

MEMROC: Multi-Eye to Mobile Robot Calibration

Davide Allegro¹, Matteo Terreran^{1†} and Stefano Ghidoni¹

Abstract—This paper presents MEMROC (Multi-Eye to Mobile Robot Calibration), a novel motion-based calibration method that simplifies the process of accurately calibrating multiple cameras relative to a mobile robot’s reference frame. MEMROC utilizes a known calibration pattern to facilitate accurate calibration with a lower number of images during the optimization process. Additionally, it leverages robust ground plane detection for comprehensive 6-DoF extrinsic calibration, overcoming a critical limitation of many existing methods that struggle to estimate the complete camera pose. The proposed method addresses the need for frequent recalibration in dynamic environments, where cameras may shift slightly or alter their positions due to daily usage, operational adjustments, or vibrations from mobile robot movements. MEMROC exhibits remarkable robustness to noisy odometry data, requiring minimal calibration input data. This combination makes it highly suitable for daily operations involving mobile robots. A comprehensive set of experiments on both synthetic and real data proves MEMROC’s efficiency, surpassing existing state-of-the-art methods in terms of accuracy, robustness, and ease of use. To facilitate further research, we have made our code publicly available¹.

I. INTRODUCTION

Mobile robots rely on one or more onboard sensors to perceive the environment and determine their location. Visual cameras are a popular choice due to their affordability and ability to capture extensive information from the scene: this led to a widespread use of such sensors in robotic navigation [1], [2] and autonomous driving [3], [4]. When the robot moves, onboard cameras frame the environment from multiple points of view, that can be related by means of pure vision algorithms or considering the robot odometry. Accurate alignment of visual data with respect to the robot reference frame is crucial for safe navigation [5] and for enhancing perception, localization, and the robot’s ability to interact effectively with its environment [6], [7]. To ensure a robot effectively utilizes visual information of its surroundings, it is required that the data collected by its cameras are expressed with respect to the robot reference frame—typically situated at the central point among its four wheels on the ground. This is why accurately calibrating the various cameras with respect to the mobile robot is essential.

Traditional calibration methods rely on specific infrastructures for precise sensor installation [8], [9], the need for environments equipped with numerous markers [10], [11], the requirement for supplementary sensors such as lidars to determine feature point correspondences among

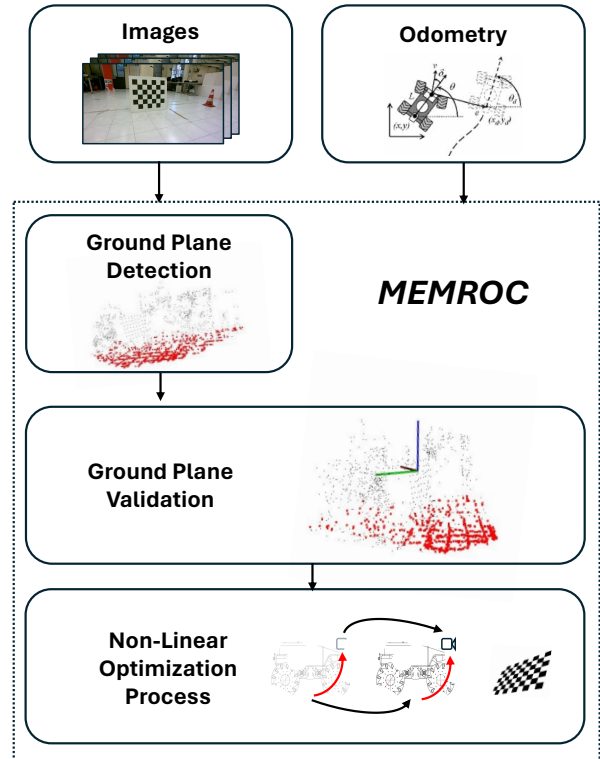


Fig. 1. MEMROC Overview. The system is structured into three primary components: the Ground Plane Detection module, responsible for identifying the ground plane; the Ground Plane Validation module, which ensures the ground plane’s alignment with the mobile robot; and the Motion-Based Calibration Process, the concluding stage that calibrates the system.

sensors [12], or the necessity of manual image collection processes [13]. These approaches are not suitable for frequent use, especially in dynamic real-world environment, which can cause minor camera displacements due to everyday operation, adjustments, or vibrations from robot movement. Frequent recalibration is necessary in such scenarios, and existing methods are often cumbersome and time-consuming. Furthermore, many approaches focus on calibrating cameras relative to each other requiring overlapping fields of view and neglecting the fundamental objective of expressing sensor data with the robot’s reference frame [14], [15], [16]. Motion-based calibration methods tackle these limitations by exploiting robot motion for calibrating cameras with respect to the robot. However, a key limitation arises from the inherent planar constraints of most mobile robots, preventing them from estimating the full 6-DoF of the camera’s extrinsic parameters solely through incremental movements [17], [18]. To overcome such limitation, Zuñiga’s method [19] estimates

† Corresponding Author matteo.terreran@dei.unipd.it
¹Intelligent Autonomous System Lab, Department of Information Engineering, University of Padova, Padua, Italy. {allegroDav, terreran, ghidoni}@dei.unipd.it
¹<https://github.com/davidea97/MEMROC.git>

all the 6-DoF through the ground plane detection, but it requires a large amount of data and a refinement procedure to be accurate; it also exhibits high sensitivity to noisy odometry data—a common issue due to odometry drift that limits its robustness [19]. Learning-based calibration methods are faster but might offer a limited accuracy and adaptability to changing environments, as they often require adherence to specific vehicle trajectories or the presence of well-defined road features for successful calibration, hindering their real-world applicability [20]. The limitations of existing methods highlight the need for faster, easier, and more robust sensor calibration techniques that consider the robot’s reference frame. This would enable robots to maintain optimal performance through frequent recalibration and ultimately enhance their ability to interact effectively with their surroundings.

