

# Automatic Target-Less Camera-LiDAR Calibration From Motion and Deep Point Correspondences

Kürsat Petek<sup>1\*</sup>, Niclas Vödisch<sup>1\*</sup>, Johannes Meyer<sup>1</sup>, Daniele Cattaneo<sup>1</sup>, Abhinav Valada<sup>1</sup>, Wolfram Burgard<sup>2</sup>

**Abstract**—Sensor setups of robotic platforms commonly include both camera and LiDAR as they provide complementary information. However, fusing these two modalities typically requires a highly accurate calibration between them. In this paper, we propose MDPCalib which is a novel method for camera-LiDAR calibration that requires neither human supervision nor any specific target objects. Instead, we utilize sensor motion estimates from visual and LiDAR odometry as well as deep learning-based 2D-pixel-to-3D-point correspondences that are obtained without in-domain retraining. We represent camera-LiDAR calibration as an optimization problem and minimize the costs induced by constraints from sensor motion and point correspondences. In extensive experiments, we demonstrate that our approach yields highly accurate extrinsic calibration parameters and is robust to random initialization. Additionally, our approach generalizes to a wide range of sensor setups, which we demonstrate by employing it on various robotic platforms including a self-driving perception car, a quadruped robot, and a UAV. To make our calibration method publicly accessible, we release the code on our project website at <http://calibration.cs.uni-freiburg.de>.

**Index Terms**—Calibration and Identification; Deep Learning Methods; Sensor Fusion

## I. INTRODUCTION

SENSOR fusion for robotic systems has been extensively investigated [1], [2] as it promises to efficiently combine complementary information from different modalities, e.g., to increase robustness in case of sensor failures [3] and towards weather conditions [4]. However, the effectiveness of fusion approaches depends heavily on the extrinsic calibration between the sensors such as cameras and LiDAR.

Due to the importance of the task, camera-LiDAR calibration has been widely studied by the research community. Previously proposed methods can generally be classified into target-based and target-less approaches. Approaches of the first category often rely on artificial patterns such as checkerboards [5], [6] and require manual labor or another kind of human supervision to associate 2D points from the image space with 3D points in the LiDAR point cloud [7]. While a

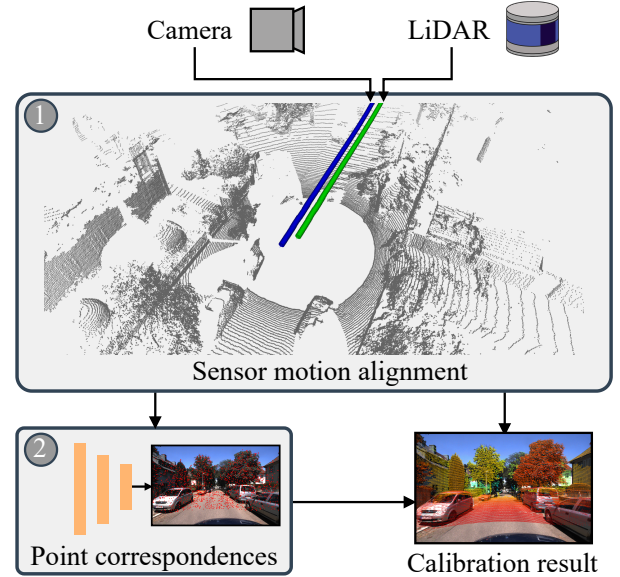


Fig. 1. Our proposed method, called MDPCalib, for camera-LiDAR calibration comprises two steps: We first initialize the extrinsic parameters by aligning the motion of both sensors. Afterward, we refine the calibration results by leveraging deep learning-based 2D-to-3D point correspondences.

substantial effort has gone into the detection of the target and the automation of the matching process [8], [9], calibration often still needs special data collection. Some target-less calibration methods intend to overcome this problem, e.g., by inferring the extrinsic transform from the sensor motion [10] or by matching vision-based structure-from-motion models with accumulated point clouds from the LiDAR [11]. Although these approaches are generally more widely applicable as they enable sensor calibration from normal robot operation, they often still require an initial set of parameters.

In this work, we propose the novel MDPCalib to eliminate this drawback by fully automating the calibration procedure. Given recorded data from normal robot operation, we employ visual and LiDAR odometry to generate two paths that can be aligned via non-linear optimization for coarse sensor registration. Afterward, we use the coarse parameters to initialize a learning-based 2D-to-3D point correspondence algorithm that outputs dense matches between the image and the point cloud spaces. In the final step, we jointly optimize with respect to both sensor motion and point correspondence, thereby effectively fusing the complementary information. Robust loss functions account for outlying observations from either source. In summary, the main contributions are:

- 1) We introduce MDPCalib for automatic target-less camera-LiDAR calibration that requires neither human

\* Equal contribution.

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. This work was funded by the German Research Foundation Emmy Noether Program grant No 468878300 and an academic grant from NVIDIA.

<sup>1</sup> Kürsat Petek, Niclas Vödisch, Johannes Meyer, Daniele Cattaneo, and Abhinav Valada are with the Department of Computer Science, University of Freiburg, Germany.

<sup>2</sup> Wolfram Burgard is with the Department of Engineering, University of Technology Nuremberg, Germany.

