

# PeLiCal: Targetless Extrinsic Calibration via Penetrating Lines for RGB-D Cameras with Limited Co-visibility

Jaeho Shin<sup>1</sup>, Seungsang Yun<sup>1</sup> and Ayoung Kim<sup>1\*</sup>

**Abstract**—RGB-D cameras are crucial in robotic perception, given their ability to produce images augmented with depth data. However, their limited field of view (FOV) often requires multiple cameras to cover a broader area. In multi-camera RGB-D setups, the goal is typically to reduce camera overlap, optimizing spatial coverage with as few cameras as possible. The extrinsic calibration of these systems introduces additional complexities. Existing methods for extrinsic calibration either necessitate specific tools or highly depend on the accuracy of camera motion estimation. To address these issues, we present PeLiCal, a novel line-based calibration approach for RGB-D camera systems exhibiting limited overlap. Our method leverages long line features from surroundings, and filters out outliers with a novel convergence voting algorithm, achieving targetless, real-time, and outlier-robust performance compared to existing methods. We open source our implementation on <https://github.com/joomeok/PeLiCal.git>.

## I. INTRODUCTION

In robotic perception, RGB-D cameras serve as pivotal sensors due to their capacity to capture images accompanied by corresponding depth values. Due to the limited FOV of the sensor, common practice is to employ a multi-camera configuration to ensure comprehensive coverage of the surrounding environment.

In systems deploying multiple RGB-D cameras, unlike RGB camera setups designed for shared FOV (e.g., binocular camera configurations), the system is set up to reduce the overlap in the FOV. This strategy aims to cover the nearby environment using a minimum number of cameras. However, such configuration makes it challenging to employ a calibration pattern observed simultaneously within the shared FOV when calibrating the extrinsic parameters of the system. Several approaches have been proposed to address these challenges in the extrinsic calibration of camera systems with limited co-visibility.

Existing works bifurcate into two distinct approaches: one leveraging specialized equipment and the other adopting a hand-eye-like calibration by estimating each camera motion [1]. In the former case, Li et al. [2] utilized a target with unique features to easily estimate the location of points in the partially detected pattern board. In the other case, Kumar et al. [3] introduced additional apparatuses like mirrors. Unfortunately, these approaches are impractical because employing such specialized tools for calibration is burdensome. In the latter case, the calibration results are

<sup>†</sup>This work was supported by IITP (No.2022-0-00480), KIAT (P0020536) and Interdisciplinary Research Initiatives Program SNU (2023).

<sup>1</sup>J. Shin, S. Yun and A. Kim are with the Dept. of Mechanical Engineering, SNU, Seoul, S. Korea [leah100, seungsang, ayoungk]@snu.ac.kr




Fig. 1: According to the inlier ratio of RANSAC-based fitting from the depth images, each line pair is classified into a *full 3D* case or *perspective-n-line (PnL)* case, and different constraints are employed (top). Two checkerboard planes from each camera are merged by extrinsic parameters estimated from PeLiCal (bottom). Our method uses long lines matched from both images, and on the condition that the cameras are set up to observe the same feature, it can reliably determine the pose in cases with sufficient overlap (left) as well as in cases without any overlap (right).

highly dependent on the accuracy of the estimated camera motion.

Overcoming these challenges, we introduce a targetless but robust methodology for extrinsic calibration, specifically for RGB-D camera systems with limited overlap. Without requiring a large calibration target, which is commonly observable, we focus on line features intersecting two cameras. Accordingly, our algorithm performs reliably in real-time, without a pattern board, additional external devices, or inter-frame motion estimation, even if co-visibility is limited. Additionally, measurements that induce estimations closely matching the conditions of the rotation matrix and which are confirmed as convergent by a voting algorithm are incorporated, enhancing robustness against outliers. The key contributions of our study are encapsulated as follows:

- We present a robust, targetless approach for calibrating the extrinsic parameters of RGB-D cameras across diverse FOV settings using penetrating line matching, available from surroundings.
- Our algorithm selectively incorporates informative

scenes by projecting the estimated rotation matrix on the  $SO(3)$  manifold and calculating its shortest distance. Furthermore, the existence of precise translation is validated by examining the convergence of 3D geometric constraints, derived from the transformation of Plücker coordinates.

- The calibration accuracy of our method is validated from various configurations of FOV using a specially designed device. The algorithm shows superior and stable accuracy with existing calibration tools, especially in settings with an increased baseline in stereo setups.