This paper presents MEMROC: Multi-Eye to Mobile Robot Calibration. This is a novel motion-based calibration method for mobile robots which efficiently determines the full 6-DoF calibration parameters of multiple cameras relative to the robot incorporating a robust mechanism for the ground plane detection and its validation. Notably, the overlap of camera fields of view is *not* a prerequisite for this method; however, the optimization process is capable of taking advantage of the overlap, if available, to refine the geometric constraints between cameras, ensuring enhanced accuracy. After comprehensive experiments on both synthetic and real-world data, MEMROC demonstrated exceptional accuracy while requiring minimal data inputs. This was achieved through the observation of a calibration pattern, such as a checkerboard, from fewer than 20 viewpoints, and capturing the ground plane at least once. Unlike conventional approaches, MEMROC’s non-linear optimization technique shows remarkable robustness against noisy odometry data. This capability greatly increases the usability of autonomous mobile robots in daily environments.

Summarizing, this paper presents the following main contributions:

- 1) The concept of MEMROC, a novel motion-based calibration method for the estimation of the full 6-DoF pose of multiple cameras with respect to the robot;
- 2) A comprehensive performance evaluation of the proposed method by means of simulations and real-world experiments in both indoor and outdoor settings, demonstrating its robustness and accuracy;
- 3) A thorough comparative analysis focusing on how the precision of MEMROC is minimally dependent on the number of images used for calibration, showing its superior precision with minimal data requirements against other state-of-the-art methods;
- 4) An extensive, publicly available dataset, including 3000 synthetic and 3000 real images collected with a mobile robot, intended to facilitate the evaluation of future camera to robot calibration methods.

The remainder of the paper is organized as follows: Section II reviews state-of-the-art methods on camera calibration for mobile robots. Section III delves into the theoretical

foundation of our method, while Section IV outlines the dataset collection. Section V provides a detailed analysis of results on the collected dataset against state-of-the-art methods, in particular Sections V-A and V-B report outcomes in synthetic and real-world scenarios, respectively. Finally, Section VI concludes and outlines future research directions.

II. RELATED WORKS

Multi-sensor calibration on mobile robots can be categorized into three main approaches: appearance-based, motion-based, and learning-based methods. Appearance-based methods leverage prior knowledge about the environment. For instance, Gómez et al. [21] employed scene corners (orthogonal trihedrons) to calibrate a 2D laser scanner and a camera. Similarly, Gao et al. [12] proposed a multi-LiDAR calibration method using point constraints from retro-reflective targets placed in the scene. Additionally, Choi et al. [22] utilized the appearance of two orthogonal planes to determine the spatial offset between dual 2D LiDARs. However, appearance-based calibration methods are typically designed for calibrating specific sensor pairs, and applying them to a multi-sensor system with various combinations can become complex and impractical, as noted by Xie et al. [10]. More recently, Meyer et al. [20] proposed a learning-based method that leverages road features like lane markings or edges to calibrate yaw, and ground surface orientation for pitch and roll. However, this approach requires the vehicle to travel along a straight path with well-defined road markings. In contrast, Yan et al. [23] introduced SensorX2Car, which utilizes vanishing points and the horizon line to estimate the 3-DoF rotation between the sensor and the robot. This method, however, relies on the assumption that the vehicle’s orientation always aligns with its trajectory, which is not true in the case of cameras mounted on the side of the vehicles [13].

Motion-based approaches draw inspiration from the well-established hand-eye calibration problem [24]. They estimate the extrinsic parameters (relative poses) between sensors by analyzing their combined motions. This typically involves solving an equation of the form $AX = XB$, where A and B represent the motion of two sensors installed on the same robot and X represents the unknown transformation. Censi et al. [25] leverage this approach by combining wheel odometry and exteroceptive sensor data for multi-sensor calibration on mobile robots. Similarly, Kummerle et al. [26] propose an algorithm for visual-odometry calibration using motion data. Della Corte et al. [17] go a step further by estimating sensor time delays alongside extrinsic parameters within their motion-based calibration framework. However, a limitation of these methods is their inability to estimate all 6-DoF describing the sensor-to-robot transformations due to the unobservability of the z-coordinate [18]. Works like Heng et al.’s CamOdoCal [27] address this problem by estimating full 6-DoF for multiple cameras for identifying feature point correspondences between the current frame of each camera and the frames from every other camera’s historical data, resulting in extensive and time-consuming calibration pro-

