

E-Graph: Minimal Solution for Rigid Rotation with Extensibility Graphs

Yanyan Li^{1,2}[0000–0001–7292–9175] and Federico Tombari^{1,3}[0000–0001–5598–5212]

¹ Technical University of Munich, Munich, Germany

² Meta-Bounds Tech, Shenzhen, China

³ Google, Zurich, Switzerland

yanyan.li@tum.de, tombari@in.tum.de

Abstract. Minimal solutions for relative rotation and translation estimation tasks have been explored in different scenarios, typically relying on the so-called co-visibility graphs. However, how to build direct rotation relationships between two frames without overlap is still an open topic, which, if solved, could greatly improve the accuracy of visual odometry. In this paper, a new minimal solution is proposed to solve relative rotation estimation between two images without overlapping areas by exploiting a new graph structure, which we call Extensibility Graph (E-Graph). Differently from a co-visibility graph, high-level landmarks, including vanishing directions and plane normals, are stored in our E-Graph, which are geometrically extensible. Based on E-Graph, the rotation estimation problem becomes simpler and more elegant, as it can deal with pure rotational motion and requires fewer assumptions, e.g. Manhattan/Atlanta World, planar/vertical motion. Finally, we embed our rotation estimation strategy into a complete camera tracking and mapping system which obtains 6-DoF camera poses and a dense 3D mesh model. Extensive experiments on public benchmarks demonstrate that the proposed method achieves state-of-the-art tracking performance.

1 Introduction

Camera pose estimation is a long-standing problem in computer vision as a key step in algorithms for visual odometry, Simultaneous Localization and Mapping (SLAM) and related applications in robotics, augmented reality, autonomous driving (to name a few). As part of the camera pose estimation problem, the minimal case [40] provides an estimate of whether the problem can be solved and how many elements are required to obtain a reliable estimate. According to the input data type and scenarios, different solutions [1, 28, 15, 9] were proposed, most of which became very popular in the computer vision and robotic community, such as the seven-point [1] and five-point [28] approaches. A typical limitation of traditional pose estimation solutions based on the minimal case [1, 28, 31, 9] is that both rotation and translation estimation rely on the co-visibility features between two frames, this having as a consequence that the length of an edge between two nodes is often relatively short. Therefore, tracking errors tend

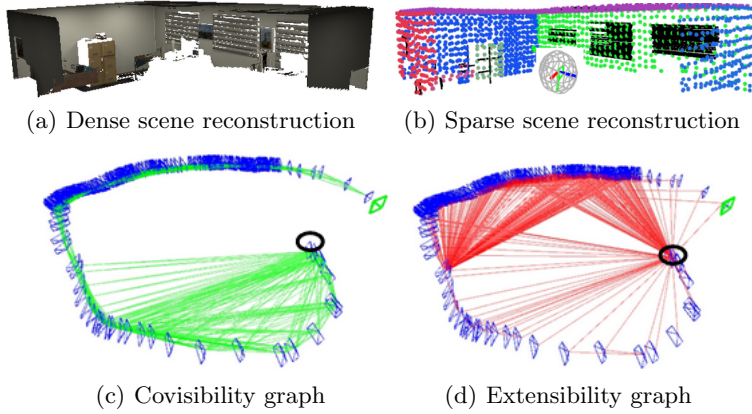


Fig. 1. Dense (a) and sparse (b) scene reconstruction of the office-room scene from the ICL dataset [10] obtained by the proposed method. (c) and (d): keyframes (in blue) and connected frames are linked with green and red lines, respectively, to build up the proposed covisibility and extensibility graphs. The black ellipses denote the start points of the camera trajectory.

to accumulate easily based on a frame-to-frame or frame-to-keyframe strategy. To solve this issue, more advanced tracking systems [26, 3] with optimization solutions, including local and global bundle adjustment approaches, were exploited to refine poses from minimal solutions. Loop Closure is a common algorithm used in feature-based [24] and direct [7] methods to remove drift. However, it also requires the camera to revisit the same place, which is a limiting assumption in many scenarios.

Compared with point features, lines and planes require more computation to be extracted and described. Early multi-feature SLAM systems [8] use them to increase the number of features to combat low-textured scenes. After that, coplanar, parallel and perpendicular relationships were explored [39, 18, 20] to add more constraints in the optimization module, still following a similar tracking strategy as ORBSLAM [25] or DSO [36] for the initial pose estimation.

Different to the tightly coupled estimation strategy, some works [43] proposed to decouple the 6-DoF pose estimation into rotation and translation estimation aiming to achieve a more accurate rotation estimation, based on the idea that pose drift is mainly caused by the rotation component [14]. At the same time, based on an estimated rotation matrix [31], only two points are required to compute the translation motion, leading to more robustness in low-textured regions.

The Manhattan World (MW) [43] and Atlanta World (AW) [13] assumptions introduce stronger constraints since they require a single orthogonal scene, or a scene with a unified vertical direction. Unlike loop closure that removes drift by detecting trajectory loops, the assumption of MW and AW is introduced for indoor tracking scenarios [19, 14] to improve the accuracy of camera

pose estimation, since most indoor artificial environments follow this assumption. MW and AW improve accuracy when the main structure of the scene has orthogonal elements. However, since this assumption requires the observation of vertical/orthogonal environmental features (such as straight lines or planes), the SLAM system using this method is also limited in the types of scenarios it can be successfully applied to.

In this paper we propose a rigid rotation estimation approach based on a novel graph structure, which we dub Extensibility Graph (E-Graph), for landmark association in RGB-D data. Our approach is designed to reduce drift and improve the overall trajectory accuracy in spite of loop closure or MW/AW assumptions. Benefiting of E-Graph, the drift-free rotation estimation problem is simplified to the alignment problem of rotating coordinate systems. Importantly, our rotation step does not need overlaps between two frames by making use of vanishing directions of lines and plane normals in the scene, hence can relate a higher number of keyframes with respect to standard co-visibility graphs, with benefits in terms of accuracy and robustness in presence of pure rotational motions.