## II. RELATED WORK

### A. Line Based Camera Pose Estimation

Leveraging line features has been widely adopted to estimate the camera pose. Among these approaches, the algorithm that relies on matching 3D to 2D pairs is the PnL. Most studies estimate the camera pose by formulating equations based on the associated 3D and 2D line features. Hartley in [4] proposed the method to calculate the essential matrix from a set of nine lines. In the context of associating 3D to 2D line pairs, a commonly invoked constraint is that the preimage of the 2D line should encapsulate the transformed 3D line as in [5]. Příbyl et al. employ representation of lines using plücker coordinates to estimate the camera pose in [6]. Xu et al. outlined the computation of the rotation matrix by decomposing it into two distinct rotation angles incorporating an auxiliary model frame, as detailed in [7]. The geometric condition of the line presented in the previous research is used as a theoretical background for our method.

### B. Non-overlapping FOV Camera Calibration

In a study where the FOV of the cameras exhibit no overlap, the methodology proposed by Kumar et al. in [3] employs reflected images of a chessboard via a planar mirror to determine extrinsic parameters. In [2], Li et al. proposes a calibration method using a calibration pattern enriched with unique features, which enables accurate localization of features, even when the pattern board is partially observed. However, these techniques are impractical because specialized equipment or a substantially large pattern board is required. Several algorithms have been developed to alleviate this challenge that harnesses camera motion estimation to retrieve the 6-DOF extrinsic parameters as in [1, 8]. However, these motion-centric calibration approaches are intrinsically reliant on the precision of the camera's motion.

The investigation conducted by Perez-Yus et al. closely corresponds to our algorithm since it exploits 3D lines to calibrate the limited co-visibility cameras in [9]. Their method entails extracting 2D and 3D lines and projecting 3D lines onto 2D images using initial extrinsic parameters to establish correspondences among line features. Subsequently, the extrinsic parameters are refined repeatedly by matching the feature based on the estimated pose until convergence. Given the dependence on the initial parameter for associating line features, achieving accurate calibration parameters is infeasible when the initial value is erroneous. In contrast,




Fig. 2: Line features (depicted in red) can be matched between two cameras even when they lack overlapping FOV areas. In contrast to point features, long lines can identify corresponding pairs under these conditions, allowing for determining geometric constraints.

our proposed algorithm calculates initial parameters using robustly matched lines, followed by feasibility evaluation of the solution and optimization. As a result, consistent calibration parameters can be obtained regardless of the initial conditions.

## III. METHODOLOGY

### A. Notation

In accordance with our notation conventions, we represent scalars using italicized characters, vectors with boldface lowercase characters, and matrices with boldface uppercase characters. The subscripts  $t$  and  $s$  are utilized to denote parameters associated with the target and source cameras, respectively. For clarity in notation, a 3D line is represented as  $\mathbf{L}$ , with its corresponding Plücker coordinates expressed as  $\mathcal{L} = [\mathbf{d}, \mathbf{m}]$ . Correspondingly, its associated dual Plücker matrix is represented as  $\mathbf{L}^*$ . Within the 2D image plane, the notation  $\mathbf{l}$  represents a 2D line. The matrix  $\mathbf{K}$  denotes the line projection matrix, while  $\mathbf{K}$  corresponds to the intrinsic camera matrix. Given a 3D vector, the operator  $(\cdot)^\wedge$  denotes its conversion to a skew-symmetric matrix. Furthermore, a vector with an overbar on it signifies its representation in homogeneous coordinates.

### B. Merged Quadratic System

To leverage 3D lines within our algorithm, we initiate by extracting and matching line features from a pair of RGB images using a novel deep matcher designed explicitly for line features [10], enabling line matching between significantly different scenes as shown in Fig. 2. Based on the matching results, the 3D coordinates of the corresponding line are established via a line fitting method using random sample consensus (RANSAC), exploiting depth images. As depicted in Fig. 1, when the proportion of inliers exceeds a predetermined threshold in both images, the feature is categorized as the *full 3D case*. In this scenario, a “3D point on 3D line” constraint is incorporated into an equation system. On the other hand, if only the line associated with the source camera aligns accurately, the feature is classified as *PnL case*, and we apply a constraint as detailed in [11].




Fig. 3: Pipeline of the proposed calibration algorithm. The procedure iteratively continues until the optimized solution reaches convergence and its associated cost drops below a predefined threshold.