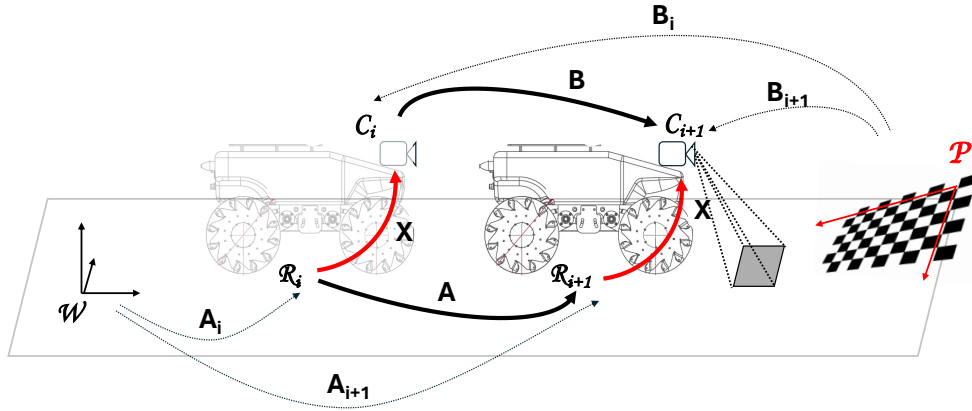


Fig. 2. Formulation of the motion-based calibration method. A mobile robot R equipped with a camera C moves from pose i to $i+1$ towards a calibration pattern P in the scene. A_i represents the odometry of the mobile robot with respect to the fixed frame W (i.e., the robot starting position), while B_i represents the camera pose with respect to the calibration pattern. X denotes the camera to robot transformation to be estimated.

cesses. Zuñiga’s RobotAutocalib [19] estimates full 6-DoF through ground plane detection, but suffers from inconsistencies due to its two-step approach (separately x , y , yaw and then the remaining 3-DoF). Furthermore, it needs a large number of images (often exceeding 50) to accurately estimate the ground plane using Structure-from-Motion (SfM) pipelines. Additionally, it only considers constraints between individual sensors and a reference, neglecting consistency across different sensor transformations.

III. METHOD

In this section, our MEMROC approach for full 6-DoF camera to robot calibration is presented and described. Our method follows a motion-based approach where calibration is obtained by minimizing the general $\|AX - XB\|^2$ cost function, introducing two main elements: a known calibration pattern fixed in the scene and a ground plane estimation module. The use of a known pattern (e.g. checkerboard) allows to overcome the inherent scale ambiguity in motion-based approaches, leading to several advantages: (i) more accurate camera motion estimation and (ii) a lower number of images needed in the optimization process compared to methods relying solely on visual features. Ground plane estimation, on the other hand, allows to address z-coordinate unobservability and achieve full 6-DoF calibration. It provides a set of measurements of the camera’s height relative to the ground which are then used as constraints in the optimization process.

A schematic overview of MEMROC is shown in Figure 1, representing its main modules. First, a set of pairs of robot poses and images is collected by moving the robot in front of a calibration pattern as depicted in Figure 2; robot poses A_i are obtained from odometry, while images allow to estimate camera poses B_i with respect to the calibration pattern by means of PnP (Perspective-n-Point) algorithm. Images are also used to compute 3D correspondences based on feature matching for ground plane detection and validation based on RANSAC and Singular Value Decomposition. Finally, robot and camera poses are used in the motion-based optimization

process with ground plane measurements as constraint.

A. Ground plane detection

The general idea is to estimate the ground plane equation from camera measurements, in order to compute its height from the ground and use such information in the optimization process. In particular, assuming a 3D representation of the scene in the camera frame is available (e.g., pointcloud), we identify the largest plane in such representation using the RANSAC (RANDOM SAMPLE CONSENSUS) method [28].

When considering RGB-D cameras, the 3D representation of the scene can be easily obtained from the pointcloud provided by the camera. However, the quality of the pointcloud can be very inaccurate in outdoor scenarios. Therefore, we consider an alternative approach to obtain 3D points of the scene based on feature matching between consecutive frames. In particular, given two consecutive images I_i and I_{i+1} , for each one of them a set of keypoints and descriptors is extracted (e.g., SIFT or SuperPoint [29]). Let $\mathcal{M} = \{(p_{1,i}, p_{1,i+1}), \dots, (p_{M,i}, p_{M,i+1})\}$ the set of matches between the two images computed by a feature matcher (e.g., LightGlue [30]). Based on the matches \mathcal{M} and the absolute camera poses B_i, B_{i+1} provided by the PnP algorithm, it is possible to compute the 3D coordinates of the matches by means of stereo triangulation using direct linear transform (DLT) [31]. Note that by using the absolute camera poses, the final result is a set of 3D points expressed with respect to the pattern reference frame \mathcal{P} , as depicted in Figure 1. Such points constitutes the 3D scene representation used for estimating the ground plane with RANSAC (red points in Figure 1), which provides the equation of the largest plane in the scene expressed in the general equation form

$$ax + by + cz + d = 0 \quad (1)$$

where a, b and c are the components of the normal vector $\vec{n} = (a, b, c)$ perpendicular to the plane and d is the minimum distance from the plane, which also represent the height of the camera from the ground plane to be used in the constrained calibration process.

B. Ground plane validation

The ground plane detection algorithm described above aims to compute the plane equation of the largest plane in the reconstructed scene. However, such plane could it be different from the ground plane, especially in crowded indoor scenarios (e.g., vertical walls, large flat objects) so a validation procedure to filter out wrong detections is needed.

When moving the robot, since the camera is rigidly attached to the robot, the robot and the camera perform the same trajectory but expressed in two different fixed reference frames (i.e., world and pattern respectively). Let $a = \{a_1, a_2, \dots, a_n\}$ and $b = \{b_1, b_2, \dots, b_n\}$ two sets of corresponding 3D points, obtained from the translation part of matrices A_i and B_i respectively. By means of Singular Value Decomposition (SVD) [32], it is possible to find a rigid transformation that optimally aligns the two sets in the least squares sense:

$$(R, t) = \underset{R \in SO(3), t \in \mathbb{R}^3}{\operatorname{argmin}} \sum_{i=1}^n \|(Rb_i + t) - a_i\|^2 \quad (2)$$

The obtained rigid transformation describes the roto-translation (R_P^W, t_P^W) between the world and the pattern reference frame. Although the estimated translation is affected by a bias term (it assumes that the 3D points are coincident, while a and b are only parallel planes which differ for an offset term), the estimated rotation can be used to transform the 3D points of the scene obtained Section III-A in the world reference frame, $p^W = R_P^W p^P$. By introducing such operation before the RANSAC-based plane detection, the estimated plane equation is then expressed in world coordinates and can be easily verify the correct alignment of the normal of the plane with the z-axis of the world reference frame. If the estimated plane is not parallel to the ground, it is rejected.