Digital Object Identifier (DOI): 10.1109/LRA.2024.3468090

initialization nor special data recording.

- 2) We propose to formulate extrinsic calibration as an optimization problem constrained by sensor motion and deep learning-based point correspondences.
- 3) We extensively demonstrate the general applicability of MDPCalib to both public and in-house datasets.
- 4) We release our code along with a detailed user guide on <http://calibration.cs.uni-freiburg.de>.

## II. RELATED WORK

Methods for extrinsic camera-LiDAR calibration can generally be categorized into target-based and target-less approaches. While the former have been well investigated for several years, recent advances in deep learning have enabled the rise of target-less calibration removing the need for explicit target objects. In this section, we present an overview of approaches from both categories.

*Target-Based Calibration:* Inspired by the estimation of camera parameters with a checkerboard pattern [12], Zhang *et al.* [6] were the first to propose using a similar target also for extrinsic calibration between a camera and a LiDAR. Since then, many different styles of patterns have been described to further optimize this calibration procedure both in terms of accuracy and applicability. For instance, Dhall *et al.* [5] exploit ArUco markers with known sizes to get accurate estimates of the corner points of the pattern in 3D, which are then registered to the corner point detected by the LiDAR using the ICP algorithm. Similarly, Kim *et al.* [9] fit points on a checkerboard detected by a camera to the corresponding plane in the LiDAR point cloud. However, not all target-based approaches rely on a checkerboard pattern or its variants to establish point correspondences. For example, both Velas *et al.* [13] and Guindel *et al.* [8] utilize wooden boards with holes to obtain point correspondences.

Similar to our method, Ou *et al.* [14] propose to formulate camera-LiDAR registration as a graph optimization problem. In particular, their method first extracts corner points in both modalities to perform an initial calibration obtained by a perspective-n-point (PnP) algorithm. This allows reprojecting the LiDAR points onto the image plane and computing the reprojection error, which is then used as a cost term in the graph-based formulation that can be efficiently solved using graph optimization methods [15].

*Target-Less Calibration:* Target-less calibration aims at performing camera-LiDAR registration without specifically designed target objects. For robotics, this opens an avenue for flexible and potentially online recalibration and enables applications to large fleets by reducing human supervision.

Correspondence-based methods replace the artificial targets with patterns that can be perceived in structured environments such as urban areas. For instance, Yuan *et al.* [16] and Yin *et al.* [17] match edges obtained from both images and LiDAR point clouds. Tu *et al.* [11] extract features of structure-from-motion (SfM) points of camera data and LiDAR points, followed by optimizing these correspondences jointly with camera intrinsics as well as camera and LiDAR

poses. Koide *et al.* [18] build on the commonly used normalized information distance that poses a distance metric between the image and projected LiDAR points to measure the amount of mutual information. Finally, Caselitz *et al.* [19] propose a method for determining the pose of an RGB camera with respect to a 3D point cloud generated from LiDAR data by matching geometric clues. On the other hand, correspondence-free methods often rely on leveraging output data of auxiliary tasks such as monocular depth prediction [20], [21] or sensor motion estimation [10], [17], [22]. Both Zhang *et al.* [10] and Yin *et al.* [17] match trajectories from visual and LiDAR odometry and obtain extrinsic parameters via optimization. The latter then utilize these parameters to initialize an edge-driven refinement stage. Finally, a more direct approach is enabled by exploiting deep learning-based correspondences between RGB images and LiDAR point clouds. Both RegNet [23] and CMRNext [24] involve training multiple CNNs on varying levels of decalibration and employ these networks during test time in a hierarchical manner. In detail, CMRNext frames the point-to-pixel matching problem as an optical flow estimation task. Building upon this approach, LCCNet [25] proposes the construction of a cost volume that stores matching costs. However, as it predicts the 6-DoF rigid-body transformation, it suffers from a high dependency on the training setup and is thus less generalizable.

In our proposed method, we combine the advantages of both correspondence-free and correspondence-based approaches by performing coarse initialization based on sensor motion followed by fine registration incorporating deep learning-based point correspondences [24]. Particularly, in contrast to previous works [17], we utilize robust cost functions and jointly optimize with respect to both sensor motion and point correspondences, accounting for their complimentary information and increasing robustness towards outlying observations.

## III. TECHNICAL APPROACH

In this section, we present our MDPCalib approach for automatic target-less camera-LiDAR calibration. As illustrated in Fig. 2, MDPCalib comprises two consecutive steps by combining a coarse initialization from sensor motion with a fine-tuning stage that takes learning-based correspondences between image pixels and LiDAR points into account. We first provide the relevant mathematical background, then introduce the general problem formulation, and finally give a detailed description of both registration steps.