1) *Full 3D Case*: In the first scenario, we utilize a constraint that transformed 3D points on the source line should lie on the 3D target line. This representation begins from the property of dual Plücker matrix. Given a 3D point  $\mathbf{X}$ , the subsequent equation holds:

$$\begin{aligned}
 \mathbf{L}_t^* \bar{\mathbf{X}}_{tj} &= \begin{pmatrix} -\mathbf{d}_t^\wedge & \mathbf{m}_t \\ -\mathbf{m}_t^T & 0 \end{pmatrix} \begin{pmatrix} \mathbf{R}\mathbf{X}_{sj} + \mathbf{t} \\ 1 \end{pmatrix} \\
 &= \begin{pmatrix} -\mathbf{d}_t^\wedge \mathbf{R}\mathbf{X}_{sj} - \mathbf{d}_t^\wedge \mathbf{t} + \mathbf{m}_t \\ -\mathbf{m}_t^T \mathbf{R}\mathbf{X}_{sj} - \mathbf{m}_t^T \mathbf{t} \end{pmatrix} \\
 &= \begin{pmatrix} -\mathbf{d}_t^\wedge \mathbf{R}\mathbf{X}_{sj} + \mathbf{m}_t \\ -\mathbf{m}_t^T \mathbf{R}\mathbf{X}_{sj} \end{pmatrix} + \begin{pmatrix} -\mathbf{d}_t^\wedge \mathbf{t} \\ -\mathbf{m}_t^T \mathbf{t} \end{pmatrix} \\
 &= \mathbf{0}_{4 \times 1}, \quad (j = 1, 2)
 \end{aligned} \tag{1}$$

where  $j$  represents each endpoint, respectively. Thus, two constraints are derived from a single line pair  $\mathbf{L}_t \leftrightarrow \mathbf{L}_s$ .

2) *PnL Case*: For the second case, we adopt the constraint, also utilized in [11]. Since the projection of transformed 3D point  $\hat{\mathbf{x}}_t$  lies on 2D line  $\mathbf{l}_t$  of the target image,  $\mathbf{l}_t^T \hat{\mathbf{x}}_{tj} = 0$  holds. Consequently, from a singular correspondence  $\mathbf{l}_t \leftrightarrow \mathbf{L}_s$ , we derive two constraints:

$$(\mathbf{K}_t^T \mathbf{l}_t)^T (\mathbf{R}\mathbf{X}_{sj} + \mathbf{t}) = 0, \quad (j = 1, 2) \tag{2}$$

3) *System Solver and Solution Refinement*: To merge (1) and (2) into a unified system, we represent the rotation matrix  $\mathbf{R}$  using the CGR parameterization [12], defined as:

$$\mathbf{R} = \frac{\bar{\mathbf{R}}}{1 + \mathbf{s}^T \mathbf{s}}, \quad \bar{\mathbf{R}} = (1 - \mathbf{s}^T \mathbf{s})\mathbf{I}_3 + 2[\mathbf{s}]_\times + 2\mathbf{s}^T \mathbf{s}, \tag{3}$$

where  $\mathbf{s} = [s_1, s_2, s_3]$ . Multiplying  $(1 + \mathbf{s}^T \mathbf{s})$  to the both sides of (1) and (2), we get following equations which can be readily aggregated:

$$\begin{aligned}
 \begin{pmatrix} -\mathbf{d}_t^\wedge \bar{\mathbf{R}}\mathbf{X}_{sj} + (1 + \mathbf{s}^T \mathbf{s})\mathbf{m}_t \\ -\mathbf{m}_t^T \bar{\mathbf{R}}\mathbf{X}_{sj} \end{pmatrix} + \begin{pmatrix} -\mathbf{d}_t^\wedge \boldsymbol{\tau} \\ -\mathbf{m}_t^T \boldsymbol{\tau} \end{pmatrix} &= \mathbf{0}_{4 \times 1}, \\
 (\mathbf{K}_t^T \mathbf{l}_t)^T \bar{\mathbf{R}}\mathbf{X}_{sj} + (\mathbf{K}_t^T \mathbf{l}_t) \boldsymbol{\tau} &= 0, \quad (j = 1, 2),
 \end{aligned} \tag{4}$$

where  $\boldsymbol{\tau} = (1 + \mathbf{s}^T \mathbf{s})\mathbf{t}$ . Given  $M$  and  $N$  measurements included in each case, we stack them vertically to have the following quadratic system:

$$\mathbf{A}\mathbf{r} + \mathbf{B}\boldsymbol{\tau} = \mathbf{0}_{(8M+2N) \times 1}, \tag{5}$$

where  $\mathbf{r} = [s_1^2, s_2^2, s_3^2, s_1 s_2, s_1 s_3, s_2 s_3, s_1, s_2, s_3, 1]^T$ . Subsequently, (5) is solved using the technique presented in [11] to retrieve the rotation matrix and translation vector. Notably, the difference is that the system based on the CGR parameter is processed using the RE3Q3 solver [13].