C. Camera to robot calibration

As illustrated in Figure 2, for each robot movement, we collect transformations A_i and B_i representing the robot's absolute pose at time step i relative to its initial position (i.e., world frame W), and the pose of the camera with respect to the fixed calibration pattern reference frame. However, robot poses obtained by odometry are generally noisy and affected by drift errors. This makes such data unsuitable for the calibration process, making it preferable to consider relative movements of the robot to limit the effect of drift in odometry. Therefore, in MEMROC, we aim to minimize the cost function defined as

$$c = \sum_{i=0}^{N-1} \|AX - XB\|^2, \quad (3)$$

where A represents the robot's incremental transformation within a 2D plane, B denotes the camera's incremental motion, and X is the camera to robot transformation matrix we aim to determine through the calibration process. Given N samples of incremental motion, we define our calibration

methodology through the objective of minimizing a specifically formulated cost function, detailed in equation (5).

$$c_{Zi} = \chi_i \|[T_C^R]_z - \bar{z}_i\|^2 \quad (4)$$

$$c = \sum_{i=0}^{N-1} \left(\left\| T_{R_{i+1}}^{R_i} T_C^R - T_C^R T_{C_{i+1}}^{C_i} \right\|^2 + c_{Zi} \right) \quad (5)$$

In this formulation, $T_{R_{i+1}}^{R_i}$ and $T_{C_{i+1}}^{C_i}$ symbolize the incremental motions of the robot and the camera, respectively, capturing their transition from one state to the next. Our objective is to accurately determine the transformation T_C^R , representing the roto-translation between the camera and the robot. An additional term c_{Zi} in the cost function aims to reduce the difference between the z-axis coordinate of the desired transformation and the camera's elevation, \bar{z}_i , as identified through ground plane detection (i.e., $\bar{z}_i = d$). The variable χ_i , which can either be 0 or 1, signifies whether the i^{th} pose yields a valid measure of the camera's height above the ground plane. This approach allows for the precise estimation of the camera's 6-DoF pose relative to the robot, moving beyond simple ground plane detection to integrate a comprehensive set of spatial data into the calibration effort. Remarkably, a single z_i measurement suffices for this calibration. Moreover, this optimization process calibrates each camera independently with respect to the robot, thereby eliminating the necessity for overlapping fields of view among sensors. This strategy not only simplifies the calibration process but also expands its versatility to support various sensor arrangements without sacrificing precision.

D. Multi-sensor joint optimization

When multiple cameras are mounted on the vehicle, each camera can be calibrated independently by means of the optimization process described in Section III-C. However, in the case that the cameras have overlapping fields of view, they can simultaneously detect the calibration pattern providing additional information on the relative pose between them. In such a scenario, it is possible to include mutual relations between cameras in the cost function to optimize not only the extrinsic parameters of each camera relative to the robot frame, but also the relative poses between pairs of cameras.

Consider the scheme in Figure 2 and assume that two cameras C_1 and C_2 are mounted on the vehicle in the front direction. When the robot moves from pose i to pose $i+1$, if both cameras can detect the calibration pattern, the relation between their incremental motion can be expressed as:

$$T_{C_{1,i+1}}^{C_{1,i}} T_{R_{i+1}}^{R_i} T_{C_2}^R = T_{R_{i+1}}^{R_i} T_{C_2}^R T_{C_{2,i+1}}^{C_{2,i}} \quad (6)$$

which can be easily reformulated in the general $AX = XB$ equation by assuming $X = T_{R_{i+1}}^{R_i} T_{C_2}^R = (T_{C_1}^{R_i})^{-1} T_{C_2}^R$. Note that such unknown matrix can be expressed as the product of the two camera to robot transformations $T_{C_1}^R$ and $T_{C_2}^R$ to be estimated in the calibration process, therefore providing an additional constraint to be exploited in the optimization.

Let denote by \mathcal{S}_i the set of camera pairs (j, k) which detects the calibration pattern at step i , and $\mathbb{1}_{\mathcal{S}_i}$ the corresponding indicator function. With such notation, the constraint describing the joint detection by multiple M cameras at step i can be expressed as:

$$c_{Ji} = \sum_{j=1}^M \sum_{k=j}^M \mathbb{1}_{\mathcal{S}_i} \left\| T_{C_{j,i+1}}^{C_{j,i}} T_R^{C_j} T_{C_k}^R - T_R^{C_j} T_{C_k}^R T_{C_{k,i+1}}^{C_{k,i}} \right\|^2 \quad (7)$$

Finally, including the joint constraint (7) in the cost function (5), the overall cost function can be expressed as:

$$c = \sum_{i=0}^{N-1} \left(\sum_{j=1}^M \left(\left\| T_{R_{i+1}}^{R_i} T_C^R - T_C^R T_{C_{j,i+1}}^{C_{j,i}} \right\|^2 + c_{Zi} \right) + c_{Ji} \right) \quad (8)$$

In the following, Joint-MEMROC will be used to denote the calibration process with joint optimization. By exploiting the mutual constraints between camera poses introduced by overlapping fields of view, our extensive experiments will demonstrate that Joint-MEMROC significantly enhances calibration accuracy. This improvement comes from incorporating additional geometric information derived from overlapping areas, enriching the calibration process by providing a more detailed understanding of the sensor configuration.