### A. Mathematical Preliminaries

We interpret camera-LiDAR calibration as an optimization problem that determines the most likely transformation between the coordinate frames of a camera and a LiDAR given a set of sensor measurements. Mathematically, this can be represented by a conditional probability distribution:

$$P(x \mid x_0, z_{0:k}), \quad (1)$$

where  $x$  denotes the state vector with an initial guess  $x_0$  and  $z_{0:k}$  refers to a set of observations, i.e., sensor data. Instead of calculating the exact probability distribution, we perform

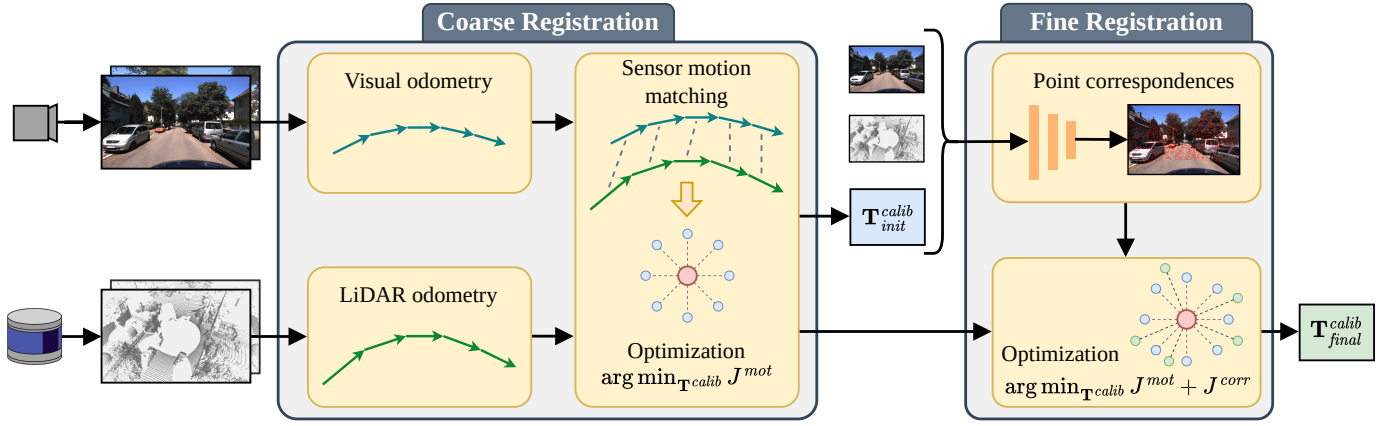


Fig. 2. Our proposed method for camera-LiDAR calibration processes two input streams of RGB images and 3D point clouds. The first step comprises a coarse registration based on sensor motion estimated with visual and LiDAR odometry. These motion estimates yield time-synchronized matches serving as constraints in an optimization problem. Given the obtained initial calibration parameters, a neural network is used to find 2D pixel to 3D point correspondences that result in additional constraints. The second step consists of joint optimization with respect to both sensor motion and point correspondences yielding the overall extrinsic calibration parameters.

maximum a posteriori (MAP) estimation assuming Gaussian distributions and independent and identically distributed measurements. This yields the optimal state  $x^*$ :

$$x^* = \arg \max_x P(x \mid x_0, z_{0:k}). \quad (2)$$

In practice, we solve for  $x^*$  by optimizing a robustified nonlinear squares problem of the form:

$$x^* = \arg \min_x \sum_i \rho_i(\|f_i(x, z_i)\|_2^2) \quad (3a)$$

$$= \arg \min_x \sum_i \rho_i(e_i(x, z_i)^T \Omega_i e_i(x, z_i)), \quad (3b)$$

where  $\rho_i$  denotes the robustifier and  $f_i$  refers to a cost function applied to observation  $z_i$ . For efficiency, we compute the squared Frobenius norm  $\|\cdot\|_2^2$  via vector multiplication of the induced error vectors  $e_i \in \mathbb{R}^{d \times 1}$  with a diagonal weighting matrix  $\Omega_i \in \mathbb{R}^{d \times d}$ . We utilize the Cauchy loss as robustifier due to its high tolerance to outliers in the observations [26], resulting in a significant increase in calibration accuracy.

To obtain the MAP estimate  $x^*$ , we leverage an optimization formulation [15] with a single state  $x$  and multiple constraints defined by the error terms  $e_i$ . In particular, we define the state  $x$  as the calibration parameters  $\mathbf{T}^{calib}$  and decompose the sum into three error types:

$$\mathbf{T}_{final}^{calib} = \arg \min_{\mathbf{T}^{calib}} \underbrace{\sum_{i=0}^{k-1} \dots}_{J^{rot}} + \underbrace{\sum_{i=k}^{2k-1} \dots}_{J^{trans}} + \underbrace{\sum_{i=2k}^{2k+m} \dots}_{J^{corr}}, \quad (4)$$

with  $k$  observations for sensor motion with rotational and translational costs  $J^{rot}$  and  $J^{trans}$  and  $m$  observations for the point correspondences cost  $J^{corr}$ .