The system's initial value is refined by optimizing a cost function combining the line reprojection error  $e_l$ , as described in [14], and the 3D point-to-line error  $e_L$  [15]:

$$\mathbf{e}_l = \begin{bmatrix} \frac{\mathbf{x}_{t1}^T \hat{\mathbf{l}}}{\sqrt{l_1^2 + l_2^2}} & \frac{\mathbf{x}_{t2}^T \hat{\mathbf{l}}}{\sqrt{l_1^2 + l_2^2}} \end{bmatrix}^T, \tag{6}$$

$$\mathbf{e}_L = (\mathbf{I}_3 - \mathbf{d}_t \mathbf{d}_t^T)(\mathbf{R}\mathbf{X}_{sj} + \mathbf{t} - \mathbf{X}_{tj}), \quad (j = 1, 2) \tag{7}$$

where  $\hat{\mathbf{l}}$  is the projected line on the target image. Using (6) and (7), we define our cost function as follows:

$$\min_{\mathbf{R}, \mathbf{t}} \left( \sum_{i=1}^M \mathbf{e}_{L_i}^T \mathbf{e}_{L_i} + \sum_{j=1}^N \mathbf{e}_{l_j}^T \mathbf{e}_{l_j} \right). \tag{8}$$

We employ Levenberg-Marquardt (LM) algorithm for optimization to refine the initial estimates of the parameters iteratively.

### C. Scene Selection

This section introduces algorithms to identify informative lines for estimating the rotation matrix and translation. It is necessary to assess whether the resulting measurement provides sufficient information for an accurate estimation. Determining the rotation matrix by solving the quadratic system as section III-B can be computationally demanding, considering that PeLiCal incorporates real-time features. Thus, our algorithm employs a lightweight approach based on a linear equation directly derived from the direction of the 3D line. Additionally, we evaluate the presence of a translation vector corresponding to the optimized rotation matrix in the form of a 3-dimensional line in the  $\mathbb{R}^3$  space.

1) *Selecting Lines for Rotation Estimation*: For the *full 3D case*, we derive the following constraint relating the direction vectors from each camera:

$$\mathbf{d}_t = \mathbf{R}\mathbf{d}_s. \tag{9}$$

---

**Algorithm 1** Convergence Voting.

---

```

1: Input:  $\epsilon_d$  and  $l_1, \dots, l_N$  from (14), (15);
2: Output: Converged inliers;
3: % Compute equidistance points
4:  $\mathbf{m} = \text{EquiDistance}(\{l_1, \dots, l_N\})$ 
5: % Initialize maximum inlier set
6:  $\mathcal{I} = \emptyset$ 
7: % Voting
8: for  $i = 1, \dots, \binom{N}{2}$  do
9:    $\mathcal{K}_i = \emptyset$ 
10:  for  $j = 1, \dots, N$  do
11:    if  $\text{PointToLineDistance}(\mathbf{m}_i, l_j) < \epsilon_d$  then
12:       $\mathcal{K}_i = \mathcal{K}_i \cup \{j\}$  % add to inlier set
13:    end if
14:  end for
15:  if  $n(\mathcal{K}_i) > n(\mathcal{I})$  then % update maximum inlier set
16:     $\mathcal{I} = \mathcal{K}_i$ 
17:  end if
18: end for
19: return  $\mathcal{I}$ 

```

---

In another case, we establish an additional constraint ensuring that the preimage of the 2D target line should contain the transformed 3D source line:

$$(\mathbf{P}_t^T \mathbf{l}_t)_{1:3}^T \mathbf{R} \mathbf{d}_s = 0, \quad (10)$$

where  $\mathbf{P}$  represents the projection matrix of the target camera. Incorporating both equations, we derive the following linear equation:

$$\mathbf{C} \text{vec}(\mathbf{R}) = \mathbf{b}, \quad (11)$$

where  $\mathbf{C}$  is  $(3M + N) \times 9$  matrix and  $\mathbf{b}$  is  $(3M + N)$ -dimensional vector derived by stacking multiple equations of (9), (10). After solving the equation and obtaining a 9-dimensional vector as a solution, we reshape it and attain a square matrix denoted as  $\mathbf{M}$  whose Singular Value Decomposition (SVD) is  $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ . When a new line measurement is introduced, we solve (11), subsequently projecting the resulting matrix onto the  $SO(3)$  manifold using a special orthogonalization technique as detailed in [16]. This projected matrix is denoted as  $\mathbf{R}'$  and can be expressed as  $\mathbf{U}\mathbf{\Sigma}'\mathbf{V}^T$ , where  $\mathbf{\Sigma}'$  is a diagonal matrix with diagonal elements  $(1, 1, \det(\mathbf{U}\mathbf{V}^T))$ .

Assume we define a set of matrices, each having an SVD representation  $\mathbf{U}\mathbf{D}\mathbf{V}^T$ , with  $\mathbf{D}$  being an arbitrary diagonal matrix, then, it can be proved that the minimum distance between any two elements within this set is equivalent to the Frobenius norm of the difference between the two diagonal matrices. Consequently, the distance between two matrices,  $\mathbf{M}$  and  $\mathbf{R}'$ , is computed as follows:

$$d_{SO(3)} = \|\mathbf{\Sigma} - \mathbf{\Sigma}'\|_F, \quad (12)$$

where  $\|\cdot\|_F$  means Frobenius norm. In the calibration process, only a new line pair that decreases the  $d_{SO(3)}$  distance is integrated.