IV. DATASET

We collected a dataset² comprising 3000 synthetic images and 3000 real images, to comprehensively validate the proposed method for calibrating a multi-camera system fixed on a mobile robot. The synthetic images were generated using the Gazebo simulation toolbox, wherein we utilized a Husky Unmanned Ground Vehicle (UGV) from Clearpath Robotics³ equipped with three simulated Intel RealSense Depth D455 sensors. We conducted 10 independent image acquisition runs across diverse environments, each featuring distinct background and floor textures, as shown in Figure 3. For each run, the mobile robot was moved in 100 different poses, at varying distances and orientations in front of a 5×4 checkerboard pattern with 10 cm squares. During each pose, all three sensors simultaneously captured images, their corresponding 3D point clouds, and the robot's odometry (i.e., its position and orientation relative to the starting position). The synthetic dataset includes ground truth information consisting of:

- The geometric transformations of each sensor relative to the mobile robot's reference frame;
- The relative rototranslations among the various sensors of the multi-camera system.

The real part of the dataset consists of images captured by the mobile robot's three onboard cameras, all front-facing, from Robotnik⁴. The robot is equipped with two Intel RealSense LiDAR L515 sensors and one Intel RealSense

Depth D455 sensor, each designed with overlapping fields of view (FoV). While D455 cameras can be used both outdoor and indoor, L515 cameras are designed especially for indoor applications and provides very low-quality point cloud when used outside; in the real dataset we consider both sensor to prove the robustness and generality of our approach.

The real-world data consists of images captured in 10 diverse scenarios: 6 indoor and 4 outdoor environments. We collected the data by driving the mobile robot in 100 poses at various distances and orientations with respect to a 6×5 checkerboard with 11 cm squares. During data acquisition, all three sensors captured images simultaneously together with their corresponding 3D point clouds and the robot's odometry.

The indoor data encompasses two distinct environments showcasing variations in ground surfaces, background textures, and the inclination and positioning of the calibration pattern. This diversity mainly aims to challenge the calibration methods with different visual features. The outdoor data focuses on capturing images with diverse backgrounds and varying checkerboard inclinations. This data is particularly valuable for testing the robustness of calibration methods against potential lighting variations and when the robot's odometry becomes very noisy due to the inherent drift experienced on uneven terrains.

Crucially, for both the synthetic and real datasets, we maintained the relative positions of the cameras throughout the entire data collection process. This allows us to evaluate the consistency of calibration methods across multiple runs.

This dataset is valuable for comparing various multi-camera calibration methods. The availability of ground truth information enables to assess how well they perform with or without the calibration patterns, their accuracy with varying image quantities and their effectiveness in different scenarios.

Furthermore, in our dataset a checkerboard pattern is visible in all images: it can be used not only for calibration but also as ground truth for the evaluation of visual odometry methods [6], [33]. To the best of our knowledge, no other publicly available dataset offers this combination of features, making it a beneficial resource for the research community.

V. EXPERIMENTS AND RESULTS

We conducted several experiments on the collected data to demonstrate the effectiveness and robustness of our calibration method in practice. First, we validated our method using synthetic data, whose ground truth parameters are known (Section V-A). Next, we assessed the calibration accuracy in the real-world environments (Section V-B). We considered both datasets for comparing the proposed methods against the state-of-the-art calibration approaches, namely the motion-based method proposed by Zuñiga et al. [19] and the learning-based method proposed by Yan et al. [23]. They both calibrate sensors with respect to the vehicle reference frame, but the latter calibration method only provides rotation calibration parameters, therefore only rotation errors will be reported in the resulting tables. To compare calibration performance, we also assessed the accuracy of estimating

²<https://github.com/davidea97/MEMROC.git>

³<https://clearpathrobotics.com/>

⁴<https://robotnik.eu/>



Fig. 3. Set of 4 images representing 4 different scenarios where the dataset was collected. The first two on the left show two different real-world scenarios, the other two simulated scenarios.

the transformations between the various sensors in the multi-camera system. To establish a reliable baseline for our measurements in real-world calibration parameter experiments, we utilized the stereo calibration method developed by Li et al. [34] as the ground truth reference. This method will also be used as a comparison in the simulated dataset to show its correctness and accuracy.

A. Experiments on synthetic dataset

The primary goal of using synthetic data is to evaluate the correctness and accuracy of our calibration method by directly comparing it to the ground truth. This allows us to compute the translation error across the three axes and the rotation error around the three rotation axes, expressed in terms of roll, pitch, and yaw. This metric was used to evaluate first the precision of the estimated camera-to-robot transformation and then we assess the consistency of the transformations among cameras by evaluating the camera-to-camera rototranslation error. This helps us identify any potential cumulative errors that might arise during the calibration process.

Tables I and II report the average error for each calibration parameter of camera-to-robot and camera-to-camera transformations, respectively. These errors represent the mean values calculated across all 10 synthetic datasets.