### B. Coarse Registration

During the first step of MDPCalib, we perform coarse camera-LiDAR calibration by matching sensor motion. In particular, we use both vision- and LiDAR-based odometry to obtain poses  $\mathbf{P}_{0:k}^{cam}$  and  $\mathbf{P}_{0:k}^{lidar}$  capturing oriented positions of the camera and the LiDAR, respectively. To estimate the camera poses, we employ ORB-SLAM3 [27]. In contrast to

deep learning-based methods [28], [29], classical feature-based tracking approaches such as ORB-SLAM3 are more robust towards detecting lost tracks, which would lead to inconsistent constraints in the optimization problem. For the LiDAR, we perform consecutive scan matching using an adapted version of FAST-LIO2 [30] without the measurements of an inertial measurement unit (IMU). If not already done during the post-processing of the sensor output, the LiDAR scans are undistorted before being matched. Given time-synchronized sensor measurements, we now interpolate the poses of the LiDAR such that we obtain pose pairs of both sensors at each image timestamp. As detailed in Sec. III-D, we further transform the nearest point cloud to the same timestamp using the estimated LiDAR odometry as this data will be processed during the fine-tuning stage. Subsequently, we compute the pose difference between two consecutive poses yielding homogeneous transforms  $\mathbf{T}_i^{sensor}$  with  $sensor \in \{cam, lidar\}$ . Using these transforms, we solve the following equation to find the extrinsic calibration parameters  $\mathbf{T}^{calib}$ :

$$\mathbf{T}^{cam} \mathbf{T}^{calib} = \mathbf{T}^{calib} \mathbf{T}^{lidar}. \quad (5)$$

As derived by Shiu *et al.* [31], such an equation can be decomposed into solving for the rotational and translational components separately:

$$\mathbf{R}^{cam} \mathbf{R}^{calib} = \mathbf{R}^{calib} \mathbf{R}^{lidar}, \quad (6a)$$

$$\mathbf{R}^{cam} \mathbf{t}^{calib} + s \mathbf{t}^{cam} = \mathbf{R}^{calib} \mathbf{t}^{lidar} + \mathbf{t}^{calib}. \quad (6b)$$

$\mathbf{R} \in \mathbb{R}^{3 \times 3}$  denotes a rotation matrix and  $\mathbf{t} \in \mathbb{R}^{3 \times 1}$  is a translation vector. Since ORB-SLAM3 does not produce metrically aware odometry estimates, we add a scaling factor  $s$ . Next, we apply Eq. (5) to all paired pose differences  $(\mathbf{T}_i^{cam}, \mathbf{T}_i^{lidar})$  with  $i = [1, k]$  and define  $J^{rot}$  and  $J^{trans}$  as follows:

$$J^{rot} = \sum_i \rho_i(e_{i,rot}^T \Omega_{rot} e_{i,rot}), \quad (7a)$$

$$J^{trans} = \sum_i \rho_i(e_{i,trans}^T \Omega_{trans} e_{i,trans}), \quad (7b)$$




Fig. 3. In the fine-tuning stage, we employ CMRNext [24] to find 2D pixel to 3D point correspondences. First, a LiDAR point is projected onto the image space using the coarse calibration parameters. Second, CMRNext predicts a 2D offset to correct the projection. Finally, during optimization, the calibration parameters are adjusted to match the corrected projection.

with error functions  $e_{rot}$  and  $e_{trans}$  induced by Eq. (6a) and Eq. (6b), respectively:

$$\mathbf{e}_{i,rot} = [(\mathbf{R}^{calib} \mathbf{R}_i^{lidar})^{-1} (\mathbf{R}_i^{cam} \mathbf{R}^{calib}) - \mathbf{I}]_{vec}, \quad (8a)$$

$$\mathbf{e}_{i,trans} = (\mathbf{R}_i^{cam} - \mathbf{I}) \mathbf{t}^{calib} + \mathbf{s}_i \mathbf{t}_i^{cam} - \mathbf{R}^{calib} \mathbf{t}_i^{lidar}, \quad (8b)$$

where  $\mathbf{I}$  denotes the identity matrix. We further define the operator  $[\cdot]_{vec}$  that reshapes a matrix  $\mathbf{M} \in \mathbb{R}^{d \times d}$  to  $\mathbb{R}^{d^2 \times 1}$  by stacking its columns. Note that we use the same information matrices  $\Omega_{rot}$  and  $\Omega_{trans}$  for all data pairs. In Fig. 4, we refer to the tuple of both error functions as *odometry constraints*. Finally, we perform MAP estimation using the combined sensor motion cost  $J^{mot}$  to obtain the initial calibration parameters  $\mathbf{T}^{calib}_{init}$  and scaling factors  $s_{1:k}^*$ :

$$\mathbf{T}^{calib}_{init} = \arg \min_{\mathbf{T}^{calib}} J^{rot} + J^{trans} = \arg \min_{\mathbf{T}^{calib}} J^{mot}. \quad (9)$$

### C. Fine Registration

The second step of MDPCalib extends the coarse registration from sensor motion with matching correspondences between image pixels and LiDAR points. In contrast to prior works [17], [22], we propose to perform joint optimization with respect to both constraints explicitly exploiting their complementary information. We demonstrate the superiority of this design in Sec. IV-C. As shown in Fig. 2, we construct input triplets comprising an image  $Img_j^{cam}$ , a synchronized point cloud  $Pcl_j^{lidar}$  (see Sec. III-D), and the calibration parameters  $\mathbf{T}^{calib}_{init}$  from coarse registration. This data is fed to the deep learning-based CMRNext [24] that is able to register a camera frame to a 3D point cloud. We select CMRNext due to its generalizability to new scenes and sensor models. We illustrate the detailed steps in Fig. 3: First, a 3D point is projected to a 2D pixel coordinate based on the provided initial guess. Second, CMRNext estimates an offset to correct the initial guess. This enhancement step is performed iteratively with network weights trained for decreasing offsets. Formally,

$$p_j^{cmr} = CMRNext(p_j^{lidar}, Img_j^{cam}, \mathbf{T}^{calib}_{init}), \quad (10)$$

\*We omit these in the optimization equations to improve readability.