2) *Evaluating Existence of Translation:* Despite exclusively employing measurements that bring the estimation closer to the  $SO(3)$  manifold, it remains necessary to validate whether a translation vector can be derived from it. Our algorithm uses geometric conditions derived from the transformation of Plücker coordinates, distinct from those employed in constructing the quadratic system.

In *full 3D case*, by transforming the Plücker coordinates of the source line, we obtain the following equation:

$$\begin{aligned} \mathbf{m}_t &= \mathbf{R}\mathbf{m}_s + \mathbf{t}^\wedge \mathbf{R}\mathbf{d}_s \\ &= \mathbf{R}\mathbf{m}_s - (\mathbf{R}\mathbf{d}_s)^\wedge \mathbf{t}, \end{aligned} \quad (13)$$

where a solution is represented as:

$$\mathbf{t} = (\mathbf{R}\mathbf{m}_s - \mathbf{m}_t) \times \mathbf{R}\mathbf{d}_s + k\mathbf{R}\mathbf{d}_s, \quad (14)$$

with  $k$  as an arbitrary scalar. This implies that the potential candidate for the translation vector should be situated along a 3D line, as expressed in (14). In *PnL case*, when projecting the transformed line represented with Plücker coordinates, we derive the following constraint by multiplying two endpoints of the 2D target line by it:

$$([\mathbf{R}\mathbf{d}_s]_\times \mathcal{K}^T \bar{\mathbf{x}}_j)^T \mathbf{t} = \bar{\mathbf{x}}_j^T \mathcal{K} \mathbf{R} \mathbf{m}_s. \quad (j = 1, 2) \quad (15)$$

For a single endpoint, (15) provides a potential candidate for the translation vector within a 3D plane. Exploiting both endpoints, we obtain a 3D line since the solution should lie on the intersection of two planes. Similarly to the previous case, given a single pair of *PnL case*, the potential solution for the translation vector also conforms to a 3D line. It is surprise to note that, in a noiseless scenario where an accurate rotation matrix is known, if we can acquire two measurements from either of the two cases, the translation vector can be precisely determined as the point of intersection between the two lines.

#### D. Outlier Rejection via Convergence Voting

In practical applications, noise in the line feature measurements diminish the precision of the rotation matrix. In such cases, we can still verify whether the estimated rotation matrix is sufficiently accurate to enable the existence

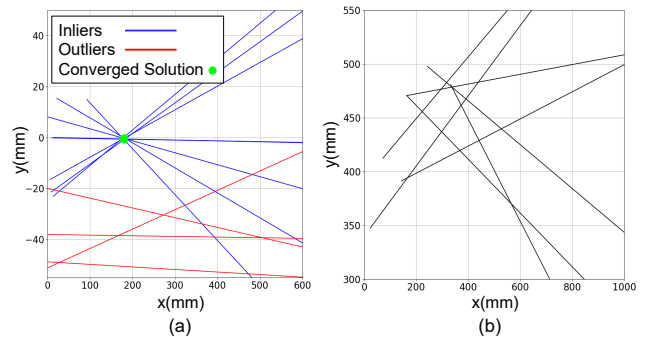


Fig. 4: The results of the convergence voting process with candidate lines projected onto the  $xy$ -plane: (a) The algorithm successfully computes a convergence point, distinguishing inliers from outliers. (b) The candidate lines for translation do not converge due to an inaccurate rotation matrix or errors in the measurement.




Fig. 5: Equipment for accurate variation of rotation and translation between cameras. In the calibration process of our algorithm, the edge surface of the desk was used as a penetrating line.

of a translation vector as the convergence point of the candidate lines. To robustly determine the presence of a solution, we propose a “convergence voting” algorithm with the pseudocode presented in Algorithm 1.

The algorithm starts by calculating equidistant points  $\mathbf{m}_i$  for every pair of provided candidate lines. These equidistant points serve as potential solutions for the translation vector. From  $\mathbf{m}_1$  to  $\mathbf{m}_N$ , for each equidistant point, a set  $\mathcal{K}_i$  consisting of lines located within a certain distance  $\epsilon_d$  is defined. After obtaining the set for all points, if the number of elements in the maximum inlier set  $\mathcal{I}$  surpasses a predefined threshold, the translation vector is considered to converge as depicted in Fig. 4. Then, we determine whether to finish the calibration depending on the residual of the cost function (8). If the residual is not small enough, only the line pair included in the  $\mathcal{I}$  is incorporated into the quadratic system. The algorithm repeats the process until two conditions are satisfied, as shown in Fig. 3.