TABLE I
AVERAGE ERROR IN CAMERA-TO-ROBOT TRANSFORMATION

Method	Translation parameters			Rotation parameters		
	t_x [cm]	t_y [cm]	t_z [cm]	r_x [deg]	r_y [deg]	r_z [deg]
MEMROC	0.61	0.25	0.15	0.08	0.13	0.16
Joint-MEMROC	0.66	0.16	0.14	0.14	0.14	0.15
RobotAutocalib [19]	0.92	0.79	0.92	1.87	1.32	1.84
SensorX2Car [23]	—	—	—	0.26	0.82	1.22

TABLE II
AVERAGE ERROR IN CAMERA-TO-CAMERA TRANSFORMATION

Method	Translation parameters			Rotation parameters		
	t_x [cm]	t_y [cm]	t_z [cm]	r_x [deg]	r_y [deg]	r_z [deg]
MEMROC	0.42	0.23	0.87	0.11	0.18	0.15
Joint-MEMROC	0.20	0.25	0.78	0.22	0.23	0.14
RobotAutocalib [19]	2.58	4.99	3.72	2.52	3.09	1.68
SensorX2Car [23]	—	—	—	0.44	0.21	2.40
Li [34]	0.26	0.27	0.34	0.12	0.14	0.17

Our method achieves higher calibration accuracy for both translation and rotation errors compared to existing state-

of-the-art approaches. Notably, RobotAutocalib struggles with camera-to-camera transformations, possibly due to its disregard for relative constraints existing among cameras. Conversely, the Li stereo calibration method performs well in this specific scenario. Yan’s method, despite being learning-based, demonstrates strong accuracy and generalizability across various scenarios, particularly for rotation parameters.

Analysis on image quantity: the synthetic dataset enables us to evaluate the impact of the number of images on the effectiveness of our multi-sensor calibration method.

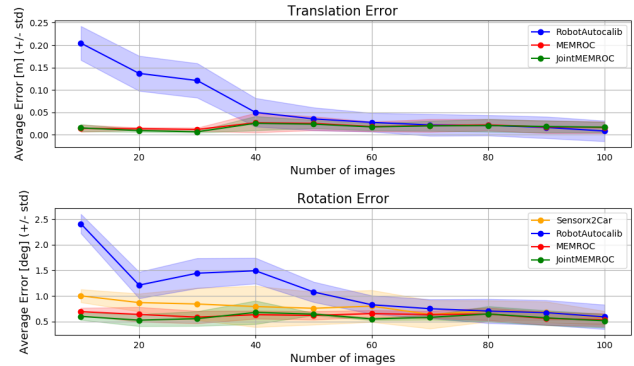


Fig. 4. Impact of image quantity on calibration methods, where the lines represent the average errors and the shaded areas indicate the standard deviation, reflecting the variability and consistency of each method.

Figure 4 illustrates the substantial impact of image quantity on the effectiveness of calibration methods. From this point of view MEMROC shows remarkable robustness and accuracy even with minimal data available for the calibration, outperforming the motion-based method by Zuñiga’s algorithm and the learning-based Yan’s approach.

Analysis with noisy odometry: the synthetic dataset enabled us to test the robustness and efficiency of our method in challenging scenarios, specifically where odometry measurements are compromised by noise. To thoroughly evaluate robustness in such conditions, we introduced Gaussian noise $N(0, \sigma_t)$ for translation, with $\sigma_t = \lambda \cdot 0.2$ cm and Gaussian noise $N(0, \sigma_r)$ for rotation, with $\sigma_r = \lambda \cdot 0.01$ rad where $\lambda \in [0, 10]$. Noise primarily affects odometry measurements

in the x and y axes for translation, and in the z-axis for rotation, reflecting the 2D nature of the motion estimates. Figure 6 illustrates how increasing noise levels in odometry measurements affect the average translation and rotation errors in camera-to-robot transformations. As a result, the comparison in this figure is limited to motion-based calibration methods, since they rely on odometry data.

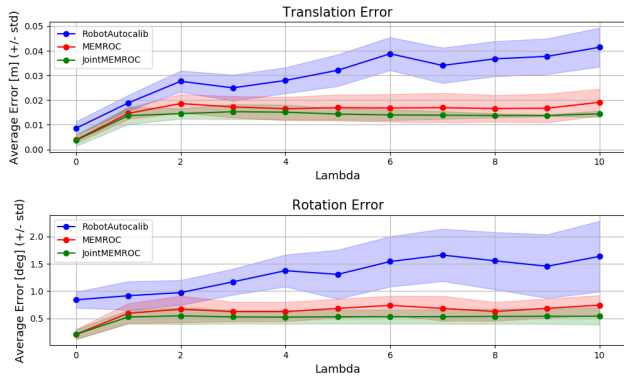


Fig. 5. Evaluation of motion-based calibration method with incremental noise in measurements provided by odometry, with $\lambda \in [0, 10]$. The lines represent the average errors and the shaded areas indicate the standard deviation.

Figure 6 demonstrates that the proposed motion-based method exhibits significantly higher accuracy and robustness compared to the RobotAutocalib method under noisy conditions. Notably, our method achieves translation errors consistently below 2 cm and rotation errors below 1° . This highlights its suitability for real-world scenarios with noisy measurements. Our experiments with synthetic data demonstrate the efficiency of the proposed method in two key ways. First, they enable the calibration of all 6-DoF for both camera-to-robot and camera-to-camera transformations. Second, they achieve high accuracy, surpassing the motion-based method by Zuñiga and the learning-based by Yan. The ability to calibrate with just a few images highlights that the system is particularly suited for real-world scenarios where a fast and user-friendly approach is crucial.