In addition, we develop a complete tracking and dense mapping system based on the proposed E-Graph and rotation estimation strategies, which we demonstrate to outperform state-of-the-art SLAM approaches [20, 38, 26, 3]. To summarize, the main contributions of this paper are as follows: i) a new perspective for reducing drift is proposed based on our novel graph structure, E-Graph, which connects keyframes across long distances; ii) a novel drift-free rotation alignment solution between two frames without overlapping areas based on E-Graph; iii) a complete SLAM system based on the two previous contributions to improve robustness and accuracy in pose estimation and mapping. The proposed approach is evaluated on common benchmarks such as ICL [10] and TUM-RGBD [33], demonstrating an improved performance compared to the state of the art.

2 Related work

By making the assumption of planar motion [9], two-view relative pose estimation is implemented based on a single affine correspondence. Point features are common geometric features used in VO and SLAM [3] systems. To remove the drift from point-based front ends, different types of back ends are explored in tracking methods. Loop closing is an important module to remove drift, which happens when the system recognizes that a place [6, 23] has been visited before. After closing the loop, associated keyframes in the covisibility graph will be adjusted. Benefiting of loop closure and optimization modules, ORB-SLAM series [26, 3] organize the keyframes efficiently, which provides robust support for tracking tasks. Different from sparse point features used in ORB-SLAM, BAD-SLAM [32] implements a direct bundle adjustment formulation supported by GPU processing.

However, in indoor environments, to cover texture-less regions that have few point features, more geometric features are merged into the front end of systems. At the early stage, methods build re-projection error functions for lines

and planes. CPA-SLAM [22] makes use of photometric and plane re-projection terms to estimate the camera pose. Based on estimated camera poses, detected planes are merged together with a global plane model. Similar to our method, CPA-SLAM and KDP-SLAM [11] can build constraints between non-overlapping frames. However those constraints are used to build heavy optimization targets instead of improving the efficiency. Furthermore, the relationship between parallel lines (vanishing points) and perpendicular planes is explored in [17, 41]. Based on the regularities of those structural features, they obtain a more accurate performance. Instead of exploring the parallel/perpendicular relationships between lines/planes, [30, 18] make use of constraints between co-planar points and lines in the optimization module.

Those regularities aim to build constraints between local features, [14, 20] introduce global constraints by modeling the environment as a special shape, like MW and AW. The MW assumption is suitable for a cuboid scenario, which is supposed to be built by orthogonal elements. Based on this assumption, those methods estimate each frame’s rotation between the frame and the Manhattan world directly, which is useful to avoid drift between frames in those scenes. L-SLAM [14] groups normal vectors of each pixel into an orthogonal coordinate by projecting them into a Gaussian Sphere [43] and tracks the coordinate axes to compute the relative rotation motion. Similar to the main idea of L-SLAM, [15] provides a RGB-D compass by using a single line and plane. Since the line lies on the plane, the underlying assumption of the system is the MW-based rotation estimation method. However, the limitation of this strategy is also very obvious, that it works only in Manhattan environments. Based on ORB-SLAM2 [26], Structure-SLAM [19, 20] merges the MW assumption with keyframe-based tracking, to improve the robustness of the system in non-MW indoor scenes, which refine decoupled camera pose by using a frame-to-model strategy. Compared with MW-based tracking methods, our approach is less sensitive to the structure of environments.

3 Minimal case in orientation estimation

Commonly, the 6-DoF Euclidean Transform $T \in SE(3)$ defines motions as a set of rotation $R \in SO(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$. Based on point correspondences, camera pose estimation can be defined as,

$$\mathbf{P}' = R\mathbf{P} + \mathbf{t} \quad (1)$$

where \mathbf{P}' and \mathbf{P} are 3D correspondences, and $[R, \mathbf{t}]$ defines the relative motion between two cameras. For monocular sensors, their image normalized representations are \mathbf{X}'_c and \mathbf{X}_c ,

$$\mathbf{X}'_c = \alpha(R\mathbf{X}_c + \gamma\mathbf{t}) \quad (2)$$

where α and γ are depth-related parameters. After multiplying (2) by $\mathbf{X}_c'^T[\mathbf{t}]_x$, we can obtain the classic essential matrix equation,

$$\mathbf{X}_c'^T E \mathbf{X}_c = 0 \quad (3)$$

where $E = [\mathbf{t}]_x R$ and $[\mathbf{t}]_x$ is the skew symmetric matrix formed by \mathbf{t} .

For RGB-D sensors, the task is simplified since the absolute depth information is directly provided by sensors. Equation (3) can be solved by using 3 non-collinear correspondences only [29], although the distance between two frames is supposed to be kept small to extract enough correspondences.