#### IV. EXPERIMENT

##### A. Evaluation Metrics

Our experiments were conducted on a setup running Ubuntu 20.04, powered by an Intel Core i7-12700@2.1 GHz CPU and an NVIDIA GeForce RTX 3080 GPU. We performed a two-fold evaluation to demonstrate the effectiveness of PeLiCal. The first evaluation focused on identifying subtle changes in pose by using an experimental setup as illustrated in Fig. 5. In the next phase, we benchmarked the precision of PeLiCal against other algorithms in a stereo setting.

We compared our approach with three calibration methods: Kalibr [17], ROS Calibrator, which is re-implementation of [18] and CamMap [19], an extrinsic calibration strategy designed to align maps formulated by ORB-SLAM3 [20]. In real-world settings, obtaining a definitive ground truth for calibration is often challenging. As a result, our initial experiment aimed to evaluate the accuracy of the pose changes detected by our algorithm, avoiding the need for such ground truth. Likewise, due to this inherent challenge, our comparison against other techniques was based on the criteria proposed by Perez et al. [9].

1) *Verification of Pose Variance Estimation:* Each of the RealSense D435i cameras is mounted on an aluminum structure in our device. The target camera features a protractor to




Fig. 6: The reconstructed point cloud representing the checkerboard plane, denoted as  $\pi_1$  and  $\pi_2$ , is merged into a singular coordinate using the predicted pose. The value  $l$  is derived by averaging the ratios of  $L_1$  to  $L_6$  for the number of squares in a row, then determining the discrepancy from the actual length of 108 mm (left). The variable  $d$  denotes the disparity between  $d_1$  and  $d_2$ , which correspond to the distances of  $\pi_1$  and  $\pi_2$  from the origin, respectively. Meanwhile,  $\theta$  represents the angle formed between the normals of the two planes (right).

assess rotational movements, while the source camera can be moved along the x-axis, determining translational changes with the ruler in place. While manual adjustments to the angle and spacing of the cameras may not result in the exact anticipated pose disparity, a precise algorithm should be able to discern and quantify the difference within a reasonable error margin.


From a base setup where both cameras faced forward with a 20 cm separation, we evaluated poses in 30 different scenarios. Within these tests, the angles varied between  $0^\circ$  and  $80^\circ$ , increasing in  $20^\circ$  steps, and the distances extended from 20 cm to 45 cm, with a progression of 5 cm at each step. Following this, we analyzed the accuracy of a 5 cm translational change along the x-axis at a constant angle and an angular adjustment of  $20^\circ$  at a set distance. The error for the  $i$ th rotation or translation, keeping the alternate parameter unvarying at either  $X_{cm}$  or  $\theta^\circ$ , is computed as follows:

$$\begin{aligned} \theta_{e_i}^{X_{cm}} &= \left\| (\boldsymbol{\theta}_{i+1}^{X_{cm}} - \boldsymbol{\theta}_i^{X_{cm}}) - (0, 20, 0)^T \right\|_2, \\ X_{e_i}^{\theta^\circ} &= \left| \left\| \mathbf{X}_{i+1}^{\theta^\circ} - \mathbf{X}_i^{\theta^\circ} \right\|_2 - 5 \right|. \end{aligned} \quad (16)$$

2) *Performance Comparison by Checkerboard:* First, we recorded a scene where a sizable calibration board was visible to both cameras. From each snapshot, we isolated planes, transformed them into 3D spatial coordinates, and combined them into one coordinate system using external parameters. Afterward, we identified the most distant corner points. By dividing the span between these corners by the count of squares, we derived  $l$ , representing the variance between the mean square dimension of the board and its actual size. An ideal estimation brings this value closer to zero. Moreover, we examined the deviations in the distance from the origin, denoted as  $d$ , and the angular disparity between the plane normals, represented by  $\theta$ . If mapped precisely, both metrics should ideally converge to zero, implying perfect alignment of the two planes. Detailed illustration of the metrics is provided in Fig. 6.

Fixed Distance	Rotation Variation (°)				Fixed Angle	Distance Variation (cm)				
	0→20	20→40	40→60	60→80		20→25	25→30	30→35	35→40	40→45
45 cm	0.6398	0.0707	0.9675	1.2077	80°	0.1697	0.1137	1.3918	0.6780	1.3741
40 cm	0.7586	0.1131	0.3187	0.2768	60°	0.1584	1.0096	1.1799	1.3456	0.4132
35 cm	0.1203	0.4511	0.8060	0.1947	40°	0.3156	0.1831	0.7042	0.6292	0.5332
30 cm	0.3974	0.2406	0.2435	0.3669	20°	0.0120	0.0093	0.3852	0.1017	0.1975
25 cm	0.1790	0.1944	0.4898	0.1901	0°	0.2690	0.2260	0.1161	0.1105	0.0595
20 cm	0.4953	0.4800	0.4810	1.1286						

(a) Error Measured for Rotation (°) and Translation (cm) Variation



(b) Error Heatmap

Fig. 7: Discrepancies between estimated pose variance and its actual value. As rotation and translation trend toward the top-right quadrant of the heatmap, error rates escalate. This elevation in error is primarily ascribed to a diminishing count of diverse feature matches as both cameras progressively capture different viewpoints.