B. Experiments on real-world scenarios

To comprehensively assess the effectiveness of our proposed calibration methods, we conducted a series of experiments in real-world indoor and outdoor scenarios. These experiments allowed us to analyze not only the precision of the calibration but also its robustness under non-ideal conditions. This includes situations where odometry measurements suffer from drift errors. Notably, these real-world tests were crucial for demonstrating the consistency of the results across diverse environments with varying visual-features and floor types, which can significantly impact odometry accuracy. We evaluated our calibration method against state-of-the-art approaches on 10 real-world datasets, each containing 300 images. Fig. 6 shows for simplicity only the average value

and standard deviation obtained by MEMROC and state-of-the-art methods across all runs for each calibration parameter (t_x, t_y, t_z for translation and r_x, r_y, r_z for rotation) of the Intel RealSense LiDAR L515 sensor.

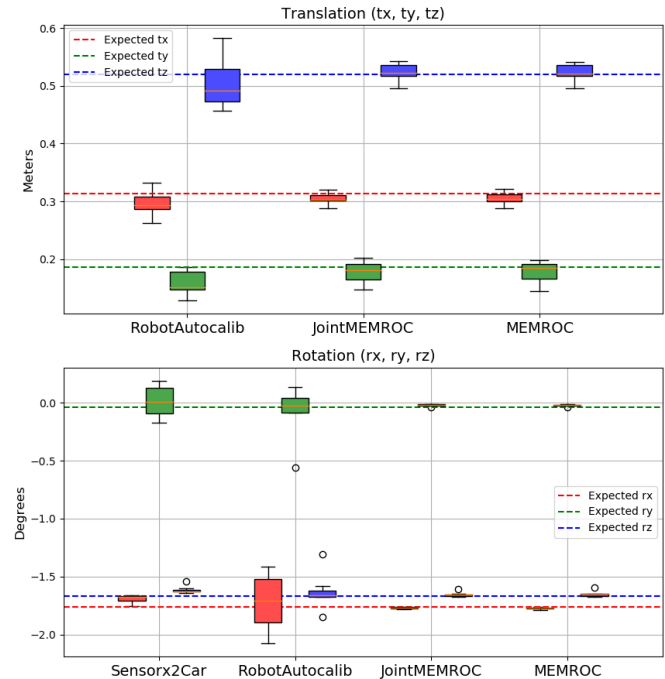


Fig. 6. Calibration parameters with average and standard deviation of the Intel RealSense LiDAR L515 sensor estimated by different calibration methods. The expected values were measured by using an AprilTag fixed at a precise distance from the mobile robot in front of the sensors.

Our experiments demonstrate that MEMROC consistently outperform existing approaches. It achieves higher accuracy on average, and it exhibits greater consistency across multiple runs of the algorithm, resulting in reliable and precise measurements.

TABLE III
AVERAGE ERROR IN CAMERA-TO-CAMERA TRANSFORMATION

Method	Translation parameters			Rotation parameters		
	t_x [cm]	t_y [cm]	t_z [cm]	r_x [deg]	r_y [deg]	r_z [deg]
MEMROC	1.97	1.15	0.34	0.79	1.02	0.72
Joint-MEMROC	0.48	0.56	0.33	0.74	0.83	0.66
RobotAutocalib [19]	2.41	4.42	2.43	3.83	4.82	4.22
SensorX2Car [23]	-	-	-	2.95	1.99	2.75

Table III presents the average errors for relative calibration parameters between the three cameras. Our method achieves superior calibration accuracy compared to both state-of-the-art motion-based and learning-based approaches. It should be emphasized that the MEMROC achieves even higher precision through joint optimization of all sensors, leveraging their overlapping fields of view. This technique is particularly effective for multi-sensor calibration.

VI. CONCLUSIONS

In this paper we presented MEMROC, a novel method for calibrating multiple sensors with respect to a mobile robot. Utilizing robust ground plane detection, MEMROC

addresses z-unobservability in planar motion, enabling the estimation of cameras' full 6-DoF pose with just a few images of a simple calibration pattern. Moreover, when sensors have overlapping fields of view, MEMROC exploits this condition optimizing also the geometric constraints existing between sensors. A comprehensive analysis of our method against state-of-the-art methods demonstrates that it is highly accurate and more robust than other methods in non-ideal conditions, when the odometry is affected by noise. Moreover, MEMROC proves to be suitable for daily use and for real-world scenarios as it requires only minimal data to achieve higher calibration accuracy. Currently, we are investigating how to completely remove the checkerboard while maintaining the same performance level. We also plan to extend the technique to sensors like IMU and Laser 2D.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101006732 (DrapeBot).