3.1 Minimal solution for rotation




Fig. 2. Minimal case of rotation estimation in EG

Different from traditional methods based on co-visibility graphs, the proposed method decouples rotation and translation estimation into two separate stages. Moreover, the rotation estimation task does not require feature correspondences. As shown in Figure 2, non-parallel direction vectors $\mathbf{v}_m, m \in [0, 1, \dots, n]$ are detected in the camera coordinate C_j , where $\mathbf{v}_m^j = [x v_{m,y}^j v_{m,z}^j v_m^j]^T$. In Euclidean 3D space, the size of a finite and linearly independent set of vectors is less than four. According to the Gram-Schmidt orthogonalization process, we can obtain an orthogonal set $S = [\mathbf{u}_0, \mathbf{u}_1, \mathbf{u}_2]$,

$$\begin{aligned} \mathbf{u}_0 &= \mathbf{v}_0^j \\ \mathbf{u}_1 &= \mathbf{v}_1^j - \text{proj}_{[\mathbf{v}_0^j]}(\mathbf{v}_1^j) \\ \mathbf{u}_2 &= \mathbf{v}_2^j - \text{proj}_{[\mathbf{v}_0^j]}(\mathbf{v}_2^j) - \text{proj}_{[\mathbf{v}_1^j]}(\mathbf{v}_2^j) \end{aligned} \quad (4)$$

by using the projection operator $\text{proj}_{[\mathbf{u}]}(\mathbf{v}) = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\langle \mathbf{u}, \mathbf{u} \rangle} \mathbf{u}$, where $\langle \mathbf{u}, \mathbf{v} \rangle$ shows the inner product of the vectors \mathbf{u} and \mathbf{v} . Furthermore, we obtain the normalized vectors $\mathbf{e}_0, \mathbf{e}_1$ and \mathbf{e}_2 via $\mathbf{e}_m = \frac{\mathbf{u}_m}{\|\mathbf{u}_m\|}$.

For the Euclidean space \mathbb{R}^3 , the relevant orthonormal basis set based on the detected direction vectors is $(\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2)$. In the j^{th} camera coordinate, the orthonormal set is detected as $(\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2)$, while $(\mathbf{e}_0^*, \mathbf{e}_1^*, \mathbf{e}_2^*)$ in the k^{th} camera coordinate.

Therefore, from the perspective of the orthonormal set, those j^{th} and k^{th} coordinates are represented as $[\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2]^T$ and $[\mathbf{e}_0^*, \mathbf{e}_1^*, \mathbf{e}_2^*]^T$, respectively.

Given $\begin{bmatrix} \mathbf{e}_0^T \\ \mathbf{e}_1^T \\ \mathbf{e}_2^T \end{bmatrix} [\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2]$ is the identity matrix, the matrix $[\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2]$ is an orthogonal matrix and the columns of $[\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2]^T$ are orthonormal vectors as well, which can be used to build the orthonormal basis set of the j^{th} camera coordinate. Therefore, in \mathbb{R}^3 an arbitrary vector \mathbf{x} can be represented by two orthonormal sets, $(\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2)^T$ and $(\mathbf{e}_0^*, \mathbf{e}_1^*, \mathbf{e}_2^*)^T$, independently,

$$\begin{aligned} \mathbf{x} &= (\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2)^T (x_0, x_1, x_2)^T \\ &= (\mathbf{e}_0^*, \mathbf{e}_1^*, \mathbf{e}_2^*)^T (x_0^*, x_1^*, x_2^*)^T \end{aligned} \quad (5)$$

Finally, $(x_0, x_1, x_2)^T = (\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2)(\mathbf{e}_0^*, \mathbf{e}_1^*, \mathbf{e}_2^*)^T (x_0^*, x_1^*, x_2^*)^T$ where the rotation motion $R_{c_j c_k}$ from camera k to camera j is $[\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2][\mathbf{e}_0^*, \mathbf{e}_1^*, \mathbf{e}_2^*]^T$.

Two-Observation case. In the spatial case where two linearly independent direction vectors are detected, \mathbf{u}_2 can be achieved by the cross product process of \mathbf{u}_0 and \mathbf{u}_1 . Obviously, the new set $[\mathbf{u}_0, \mathbf{u}_1, \mathbf{u}_0 \times \mathbf{u}_1]$ maintains the orthogonal property, which is the minimal solution for relative pose estimation problems.

Orthogonal-Observation case. As discussed in Section 2, the MW assumption is enforced mostly by SLAM/VO methods designed to work indoor [15, 14, 19, 38], achieving particularly good results when the MW assumption holds. When the observation vectors \mathbf{v}_m^j are orthogonal, the projection operation between different vectors is zero and the proposed method degenerates to a multi-MW case,

$$\begin{aligned} R_{c_j c_k} &= R_{c_j M_i} R_{c_k M_i}^T \\ &= \begin{bmatrix} \frac{\mathbf{v}_0^j}{\|\mathbf{v}_0^j\|}, \frac{\mathbf{v}_1^j}{\|\mathbf{v}_1^j\|}, \frac{\mathbf{v}_2^j}{\|\mathbf{v}_2^j\|} \end{bmatrix} \begin{bmatrix} \frac{\mathbf{v}_0^k}{\|\mathbf{v}_0^k\|}, \frac{\mathbf{v}_1^k}{\|\mathbf{v}_1^k\|}, \frac{\mathbf{v}_2^k}{\|\mathbf{v}_2^k\|} \end{bmatrix}^T. \end{aligned} \quad (6)$$

For single-MW scenarios, a global orthogonal set can be obtained by every frame, therefore $R_{c_j w}$, from world to camera C_j , can be computed by $R_{c_j M} R_{c_0 M}^T$, here $R_{c_0 w}$ is an identity matrix.

Compared with the visual compass [15] method making use of a combination of line and plane features from MW [14] to estimate camera rotation, our graph is more robust and flexible. Furthermore, compared to [31] that generates four rotation candidates after aligning two frames' vanishing points, our method not only leverages plane features, but also solves the ambiguity regarding the directions of the vanishing points [31].

After the relative rotation pose estimation step between two frames, in case of no overlap between them, we need to make use of their neighboring frames to compute translation vectors. Note that only two correspondences are required in translation estimation by making use of Equation 3, which is particularly suited to deal with scenes and environments characterized by different texture types compared to traditional approaches [26, 3].