### B. Calibration with Various Common FOV

The experiment outcomes are presented in the table with the heatmap in Fig. 7. In setups characterized by short distances and minimal angles where a shared FOV between the cameras is significant, the outcomes are closely aligned with the actual modifications in both parameters. Notably, in the first column with the minimum variation for each scenario, we observed a translational error of up to 0.3156 cm and an angular error of 0.7586°. However, as the angle exceeded 40° and the FOV no longer overlapped, the errors escalated due to the diminished number of line features. This is because as the angle between the camera’s central axis and the line expands, the accuracy of the depth sensor readings diminishes. Nonetheless, even under these challenging conditions, where traditional calibration techniques that require specialized tools falter, PeLiCal showed discrepancies reaching up to 1.2077° and 1.3918 cm.

### C. Stereo Calibration

In stereo camera systems, an increase in the baseline reduces the overlapping FOV between the cameras. Considering this, we evaluated the extrinsic calibration performance at baselines of 30 cm (near) and 45 cm (far) to examine the effects of a diminished shared FOV. For calibration purposes, Kalibr utilized an A3-sized AprilTag, while the ROS Calibrator employed an 8×5 checkerboard with squares measuring 45 mm each. We also tested CamMap in visually feature-rich indoor environments, such as labs or offices, selecting the setup that produced the the minimal errors.

The result is shown in Table. I with Fig. 8. The best results for each metric are emphasized in bold. In the first configuration, the ROS Calibrator most accurately determined the distance difference of the two planes from the origin. For the remaining metrics and configuration, PeLiCal predominantly outperformed others. Although CamMap effectively produced its trajectory and map over multiple sequences, it fails to estimate the given translational configuration. As a result, its associated values were not incorporated in the table. Also, while the accuracy of the ROS Calibrator and Kalibr diminished with the reduction of common FOV caused by a more increased baseline, our algorithm consistently maintained its accuracy without substantial deviation, attributing to its utilization of extended, distinct line features spanning across the images for calibration.

### V. CONCLUSION

We present a novel extrinsic calibration algorithm for RGB-D cameras leveraging line features. Our method solves the merged quadratic system derived from two conditions for initial pose estimation. Moreover, exploiting the constraints of transformed Plücker coordinates, translation candidates were represented as 3D lines. Existence of the optimal translation was validated through a convergence assessment through voting. The algorithm showed its ability to adapt to slight variations in pose, achieving performance superior to algorithms employing a calibration pattern.

TABLE I: Estimated rotation  $\Delta\mathbf{R}$  (roll, pitch, yaw), translation  $\Delta\mathbf{t}$  ( $x, y, z$ ) and metrics by merging two planes ( $l, d, \theta$ ) for each algorithm in two configurations. The results with substantial errors are replaced with a hyphen (—).

	Baseline	Near	Far
Kalibr	$\Delta\mathbf{R}$ (°)	[0.66, 0.08, -0.01]	[0.22, 0.98, 0.18]
	$\Delta\mathbf{t}$ (mm)	[296.46, 2.47, 1.36]	[441.94, 3.12, -7.96]
	[ $l, d, \theta$ ]	[1.28, 0.89, 1.09]	[0.88, 1.21, 2.21]
ROS	$\Delta\mathbf{R}$ (°)	[-0.56, 0.29, -0.08]	[-0.24, 0.55, -0.15]
	$\Delta\mathbf{t}$ (mm)	[297.79, -4.48, -2.46]	[451.54, -7.50, 0.00]
	[ $l, d, \theta$ ]	[0.69, <b>0.67</b> , 1.11]	[1.39, 1.19, 1.55]
CamMap	$\Delta\mathbf{R}$ (°)	[0.04, -0.19, 1.10]	[2.23, 0.15, -0.47]
	$\Delta\mathbf{t}$ (mm)	—	—
	[ $l, d, \theta$ ]	[—, —, 2.21]	[—, —, 2.11]
PeLiCal	$\Delta\mathbf{R}$ (°)	[0.00, 0.10, 0.00]	[0.00, -0.23, 0.00]
	$\Delta\mathbf{t}$ (mm)	[303.00, -3.14, 2.02]	[453.63, -5.33, 3.76]
	[ $l, d, \theta$ ]	[ <b>0.31</b> , 0.98, <b>0.77</b> ]	[ <b>0.07</b> , <b>0.97</b> , <b>0.65</b> ]




Fig. 8: Experimental results of stereo calibration (30cm (left), 45cm (right)) for each algorithm (Kalibr (red), ROS Calibrator (green), CamMap (cyan), PeLiCal (blue)).