REFERENCES

- [1] W. Yuan, Z. Li, and C.-Y. Su, "Multisensor-based navigation and control of a mobile service robot," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 51, no. 4, pp. 2624–2634, 2019.
- [2] K.-W. Chiang, G.-J. Tsai, H. Chang, C. Joly, and N. Ei-Sheimy, "Seamless navigation and mapping using an ins/gnss/grid-based slam semi-tightly coupled integration scheme," *Information Fusion*, vol. 50, pp. 181–196, 2019.
- [3] M. Henein, J. Zhang, R. Mahony, and V. Ila, "Dynamic slam: The need for speed," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2123–2129.
- [4] D. Fernandes, A. Silva, R. Névoa, C. Simões, D. Gonzalez, M. Guevara, P. Novais, J. Monteiro, and P. Melo-Pinto, "Point-cloud based 3d object detection and classification methods for self-driving applications: A survey and taxonomy," *Information Fusion*, vol. 68, pp. 161–191, 2021.
- [5] D. Fusaro, E. Olivastri, D. Evangelista, M. Imperoli, E. Menegatti, and A. Pretto, "Pushing the limits of learning-based traversability analysis for autonomous driving on cpu," in *International Conference on Intelligent Autonomous Systems*. Springer, 2022, pp. 529–545.
- [6] F. Zheng and Y.-H. Liu, "Visual-odometric localization and mapping for ground vehicles using se (2)-xyz constraints," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3556–3562.
- [7] G. Evangelidis and B. Micusik, "Revisiting visual-inertial structure-from-motion for odometry and slam initialization," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1415–1422, 2021.
- [8] Z. Koppanyi and C. Toth, "Experiences with acquiring highly redundant spatial data to support driverless vehicle technologies," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 4, pp. 161–168, 2018.
- [9] M. Horn, T. Wodtke, M. Buchholz, and K. Dietmayer, "Extrinsic infrastructure calibration using the hand-eye robot-world formulation," in *2023 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2023, pp. 1–8.
- [10] Y. Xie, R. Shao, P. Guli, B. Li, and L. Wang, "Infrastructure based calibration of a multi-camera and multi-lidar system using apriltags," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 605–610.
- [11] G. Antonelli, F. Caccavale, F. Grossi, and A. Marino, "Simultaneous calibration of odometry and camera for a differential drive mobile robot," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 5417–5422.
- [12] C. Gao and J. R. Spletzer, "On-line calibration of multiple lidars on a mobile vehicle platform," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 279–284.
- [13] Y. Wang, W. Jiang, K. Huang, S. Schwertfeger, and L. Kneip, "Accurate calibration of multi-perspective cameras from a generalization of the hand-eye constraint," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1244–1250.
- [14] G. Yan, Z. Liu, C. Wang, C. Shi, P. Wei, X. Cai, T. Ma, Z. Liu, Z. Zhong, Y. Liu *et al.*, "Opencalib: A multi-sensor calibration toolbox for autonomous driving," *Software Impacts*, vol. 14, p. 100393, 2022.
- [15] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1280–1286.
- [16] J. Jiao, Y. Yu, Q. Liao, H. Ye, R. Fan, and M. Liu, "Automatic calibration of multiple 3d lidars in urban environments," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 15–20.
- [17] B. Della Corte, H. Andreasson, T. Stoyanov, and G. Grisetti, "Unified motion-based calibration of mobile multi-sensor platforms with time delay estimation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 902–909, 2019.
- [18] C. X. Guo, F. M. Mirzaei, and S. I. Roumeliotis, "An analytical least-squares solution to the odometer-camera extrinsic calibration problem," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3962–3968.
- [19] D. Zuñiga-Noël, J.-R. Ruiz-Sarmiento, R. Gomez-Ojeda, and J. Gonzalez-Jimenez, "Automatic multi-sensor extrinsic calibration for mobile robots," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2862–2869, 2019.
- [20] S. W. Meyer, H. Chen, and D. M. Bevy, "Automatic extrinsic rotational calibration of lidar sensors and vehicle orientation estimation," *IFAC-PapersOnLine*, vol. 54, no. 20, pp. 424–429, 2021.
- [21] R. Gomez-Ojeda, J. Briales, E. Fernandez-Moral, and J. Gonzalez-Jimenez, "Extrinsic calibration of a 2d laser-rangefinder and a camera based on scene corners," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3611–3616.
- [22] D.-G. Choi, Y. Bok, J.-S. Kim, and I. S. Kweon, "Extrinsic calibration of 2-d lidars using two orthogonal planes," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 83–98, 2015.
- [23] G. Yan, Z. Luo, Z. Liu, and Y. Li, "Sensorx2car: Sensors-to-car calibration for autonomous driving in road scenarios," *arXiv preprint arXiv:2301.07279*, 2023.
- [24] R. Horaud and F. Dornaika, "Hand-eye calibration," *The international journal of robotics research*, vol. 14, no. 3, pp. 195–210, 1995.
- [25] A. Censi, A. Franchi, L. Marchionni, and G. Oriolo, "Simultaneous calibration of odometry and sensor parameters for mobile robots," *IEEE Transactions on Robotics*, vol. 29, no. 2, pp. 475–492, 2013.
- [26] R. Kümmerle, G. Grisetti, and W. Burgard, "Simultaneous calibration, localization, and mapping," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 3716–3721.
- [27] L. Heng, B. Li, and M. Pollefeys, "Camodocal: Automatic intrinsic and extrinsic calibration of a rig with multiple generic cameras and odometry," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1793–1800.
- [28] L. Li, F. Yang, H. Zhu, D. Li, Y. Li, and L. Tang, "An improved ransac for 3d point cloud plane segmentation based on normal distribution transformation cells," *Remote Sensing*, vol. 9, no. 5, p. 433, 2017.
- [29] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superpoint: Self-supervised interest point detection and description," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 224–236.
- [30] P. Lindenberger, P.-E. Sarlin, and M. Pollefeys, "LightGlue: Local Feature Matching at Light Speed," in *ICCV*, 2023.
- [31] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, ISBN: 0521540518, 2004.
- [32] O. Sorkine-Hornung and M. Rabinovich, "Least-squares rigid motion using svd," *Computing*, vol. 1, no. 1, pp. 1–5, 2017.
- [33] X. Wei, J. Lv, J. Sun, E. Dong, and S. Pu, "Gclo: Ground constrained lidar odometry with low-drifts for gps-denied indoor environments," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2229–2235.
- [34] B. Li, L. Heng, K. Koser, and M. Pollefeys, "A multiple-camera system calibration toolbox using a feature descriptor-based calibration pattern," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 1301–1307.