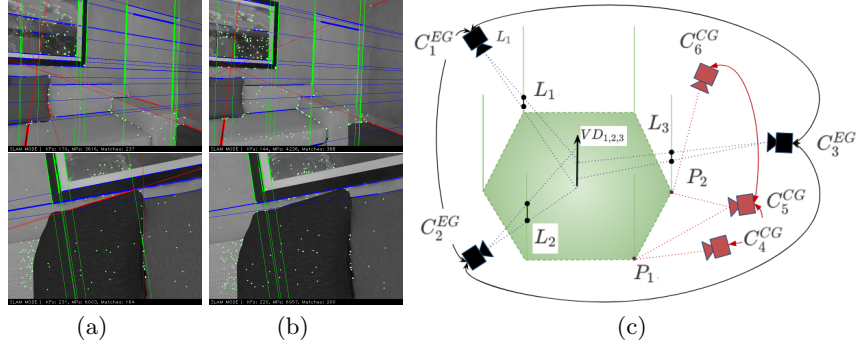


Fig. 3. Vanishing point detection and Rotation connection examples. (a) Detection results of J-Linkage. (b) Refined results by our system. (c) E-Graph (black) and co-visibility graph (red).

4 Extensibility Graph (E-Graph)

As shown in Figure 3(c), the E-Graph method builds rotation connections (edges) between frames $[C_1^{EG}, C_2^{EG}, C_3^{EG}]$ that share global directions instead of any low-level correspondences (like points and lines). At the same time, no connection between C_4^{CG} and C_6^{CG} can be made since these frames have no co-visible features within the co-visibility graph. The proposed connection strategy will be detailed in the following subsections.

4.1 Landmarks from a RGB-D frame

Similar to the co-visibility graph, the proposed graph is also a topological representation of scenes. The difference is that the proposed graph is built based on the scene structure rather than on overlapping parts between frames. The distance between connected frames in a co-visibility graph tends to be small (see Figure 3) since two frames that are distant from each other rarely overlap, leading to the pose of the current frame being estimated based on the last frame or last keyframes only. The issue can be alleviated by using global bundle adjustment and loop closure modules, although they bring in intensive computation and trajectory constraints (e.g. need to re-visit a certain area).

In our graph $\mathcal{G} = [\mathcal{N}_c, \mathcal{N}_m, \mathcal{E}]$ frames and landmarks are regarded as nodes \mathcal{N}_c and \mathcal{N}_m respectively, while \mathcal{E} represents the edges among connected frames. Note that landmarks are border-less planes and vanishing directions, e.g. $VD_{1,2,3}$, of lines detected in multiple views. In particular, an edge is established between two frames every time two or more structural elements are matched across them.

Features and landmarks. Vanishing directions are estimated from parallel lines detected by a joint 2D-3D process, where LSD [35] is used to extract 2D line features from RGB images. Meanwhile, AHP [5] is carried out to extract plane features from depth maps.

Firstly, as shown in Figure 3(a), we make use of the J-Linkage algorithm to classify detected 2D lines into different groups of parallel lines as described in [34]. However, there are still outliers left based on the 2D process. To solve this issue, we take advantage of depth maps to check the directions of lines in each group by using RANSAC to detect the best direction vector VD_n to represent the group S_n .

As for planar landmarks, we make use of the Hessian ($\boldsymbol{\pi} = (\mathbf{n}^\pi, d^\pi)$) to represent a plane detected from the i^{th} frame, where \mathbf{n}^π denotes the normal vector and d^π represents the distance between the camera center and this plane, which is transferred to world coordinates via the initial pose T_{wc_i} .

4.2 Data Association

After generating vanishing directions and planes, we now explain how to initialize and update them.

Initialization. Combined with the first keyframe Kf_0 , detected planes and optimized vanishing directions are used to initialize the E-Graph. The camera pose T_0 of Kf_0 is set as the world coordinate for landmarks in the E-Graph. Planes $\boldsymbol{\pi}_i$ measured by Kf_0 are transferred to the graph directly as,

$$\mathcal{G}_0 = [\mathcal{N}_{c_0}, \mathcal{N}_{lm_0}, \mathcal{E}_0] \quad (7)$$

where \mathcal{N}_{c_0} is Kf_0 and \mathcal{E}_0 has no edges yet. \mathcal{N}_{lm_0} contains $[\boldsymbol{\pi}_i, VD_i, PD_j]$, where VD_i and PD_j refer to two different types of 3D lines detected in the RGB-D frame: the former refers to lines that are parallel to at least another 3D line, the latter to lines that are not parallel to any line. The first type of lines can generate vanishing directions VD_i in a single view, which are stored into the graph directly, similarly to planes. In addition, lines that do not have parallel lines detected in this RGB-D frame are marked as potential vanishing direction PD_j . In case parallel lines will be detected in successive frames, these lines will also be transferred to VD_j , otherwise, they are removed from the E-Graph.

Landmarks fusion. For each new input frame we need to extract vectors \mathbf{n}^π , VD and PD from the current frame. After rotating VD_i^c to the world coordinate frame as VD_i^w , if the direction between VD_i^c is parallel to $VD_k^w, k \in [0, \dots, m]$, where m is the number of vanishing directions saved in E-Graph, VD_i^c is then associated to the graph. To solve the unsure issues [31] of vanishing directions, we will unify the direction during the association process by using

$$\tilde{VD}_i^c = \begin{cases} VD_i^c & (|norm(VD_i^w \cdot VD_k^w) - 1| < th_{vd}) \\ -VD_i^c & (|norm(VD_i^w \cdot VD_k^w) + 1| < th_{vd}) \end{cases} \quad (8)$$

where $norm(\cdot)$ shows a dot product between two normalized vectors and $|\cdot|$ is the absolute difference. th_{vd} is a threshold to check the angle distance between two vectors. To include additional graph connections, we also try to associate

PD_j^c with VD_k^w and PD_k^w . If new pairs can be made at this stage, the associated PD vectors are transferred to the vanishing directions and fused into the graph.