## REFERENCES

- [1] S. Esquível, F. Woelk, and R. Koch, "Calibration of a multi-camera rig from non-overlapping views," in *Pattern Recognition: 29th DAGM Symposium, Heidelberg, Germany, September 12-14, 2007. Proceedings 29*. Springer, 2007, pp. 82–91.
- [2] B. Li, L. Heng, K. Koser, and M. Pollefeys, "A multiple-camera system calibration toolbox using a feature descriptor-based calibration pattern," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.* IEEE, 2013, pp. 1301–1307.
- [3] R. K. Kumar, A. Ilie, J.-M. Frahm, and M. Pollefeys, "Simple calibration of non-overlapping cameras with a mirror," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recog.* IEEE, 2008, pp. 1–7.
- [4] R. I. Hartley, "Camera calibration using line correspondences," in *Proc. DARPA Image Understanding Workshop*, 1993, pp. 361–366.
- [5] F. M. Mirzaei and S. I. Roumeliotis, "Globally optimal pose estimation from line correspondences," in *Proc. IEEE Intl. Conf. on Robot. and Automat.* IEEE, 2011, pp. 5581–5588.
- [6] B. Přebyl, P. Zemčík, and M. Čadík, "Camera pose estimation from lines using plücker coordinates," *arXiv preprint arXiv:1608.02824*, 2016.
- [7] C. Xu, L. Zhang, L. Cheng, and R. Koch, "Pose estimation from line correspondences: A complete analysis and a series of solutions," *IEEE Trans. Pattern Analysis and Machine Intell.*, vol. 39, no. 6, pp. 1209–1222, 2016.
- [8] C. Zhu, Z. Zhou, Z. Xing, Y. Dong, Y. Ma, and J. Yu, "Robust plane-based calibration of multiple non-overlapping cameras," in *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE, 2016, pp. 658–666.
- [9] A. Perez-Yus, E. Fernandez-Moral, G. Lopez-Nicolas, J. J. Guerrero, and P. Rives, "Extrinsic calibration of multiple rgb-d cameras from line observations," *IEEE Robot. and Automat. Lett.*, vol. 3, no. 1, pp. 273–280, 2017.
- [10] R. Pautrat, I. Suárez, Y. Yu, M. Pollefeys, and V. Larsson, "Gluestick: Robust image matching by sticking points and lines together," *arXiv preprint arXiv:2304.02008*, 2023.
- [11] L. Zhou, D. Koppel, and M. Kaess, "A complete, accurate and efficient solution for the perspective-n-line problem," *IEEE Robot. and Automat. Lett.*, vol. 6, no. 2, pp. 699–706, 2020.
- [12] F. M. Mirzaei and S. I. Roumeliotis, "Optimal estimation of vanishing points in a manhattan world," in *Proc. IEEE Intl. Conf. on Comput. Vision.* IEEE, 2011, pp. 2454–2461.
- [13] L. Zhou, J. Ye, and M. Kaess, "A stable algebraic camera pose estimation for minimal configurations of 2d/3d point and line correspondences," in *Asian Conference on Computer Vision (ACCV)*, 2019, pp. 273–288.
- [14] X. Zuo, X. Xie, Y. Liu, and G. Huang, "Robust visual slam with point and line features," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2017, pp. 1775–1782.
- [15] L. Zhou, S. Wang, and M. Kaess, "A fast and accurate solution for pose estimation from 3d correspondences," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2020, pp. 1308–1314.
- [16] J. Levinson, C. Esteves, K. Chen, N. Snavely, A. Kanazawa, A. Rostamizadeh, and A. Makadia, "An analysis of svd for deep rotation estimation," *Advances in Neural Information Processing Sys. Conf.*, vol. 33, pp. 22 554–22 565, 2020.
- [17] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2016, pp. 4304–4311.
- [18] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Analysis and Machine Intell.*, vol. 22, no. 11, pp. 1330–1334, 2000.
- [19] J. Xu, R. Li, L. Zhao, W. Yu, Z. Liu, B. Zhang, and Y. Li, "Cammmap: Extrinsic calibration of non-overlapping cameras based on slam map alignment," *IEEE Robot. and Automat. Lett.*, vol. 7, no. 4, pp. 11 879–11 885, 2022.
- [20] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam," *IEEE Trans. Robot. and Automat.*, vol. 37, no. 6, pp. 1874–1890, 2021.