In *European Conference on Computer Vision*, pages 137–154. Springer, 2020.
27. Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
28. Cosimo Rubino, Marco Crocco, and Alessio Del Bue. 3d object localization from multi-view image detections. In *Pattern Analysis and Machine Intelligence (TPAMI), 2017 IEEE Transactions on*, 2017.
29. Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H.J. Kelly, and Andrew J. Davison. Slam++: Simultaneous localisation and mapping at the level of objects. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
30. Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020.
31. Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
32. Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, C. Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32:4–24, 2021.
33. Shen Yan, Yang Pen, Shiming Lai, Yu Liu, and Maojun Zhang. Image retrieval for structure-from-motion via graph convolutional network. *CoRR*, abs/2009.08049, 2020.
34. 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 (CVPR)*, pages 11774–11783, June 2021.

35. 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.
36. Zi Jian Yew and Gim Hee Lee. Learning iterative robust transformation synchronization. In *International Conference on 3D Vision (3DV)*, 2021.
37. Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.