Fig. 4. We interpolate the poses from LiDAR odometry to the timestamps of the camera poses and further project the point cloud of the nearest neighbor to the same time to yield synchronized image-point cloud pairs. The odometry poses are then used to align the sensor motion, whereas the sensor measurements are fed to CMRNext [24] to obtain point correspondences.

with point  $p_j^{lidar} \in Pcl_j^{lidar}$ . Third, we define a error function  $e_{corr}$  that attempts to gradually alter the calibration parameters  $\mathbf{T}^{calib}$  such that the direct 3D-to-2D projection approaches the estimate of CMRNext:

$$e_{j,corr} = proj(p_j^{lidar}, \mathbf{K}, \mathbf{T}^{calib}) - p_j^{cmr}, \quad (11)$$

where  $\mathbf{K}$  denotes the camera matrix. This error function, referred to as *point correspondence constraints* in Fig. 4, is then plugged in the corresponding cost function  $J^{corr}$ :

$$J^{corr} = \sum_j \rho_i(e_{j,corr}^T \Omega_{corr} e_{j,corr}), \quad (12)$$

with  $\Omega_{corr}$  denoting the information matrix for all  $j$ . As indicated in Eq. (4), we do not add an error term for every image-point cloud pair. For instance, we do not apply  $J^{corr}$  for all  $p_j^{lidar} \in Pcl_j^{lidar}$  but only to a subset reducing the number of partially redundant constraints in the optimization problem. Finally, we repeat the process of MAP estimation to obtain the overall calibration parameters  $\mathbf{T}^{calib}_{final}$ :

$$\mathbf{T}^{calib}_{final} = \arg \min_{\mathbf{T}^{calib}} J^{mot} + J^{corr}. \quad (13)$$

### D. Pose Synchronization

As noted in Sec. III-B, we interpolate camera and LiDAR poses to the same timestamp to yield synchronized sensor motion pairs. We further assume the usage of a global-shutter camera, i.e., all pixels are captured simultaneously. Since shifting a point cloud from one timestamp to another is significantly easier than simulating an image to be taken at a different time, we use the timestamps of the camera data as the reference. Next, we identify the two LiDAR measurements recorded before and after the image was taken, shown in Fig. 4. Based on the corresponding poses from LiDAR odometry, we compute an intermediate pose via linear and spherical linear interpolation for the translation and rotation components, respectively. Finally, we project the point cloud of the nearest neighbor to the same timestamp to obtain synchronized image-point cloud pairs that are used as input to CMRNext [24].

## IV. EXPERIMENTAL EVALUATION

In this section, we extensively evaluate our proposed MDPCalib approach for various robotic platforms and provide a comparison with previous methods. We further analyze the effects of several parameters of our method.






Fig. 5. We calibrate the sensors on three in-house robotic platforms including a self-driving perception car, a quadruped robot, and a UAV.

### A. Robotic Platforms

We apply our method to a diverse set of sensor setups on four robotic platforms including the public KITTI dataset [32] as well as three in-house configurations. In the supplementary material, we provide further results on the Argoverse 2 [33] dataset.

**KITTI Dataset:** We extensively evaluate our method on the publicly available KITTI dataset [32]. In detail, we utilize the rectified images and undistorted point clouds of sequence 00 of the odometry benchmark as data from the left camera of the other sequences was seen during the training of CMRNext [24]. The vehicle is equipped with two RGB cameras Point Grey Flea 2 and one LiDAR of type Velodyne HDL-64E. Sensor data is captured at 10 Hz.

**In-House Robots:** We further employ our method to three in-house datasets that were not seen during the training of CMRNext [24] and include new sensor models. Each robot, shown in Fig. 5, is equipped with one Ouster OS-1 LiDAR with 128 channels and different cameras. In detail, our vehicle is equipped with four FLIR Blackfly 2353C, the quadruped robot includes five AVT cameras, and the UAV uses a FLIR Blackfly S. The sensors on all robots capture data at 10 Hz. The LiDAR is time synchronized via PTP, whereas the images are stamped on arrival. To compute the calibration error, we obtain reference parameters using manually selected camera-LiDAR point correspondences [7].

### B. Main Results

We measure the calibration error as the difference between the calibration parameters produced by MDPCalib and the provided or manually obtained parameters. Throughout this section, we report both the magnitude of the rotation and translation errors as well as the errors of the individual axes. In detail, we compute the errors as follows:

$$E_t = \|\mathbf{t} - \hat{\mathbf{t}}\|_2, \quad (14a)$$

$$m = q * \hat{q}^{-1}, \quad (14b)$$

$$E_R = \text{atan2} \left( \sqrt{m_x^2 + m_y^2 + m_z^2}, m_w \right), \quad (14c)$$

where  $\mathbf{t}$  and  $\hat{\mathbf{t}}$  are ground truth and predicted translation parameters,  $q$  and  $\hat{q}$  denote the corresponding parameters of

the rotation  $\mathbf{R}$  as quaternions. Finally,  $*$  and  $^{-1}$  are the multiplicative and inverse operations for quaternions.