Since the vanishing direction is independent from translation motion, VD_i^w , the vanishing direction in the world coordinate can be obtained as

$$VD_i^w = R_{wc} VD_i^c \quad (9)$$

where R_{wc} is the rotation motion from the camera coordinate frame to the world coordinate frame.

In certain indoor scenes, e.g. a corridor or hallway, when a robot moves along the wall, an extended planar region is detected across multiple views, with most of these views encompassing no overlap. To address this issue, we extract the normal vector $[n_x^c, n_y^c, n_z^c]$ of the plane in the camera coordinate, which can be fused into the world coordinate in the same way as the vanishing directions.

Edge connection. In E-Graph, all landmarks come from keyframes that follow the decision mechanisms of a feature-based SLAM system [24, 20], which we summarize in the following. A new keyframe is detected if it satisfies one of the following two conditions: 1) 20 frames have passed from the last keyframe; 2) the current frame tracks less than 85% points and lines correspondences with the last keyframe. Furthermore, when the current frame detects a new plane or a new vanishing direction, the frame is considered as a new keyframe. In addition, new landmarks connected to this keyframe are also merged into the graph at this stage.

By sequentially processing keyframes, if more than two pairs of matched landmarks are observed between two keyframes, an edge will be created to connect the respective two graph nodes. As shown in Figure 2, C_j and C_k detect the plane π and the same vanishing point generated by L_1 and L_2 . Notably, even if these two frames do not have any correspondence, they can still be connected in our E-Graph.

5 Experiments

In this section, the proposed system is evaluated on different indoor benchmarks: ICL-NUIM [10] and TUM RGB-D [33]. ICL-NUIM [10] contains eight synthetic sequences recorded in two scenarios (living room and office room). TUM RGB-D [33] is recorded in real scenarios and includes varied sequences in terms of texture, scene size, presence of moving objects, etc.

Rotation estimation. The proposed rotation algorithm is compared with other state-of-the-art orientation estimation approaches. Compass [15] makes use of a single line and plane. OPRE [43] and GOME [12] estimate the distribution of surface normal vectors based on depth maps. OLRE [2] and ROVE [16] take advantage of vanishing directions for rotation estimation. Importantly, Compass, GOME, OLRE, OPRE, and P-SLAM [20] are all based on the MW assumption, while our method, ORB-SLAM2 [25] and ROVE are designed for general scenes.

Translation estimation. Since the rotation of the current frame is estimated from a keyframe that may not be overlapping with the current frame, we follow the 3D translation estimation model [26, 20] to estimate the translation \mathbf{t} based on the predicted rotation. In this module, re-projection errors from point-line-plane feature correspondences are used to build a target optimization function, $\mathbf{t} = \operatorname{argmin}(\sum_{j=0}^n e_{i,j}^{\pi} \Lambda^{\pi} e_{i,j}^{\pi} + e_{i,j}^L \Lambda^L e_{i,j}^L + e_{i,j}^P \Lambda^P e_{i,j}^P)$, where e^{π} , e^L and e^P are re-projection error functions for planes, lines and points, respectively. The target function is optimized by using the Levenberg-Marquardt method. The translation is compared with the following state-of-the-art methods. ORB-SLAM2 [26] and ORB-SLAM3 [3] are popular keypoint-based SLAM systems. In our experiments, for fairness of comparison the loop closure is removed to reduce the effect of the back-ends. SP-SLAM [39] additionally uses points and planes in the tracking and optimization modules based on ORB-SLAM2. P-SLAM [19] assumes the indoor environments as MW, and includes a refinement module to make the tracking process more robust. Moreover, we also compare our system with GPU-based methods, including BadSLAM [32] and BundleFusion [4].

Dense mapping. In this paper, a mapping module is implemented to reconstruct unknown environments in sparse and dense types. The sparse map is reconstructed by the point-line-plane features extracted from keyframes, which supports a frame-to-map pose refinement step. Since sparse maps cannot provide enough information for robots, our system also generates a dense mesh map incrementally based on CPU. When a new keyframe is generated from the tracking thread, we make use of the estimated camera pose and the RGB-D pair to build a dense TSDF model based on [42, 27]. After that, the marching cubes method [21] is exploited to extract the surface from voxels.

Metrics. The metrics used in our experiments include absolute trajectory error (ATE), absolute rotation error (ARE), and relative pose error (RPE) that shows the difference in relative motion between two pairs of poses to evaluate the tracking process. Our results are reported in Table 2 and obtained on an Intel Core @i7-8700 CPU @3.20GHz and without any use of GPU resources.

5.1 ICL NUIM dataset

As shown in Table 1, the proposed method outperforms other MW-based and feature-based methods in terms of average rotation error. In *office room* sequences, OPRE and P-SLAM also perform well since orthogonal planar features can be found in the environment. However, in *office room 0*, parts of the camera movement only contain a single plane and some lines, leading to performance degradation, while our method achieves robust orientation tracking by taking advantage of a set of non-parallel planes and lines.

Furthermore, we compare the translation results against two feature-based methods as shown in Table 2. The first four sequences are related to a living room scenario, while the remaining sequences are from an office scenario. All methods obtain good results in *living room 0* where the camera moves back and

Table 1. Comparison of the average value of the absolute rotation error (degrees) on ICL-NUIM and TUM RGB-D structural benchmarks. The best result for each sequence is bolded. \times shows that the method fails to track the orientation.