For KITTI, we compare our MDPCalib approach to several classical [11], [17] and learning-based [20], [21], [24], [25], [34], [35] baselines in Tab. I. If the authors released the corresponding code, we reproduce the calibration results for both the left and the right camera. Besides our MDPCalib, only the method by Yin *et al.* [17] is initialization-free. For the others, we follow the original sampling ranges. Note that the reproduced metrics of the learning-based approaches correspond to median calibration over multiple frames. For our method, we utilize 1,000 sensor motion constraints and 5% randomly sampled correspondences from each of the 100 image-point cloud pairs. Although all baselines are outperformed by our approach, they generally yield accurate results for calibrating the LiDAR to the left camera. However, it is paramount to emphasize that all learning-based methods incorporate samples of the left camera in their training data. Therefore, we also calibrate the right camera measuring the capability to generalize. Except for CMRNext, all baselines suffer from a substantial performance drop. Nonetheless, MDPCalib yields the smallest error demonstrating that our joint optimization further increases robustness to unseen sensor configuration.

We confirm this observation by employing MDPCalib, CMRNext [24], and the method by Koide *et al.* [18] to three in-house robotic platforms. To account for the missing hardware time synchronization between camera and LiDAR, we only utilize static frames for computing point correspondences. As shown in Tab. II, MDPCalib yields significantly smaller calibration errors on the vehicle and the quadruped robot. For the UAV, we hypothesize that point cloud undistortion is more challenging due to less constrained motion resulting in less accurate odometry and hence larger errors.

We report our calibration parameters for KITTI in the supplementary material. We show qualitative results by projecting the point clouds onto the images in Fig. 10 and provide further visualizations in the complementary video. Notably, the projections based on our calibration are still highly accurate at the maximum distance of 80 m.

### C. Ablation Studies

We extensively analyze the sensitivity of MDPCalib towards several design choices and adjustable parameters by conducting various ablation studies on the KITTI dataset [32]. In particular, we present results for calibrating the right RGB camera based on sequence 00. If not explicitly stated otherwise, we utilize 1,000 odometry poses and 5% of the point correspondences from 100 image-point cloud pairs. In Figs. 6 to 9, we visualize the mean and standard deviation of three runs. To improve readability, we add small offsets to the  $x$ -values. For convenience, we provide the raw numerical metrics in the supplementary material.

**Components Analysis:** We analyze the impact of the calibration stages and the components of the optimization problem in Tab. III. Although the initialization-free coarse registration step based on sensor motion reduces the rotation error to sub-degree accuracy, the translation parameters suffer from

TABLE I  
CALIBRATION ERROR ON THE KITTI DATASET

Method	Initial range	Left camera								Right camera							
		Magnitude		Translation [cm]			Rotation [°]			Magnitude		Translation [cm]			Rotation [°]		
		$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw
Tu <i>et al.</i> <sup>†</sup> [11]	±0.20 m / ±2°	4.40	0.16	–	–	–	–	–	–	–	–	–	–	–	–	–	–
Yin <i>et al.</i> <sup>†</sup> [17]	–	5.91	0.16	2.90	4.90	1.60	0.08	0.09	0.10	–	–	–	–	–	–	–	–
Borer <i>et al.</i> <sup>†</sup> [21]	±0.25 m / ±1°	9.51	0.18	9.40	1.30	0.60	0.18	<b>0.03</b>	<b>0.03</b>	–	–	–	–	–	–	–	–
CalibDepth <sup>†</sup> [20]	±1.5 m / ±20°	1.17	0.12	1.31	1.02	1.17	0.06	0.23	0.08	–	–	–	–	–	–	–	–
CMRNet [34]	±1.5 m / ±20°	1.57	0.10	1.06	0.74	<u>0.34</u>	<u>0.03</u>	<b>0.01</b>	0.08	52.92	1.49	1.59	52.87	<b>0.36</b>	0.04	<b>0.02</b>	1.49
RGGNet [35]	±0.3 m / ±20°	11.49	1.29	8.14	2.79	<u>3.97</u>	<u>0.35</u>	0.74	0.64	23.52	3.87	18.03	5.55	6.06	0.51	3.38	1.48
LCCNet [25]	±1.5 m / ±20°	<u>1.01</u>	0.12	<u>0.26</u>	<u>0.36</u>	0.35	<b>0.02</b>	0.11	<b>0.03</b>	52.51	1.47	52.48	<b>0.26</b>	0.74	<b>0.01</b>	1.47	<b>0.03</b>
CMRNext [24]	±1.5 m / ±20°	1.89	<u>0.08</u>	1.12	0.83	0.79	0.04	0.04	<u>0.04</u>	<u>7.07</u>	<u>0.23</u>	2.17	5.78	0.94	0.05	<u>0.05</u>	0.20
MDPCalib ( <i>ours</i> )	–	<b>0.18</b>	<b>0.06</b>	<b>0.07</b>	<b>0.16</b>	<b>0.01</b>	<b>0.02</b>	0.04	<u>0.04</u>	<b>2.94</b>	<b>0.14</b>	<b>0.66</b>	<u>2.78</u>	<u>0.49</u>	<u>0.03</u>	<u>0.05</u>	<u>0.13</u>

We provide results based on data from sequence 00 of the KITTI odometry benchmark [32]. Unlike many other works, we evaluate our approach and previous methods for both cameras. Bold and underlined values indicate the best and second-best scores, respectively. <sup>†</sup>: These methods did not release (English-speaking) code preventing reproducing results for both cameras.

TABLE II  
CALIBRATION ERROR ON IN-HOUSE ROBOTIC PLATFORMS

Platform	Method	Magnitude		Translation [cm]			Rotation [°]		
		$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw
Vehicle	Koide <i>et al.</i> [18]	9.76	<b>0.27</b>	9.24	<b>0.63</b>	3.06	0.07	<b>0.18</b>	0.19
	CMRNext [24]	12.20	<u>0.90</u>	<b>1.35</b>	10.83	3.60	0.45	0.82	0.38
	MDPCalib ( <i>ours</i> )	<b>4.50</b>	<b>0.27</b>	<u>4.04</u>	<u>1.71</u>	<b>1.03</b>	<b>0.00</b>	<b>0.19</b>	<b>0.18</b>
Quadruped	Koide <i>et al.</i> [18]	16.21	1.34	<u>2.10</u>	<u>15.61</u>	<b>3.84</b>	1.07	0.67	0.44
	CMRNext [24]	23.95	1.36	3.73	16.07	15.05	<u>0.40</u>	0.85	<u>0.32</u>
	MDPCalib ( <i>ours</i> )	<b>9.54</b>	<b>0.38</b>	<b>1.27</b>	<b>2.70</b>	<u>9.06</u>	<b>0.31</b>	<b>0.19</b>	<b>0.10</b>
UAV	Koide <i>et al.</i> [18]	<b>1.65</b>	<b>0.36</b>	<u>1.56</u>	<b>0.08</b>	<b>0.50</b>	<u>0.16</u>	<u>0.17</u>	<b>0.28</b>
	CMRNext [24]	12.47	0.97	<b>1.17</b>	5.31	6.83	0.27	0.34	0.36
	MDPCalib ( <i>ours</i> )	<u>5.12</u>	<u>0.51</u>	4.19	<u>0.59</u>	<u>2.88</u>	<b>0.15</b>	<b>0.13</b>	<u>0.30</u>

We obtain reference calibration parameters to compute the errors by using manually selected camera-LiDAR point correspondences [7]. Neither our MDPCalib nor the method by Koide *et al.* [18] requires an initial guess. Bold and underlined values indicate the best and second-best scores.

a lack of observability. For reference, we provide the error of the median calibration parameters when running PnP on 100% of the predicted point correspondences. In Tab. III, we refer to this as CMRNext [24]. In MDPCalib, we instead process a subset of the correspondences as constraints in the optimization formulation, which further reduces the errors. In the bottom two rows, we demonstrate the efficacy of the key ingredients of MDPCalib. First, we show that utilizing the Cauchy loss [26] to robustify the cost functions significantly improves the results by reducing the effect of outliers in the observations. Second, while previous works [22], [17] only proposed to utilize sensor motion to initialize a subsequent correspondence-based calibration scheme, we incorporate both constraints in the refinement stage leveraging their complementary information. Our experiment clearly underlines the positive impact of joint optimization with respect to both sensor motion and point correspondences.

*Runtime Analysis:* We discuss the runtime of our method with respect to the number of both motion-based and point correspondences-based constraints. We conduct these experiments on a machine with an AMD Ryzen Threadripper PRO 3975WX CPU with 128 GB and an NVIDIA A6000 GPU with 48 GB. In Fig. 6 and Fig. 7, we visualize the runtime versus the number of poses and the relative number of correspondences, respectively. As can be seen in both studies, the runtime scales approximately linearly, whereas the errors decrease. We further observe an optimal configuration, after which the runtime continues to increase without major impacts on the accuracy.

TABLE III  
COMPONENTS ANALYSIS

Component	Magnitude		Translation [cm]			Rotation [°]		
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw
Coarse registration	39.37	0.51	18.11	3.57	34.77	0.38	0.33	<b>0.09</b>
+ CMRNext [24]	6.26	0.28	<b>0.31</b>	6.25	<b>0.00</b>	<b>0.01</b>	<b>0.00</b>	0.28
Fine registration:								
Point constraints	5.89	0.21	0.78	5.81	0.50	0.04	0.06	0.19
+ Cauchy robustifier	<u>3.42</u>	0.16	0.35	<u>3.38</u>	<u>0.37</u>	<u>0.02</u>	<u>0.05</u>	0.15
+ Motion constraints	<b>2.94</b>	<b>0.14</b>	0.66	<b>2.78</b>	0.49	0.03	<u>0.05</u>	<u>0.13</u>

The last line highlighted in gray corresponds to our proposed MDPCalib. CMRNext denotes the errors of the median calibration parameters when running PnP on 100% of the predicted point correspondences after initialization with the coarse registration. Bold and underlined values indicate the best and second-best scores, respectively.

*Number of Poses:* In this study, we evaluate the effect of the number of poses from visual and LiDAR odometry on the calibration error, i.e., the length of the required data recording. In Fig. 6, we show the translation and rotation errors for an increasing number of poses. Although it generally holds that adding more motion constraints to the optimization problem results in more accurate calibration parameters, we observe that the error converges.