Sequence	Ours	Compass [15]	OPRE [43]	GOME [12]	ROVE [16]	OLRE [2]	ORB2 [26]	P-SLAM [20]
office room 0	0.11	0.37	0.18	5.12	29.11	6.71	0.40	0.57
office room 1	0.22	0.37	0.32	\times	34.98	\times	2.30	0.22
office room 2	0.39	0.38	0.33	6.67	60.54	10.91	0.51	0.29
office room 3	0.24	0.38	0.21	5.57	10.67	3.41	0.36	0.21
living room 0	0.44	0.31	\times	\times	\times	\times	0.97	0.36
living room 1	0.24	0.38	0.97	8.56	26.74	3.72	0.22	0.26
living room 2	0.36	0.34	0.49	8.15	39.71	4.21	0.83	0.44
living room 3	0.36	0.35	1.34	\times	\times	\times	0.42	0.27
<i>f3_stru_notex</i>	4.46	1.96	3.01	4.07	\times	11.22	\times	4.71
<i>f3_stru_tex</i>	0.60	2.92	3.81	4.71	13.73	8.21	0.63	2.83
<i>f3_l_cabinet</i>	1.45	2.04	36.34	3.74	28.41	38.12	2.79	2.55
<i>f3_cabinet</i>	2.47	2.48	2.42	2.59	\times	\times	5.45	1.18

forth between the two parallel walls. P-SLAM detects a good MW model, and ORB-SLAM3 also observes enough features, benefiting from paintings hanging on the wall and small furniture. Compared with the living room, the office room has many low-textured regions. The performance of feature-based algorithms is not as good as in the living room scenes, especially in *office room 1* and *office room 3*.

Table 2. Comparison in terms of translation RMSE (m) for ICL-NUIM and TUM RGB-D sequences. \times means that the system fails in the tracking process.

Sequence	Ours	P-SLAM[20]	ORB-SLAM3[3]
office room 0	0.014	0.068	0.035
office room 1	0.013	0.020	0.091
office room 2	0.020	0.011	0.010
office room 3	0.011	0.012	0.096
living room 0	0.008	0.006	0.006
living room 1	0.006	0.015	0.206
living room 2	0.017	0.020	0.018
living room 3	0.021	0.012	0.019
<i>f1_360</i>	0.114	\times	0.108
<i>f1_room</i>	0.095	\times	\times
<i>f2_rpy</i>	0.002	0.154	0.003
<i>f2_xyz</i>	0.003	0.009	0.004
<i>f3_l.o.house</i>	0.012	0.122	0.009
<i>f3_stru_notex</i>	0.017	0.025	\times
<i>f3_l_cabinet</i>	0.058	0.071	0.072

To analyze the relationship between rotation and translation results of different methods, absolute translation and rotation errors on the *office room 0* sequence are presented in Figure 4. When the camera moves to the ceiling, the number of detected features decreases, then an interesting phenomenon is

witnessed (see also Figure 4(a)): the tracking error of feature-based systems quickly and drastically increases, then gradually fades as the number of features increases. At the same time, our method and P-SLAM exhibit a more robust performance when they face this challenge. An important difference is that, while P-SLAM underperforms due to the non-rigid MW scene, our method’s performance is accurate thanks to the use of the E-Graph, which demonstrates to be more flexible than MW-based paradigms.

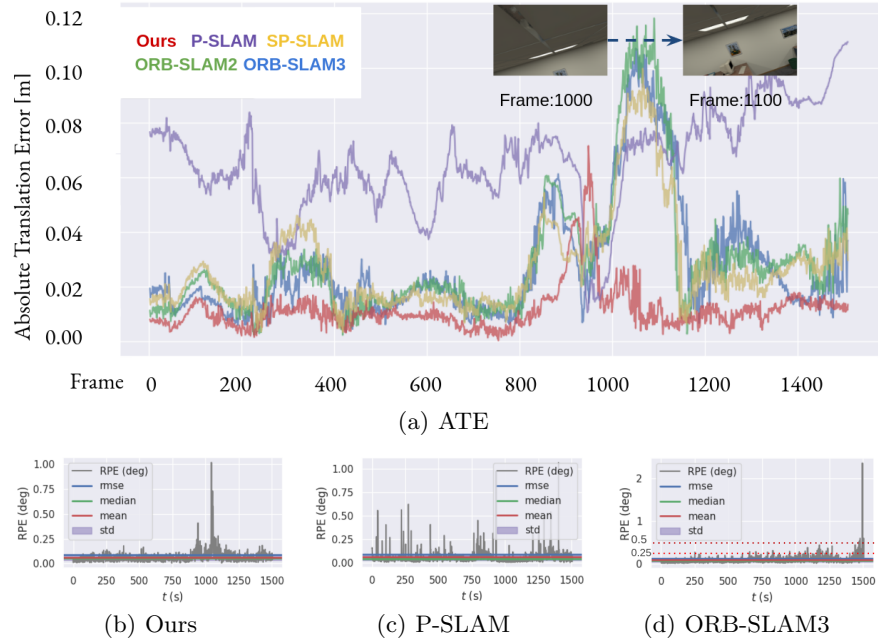


Fig. 4. Comparison of the proposed system against state-of-the-art methods in the *office room 0* sequence of ICL NUIM in terms of mean/average absolute translation errors (top) and rotation errors (bottom).

5.2 TUM RGB-D

Different types of sequences are included from the TUM RGB-D benchmark, which aims to test general indoor scenes with low-textured scenes and sharp rotational motions. *f1_360*, *f1_room*, *f2_rpy* and *f2_xyz* are recorded in real office scenes, but the camera’s rotation motion changes sharply especially in the first sequence. *f3_lo_house*, *f3_sn_near* and *f3_l_cabinet* contain more structural information, where *f3_sn_near* is built on two white corners, and *f3_l_cabinet* records several movements around the white cabinet. Table 1 shows that ROVE, OLR and ORB-SLAM2 have problems in low/non-textured regions. In *f3_l_cabinet*




Fig. 5. Scene and graphs of *f3.Lo.house*. (a) 2D image, (b) dense mesh model, (c) sparse map, (d) E-Graph, (e) co-visibility graph.

that is not a rigid MW environment, the quality of depth maps is noisy, the surface normal maps extracted by OPRE have a negative effect on rotation estimation.

Table 3. ATE RMSE results (cm) on the TUM RGB-D dataset. Results for BundleFusion and BadSLAM are taken from [32]

Sequence	Ours	BundleFusion [4]	ElasticFusion [37]	BadSLAM [32]
	CPU	GPU	GPU	GPU
f1_desk	1.0	1.6	2.0	1.7
f2_xyz	0.7	1.1	1.1	1.1
f3_office	1.4	2.2	3.6	1.7

For structural sequences listed in Table 1, P-SLAM shows stable performance. In Table 2, general scenes are added as a comparison. As listed in Table 2, the keypoint-based method [3] cannot achieve robust results in *f3.sn.near*, i.e., a textureless scenario, while the MW-based method [20] has problems when the scene structure breaks the MW assumption, by reporting a low performance in *f2.rpy* and *f3.Lo.house*, and even losing track in *f1.360* and *f1.room*. Therefore, the proposed method shows more robust performances in different types of scenarios, compared with MW-based systems [20, 15] and feature-based approaches [26, 3]. Furthermore, compared with GPU-based systems, our system only works on limited computation sources. As shown in Figure 5, *f3.Lo.house* is used to compare E-Graph and co-visibility graph. As clearly shown, E-Graph allows connecting more distant keyframes than a co-visibility graph. When two keyframes can be connected together, drifting phenomena can more easily be

limited, in a similar way to the underlying idea behind loop closure. The cabinet scene is also a difficult sequence for point-based methods (see Figure 6(b)) since point features are concentrated in a few boundary regions. However, our method can deal with this type of scene where the same plane is observed in a number of frames.




Fig. 6. Scene and graph of *f3_cabinet*. (a) E-Graph, (b) trajectory from ORB-SLAM3, (c) sparse map, (d) 2D image.

6 Conclusion

This paper proposed a new graph structure, E-Graph, to reduce tracking drift based on plane normals and vanishing directions in a scene, which can be used to build a rotation connection between two frames without visual overlap. The advantage of this idea is that rotation errors that occur between two frames have small or no effect on this relative rotation estimation step. Based on the proposed graph, a minimal solution is presented, that shows that two landmarks and two correspondences can be used to solve the relative camera pose. Therefore, the proposed method is better suited for texture-less scenes compared with traditional minimal solutions based on co-visible features. However, the proposed method also has limitations. Compared with point-based systems, our approach requires more types of features. Furthermore, since we need vanishing directions and plane vectors, the method is more suitable for man-made scenes.

Feature work. The E-Graph is a new tool to establish connections across frames and keyframes. An interesting topic for future exploration is considering a covisibility graph and our graph together to revisit pose estimation and obtain further improvements in drift removal.

Acknowledgements. We gratefully acknowledge Xin Li, Keisuke Tateno, Nicolas Brasch and Dr. Liang Zhao for the helpful discussion.

References

1. Andrew, A.M.: Multiple view geometry in computer vision. Kybernetes (2001)
2. Bazin, J.C., Pollefeys, M.: 3-line ransac for orthogonal vanishing point detection. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 4282–4287. IEEE (2012)
3. Campos, C., Elvira, R., Rodríguez, J.J.G., Montiel, J.M., Tardós, J.D.: Orb-slam3: An accurate open-source library for visual, visual-inertial and multi-map slam. arXiv preprint arXiv:2007.11898 (2020)
4. Dai, A., Nießner, M., Zollhöfer, M., Izadi, S., Theobalt, C.: Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. ACM Transactions on Graphics (ToG) **36**(4), 1 (2017)
5. Feng, C., Taguchi, Y., Kamat, V.R.: Fast plane extraction in organized point clouds using agglomerative hierarchical clustering. In: 2014 IEEE International Conference on Robotics and Automation (ICRA). pp. 6218–6225. IEEE (2014)
6. Galvez-Lpez, D., Tardos, J.D.: Bags of binary words for fast place recognition in image sequences. IEEE Transactions on Robotics **28**(5), 1188–1197 (2012). <https://doi.org/10.1109/TRO.2012.2197158>
7. Gao, X., Wang, R., Demmel, N., Cremers, D.: Ldso: Direct sparse odometry with loop closure. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2198–2204. IEEE (2018)
8. Gomez-Ojeda, R., Moreno, F.A., Zuniga-Noël, D., Scaramuzza, D., Gonzalez-Jimenez, J.: Pl-slam: A stereo slam system through the combination of points and line segments. IEEE Transactions on Robotics **35**(3), 734–746 (2019)
9. Guan, B., Zhao, J., Li, Z., Sun, F., Fraundorfer, F.: Minimal solutions for relative pose with a single affine correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1929–1938 (2020)
10. Handa, A., Whelan, T., McDonald, J., Davison, A.: A Benchmark for RGB-D Visual Odometry, 3D Reconstruction and SLAM. In: IEEE International Conference on Robotics and Automation (2014)
11. Hsiao, M., Westman, E., Zhang, G., Kaess, M.: Keyframe-based dense planar slam. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 5110–5117 (2017). <https://doi.org/10.1109/ICRA.2017.7989597>
12. Joo, K., Oh, T.H., Kim, J., Kweon, I.S.: Globally optimal manhattan frame estimation in real-time. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1763–1771 (2016)
13. Joo, K., Oh, T.H., Rameau, F., Bazin, J.C., Kweon, I.S.: Linear rgb-d slam for atlanta world. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 1077–1083. IEEE (2020)
14. Kim, P., Coltin, B., Jin Kim, H.: Linear rgb-d slam for planar environments. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 333–348 (2018)
15. Kim, P., Coltin, B., Kim, H.J.: Indoor rgb-d compass from a single line and plane. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4673–4680 (2018)
16. Lee, J.K., Yoon, K.J.: Real-time joint estimation of camera orientation and vanishing points. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1866–1874 (2015)
17. Li, H., Xing, Y., Zhao, J., Bazin, J., Liu, Z., Liu, Y.: Leveraging structural regularity of atlanta world for monocular slam. In: 2019 International

- Conference on Robotics and Automation (ICRA). pp. 2412–2418 (2019). <https://doi.org/10.1109/ICRA.2019.8793716>
18. Li, X., Li, Y., Pinar Örneke, E., Lin, J., Tombari, F.: Co-planar parametrization for stereo-slam and visual-inertial odometry. arXiv e-prints pp. arXiv–2009 (2020)
 19. Li, Y., Brasch, N., Wang, Y., Navab, N., Tombari, F.: Structure-slam: Low-drift monocular slam in indoor environments. *IEEE Robotics and Automation Letters* **5**(4), 6583–6590 (2020)
 20. Li, Y., Yunus, R., Brasch, N., Navab, N., Tombari, F.: Rgb-d slam with structural regularities. In: 2021 IEEE international conference on Robotics and automation (ICRA) (2021)
 21. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics* **21**(4), 163–169 (1987)
 22. Ma, L., Kerl, C., Stckler, J., Cremers, D.: Cpa-slam: Consistent plane-model alignment for direct rgb-d slam. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). pp. 1285–1291 (2016). <https://doi.org/10.1109/ICRA.2016.7487260>
 23. Mei, C., Sibley, G., Newman, P.: Closing loops without places. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 3738–3744. IEEE (2010)
 24. Mur-Artal, Raúl, M.J.M.M., Tardós, J.D.: ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics* **31**(5), 1147–1163 (2015). <https://doi.org/10.1109/TRO.2015.2463671>
 25. Mur-Artal, R., Montiel, J.M.M., Tardos, J.D.: Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics* **31**(5), 1147–1163 (2015)
 26. Mur-Artal, R., Tardós, J.D.: ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. *IEEE Transactions on Robotics* **33**(5), 1255–1262 (2017). <https://doi.org/10.1109/TRO.2017.2705103>
 27. Nießner, M., Zollhöfer, M., Izadi, S., Stamminger, M.: Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (ToG)* **32**(6), 1–11 (2013)
 28. Nistér, D.: An efficient solution to the five-point relative pose problem. *IEEE transactions on pattern analysis and machine intelligence* **26**(6), 756–770 (2004)
 29. Pomerleau, F., Colas, F., Siegwart, R.: A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends in Robotics* **4**(1), 1–104 (2015)
 30. Rosinol, A., Sattler, T., Pollefeys, M., Carlone, L.: Incremental visual-inertial 3d mesh generation with structural regularities. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 8220–8226. IEEE (2019)
 31. Salaün, Y., Marlet, R., Monasse, P.: Robust and accurate line-and/or point-based pose estimation without manhattan assumptions. In: European Conference on Computer Vision. pp. 801–818. Springer (2016)
 32. Schps, T., Sattler, T., Pollefeys, M.: Bad slam: Bundle adjusted direct rgb-d slam. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 134–144 (2019). <https://doi.org/10.1109/CVPR.2019.00022>
 33. Sturm, J., Engelhard, N., Endres, F., Burgard, W., Cremers, D.: A Benchmark for the Evaluation of RGB-D SLAM Systems. In: IEEE International Conference on Intelligent Robots and Systems (2012)
 34. Tardif, J.P.: Non-iterative approach for fast and accurate vanishing point detection. In: 2009 IEEE 12th International Conference on Computer Vision. pp. 1250–1257. IEEE (2009)

35. Von Gioi, R.G., Jakubowicz, J., Morel, J.M., Randall, G.: LSD: A Fast Line Segment Detector with a False Detection Control. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32**(4), 722–732 (2010)
36. Wang, R., Schworer, M., Cremers, D.: Stereo dso: Large-scale direct sparse visual odometry with stereo cameras. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3903–3911 (2017)
37. Whelan, T., Leutenegger, S., Salas-Moreno, R., Glocker, B., Davison, A.: Elastic-fusion: Dense slam without a pose graph. *Robotics: Science and Systems* (2015)
38. Yunus, R., Li, Y., Tombari, F.: Manhattanslam: Robust planar tracking and mapping leveraging mixture of manhattan frames. *arXiv preprint arXiv:2103.15068* (2021)
39. Zhang, X., Wang, W., Qi, X., Liao, Z., Wei, R.: Point-plane slam using supposed planes for indoor environments. *Sensors* **19**(17), 3795 (2019)
40. Zhao, J., Kneip, L., He, Y., Ma, J.: Minimal case relative pose computation using ray-point-ray features. *IEEE transactions on pattern analysis and machine intelligence* **42**(5), 1176–1190 (2019)
41. Zhou, H., Zou, D., Pei, L., Ying, R., Liu, P., Yu, W.: Structslam: Visual slam with building structure lines. *IEEE Transactions on Vehicular Technology* **64**(4), 1364–1375 (2015). <https://doi.org/10.1109/TVT.2015.2388780>
42. Zhou, Q.Y., Koltun, V.: Dense scene reconstruction with points of interest. *ACM Transactions on Graphics (ToG)* **32**(4), 1–8 (2013)
43. Zhou, Y., Kneip, L., Rodriguez, C., Li, H.: Divide and conquer: Efficient density-based tracking of 3d sensors in manhattan worlds. In: *Asian Conference on Computer Vision*. pp. 3–19. Springer (2016)