*Number of Point Correspondences:* Next, we repeat a similar study for the number of constraints induced by the point correspondences. For KITTI data, we measure an average of 20,000 correspondences per image-point cloud pair, i.e., our default setting of utilizing 5% generates approximately 1,000 constraints. In Fig. 7, we keep a fixed number of 100 pairs and vary the relative number of constraints. Fewer correspondences result in higher calibration errors. We confirm this observation in a second experiment visualized in Fig. 8, where we utilize 5% of the point correspondences but reduce the number of image-point cloud pairs. Note that the errors converge towards a lower limit in both experimental setups. Finally, we address the following question: Given a fixed number of correspondences, and hence runtime, is it preferable to increase the number of image-point cloud pairs or the relative amount of correspondences per pair? In Fig. 9, we plot the calibration errors for a stable number of point correspondences, i.e., when decreasing the number of pairs with respect to the default setting, we increase the relative number of correspondences by the respective amount. The experiment shows that a smaller number of correspondences from a more diverse set of data pairs is preferable.




Fig. 6. For the point correspondence constraints, we use 5% correspondences per image-point cloud pair from a total of 100 pairs.




Fig. 7. For the motion constraints, we use 1,000 matched poses. Per image-point cloud pair, we select 5% of the correspondences.

## V. CONCLUSION

In this paper, we proposed MDPCalib, a novel method for the extrinsic calibration between camera and LiDAR sensors in an automated manner without the need for dedicated calibration targets. Our approach utilizes non-linear optimization to obtain the calibration parameters by aligning sensor motions from visual and LiDAR odometry and leveraging deep learning-based correspondences between 2D pixels and 3D points. In contrast to most previous learning-based methods, our approach generalizes to sensor configurations that differ from the training setup. Importantly, MDPCalib does not require accurate parameters for initialization. In practical experiments carried out on diverse robotic platforms, we demonstrated the efficacy of our method and provided detailed evaluations of several design choices. Future versions of our approach could incorporate constraints for multiple cameras and LiDARs and include intrinsic calibration in the optimization problem. Due to its unsupervised procedure, MDPCalib could further be extended to online calibration.

## ACKNOWLEDGMENT

The authors thank Johan Vertens for designing the in-house driving platform used in this work and both Julia Hindel and Iana Zhura for assistance with the quadrupedal robot. Finally, the authors thank Wera Winterhalter, Freya Fleckenstein, and Srdjan Mилоjevic for technical discussions.

## REFERENCES

- [1] L. Andresen, A. Brandemuehl, A. Honger, B. Kuan, N. Vödisch, H. Blum, V. Reijgwart, L. Bernreiter, L. Schaupp, J. J. Chung, M. Burki, M. R. Oswald, R. Siegwart, and A. Gavel, “Accurate mapping and planning for autonomous racing,” in *Int. Conf. Intelligent Robots and Systems*, 2020, pp. 4743–4749.
- [2] T. Liang, H. Xie, K. Yu, Z. Xia, Z. Lin, Y. Wang, T. Tang, B. Wang, and Z. Tang, “BEVFusion: A simple and robust LiDAR-camera fusion framework,” in *Adv. in Neu. Inf. Proc. Sys.*, 2022, pp. 10 421–10 434.
- [3] C. Ge, J. Chen, E. Xie, Z. Wang, L. Hong, H. Lu, Z. Li, and P. Luo, “MetaBEV: Solving sensor failures for 3D detection and map segmentation,” in *Int. Conf. Comp. Vis.*, October 2023, pp. 8721–8731.




Fig. 8. For the motion constraints, we use 1,000 matched poses. We extract correspondences from a total of 100 image-point cloud pairs.




Fig. 9. For the motion constraints, we use 1,000 matched poses. We vary the pose count and the relative number of correspondences to obtain an approximately constant absolute number of correspondences.

- [4] J. Schramm, N. Vödisch, K. Petek, R. B. Kiran, S. Yogamani, W. Burgard, and A. Valada, “BEVCar: Camera-radar fusion for BEV map and object segmentation,” *arXiv preprint arXiv:2403.11761*, 2024.
- [5] A. Dhall, K. Chelani, V. Radhakrishnan, and K. M. Krishna, “LiDAR-camera calibration using 3D-3D point correspondences,” *arXiv preprint arXiv:1705.09785*, 2017.
- [6] Q. Zhang and R. Pless, “Extrinsic calibration of a camera and laser range finder (improves camera calibration),” in *Int. Conf. Intelligent Robots and Systems*, vol. 3, 2004, pp. 2301–2306.
- [7] H. Vhavle, “ROS camera lidar calibration package,” 2021, [https://github.com/heethesh/lidar\\_camera\\_calibration](https://github.com/heethesh/lidar_camera_calibration).
- [8] C. Guindel, J. Beltrán, D. Martín, and F. García, “Automatic extrinsic calibration for lidar-stereo vehicle sensor setups,” in *Int. Conf. Intel. Transp. Systems*, 2017.
- [9] E. Kim and S.-Y. Park, “Extrinsic calibration between camera and LiDAR sensors by matching multiple 3D planes,” *Sensors*, vol. 20, no. 1, p. 52, 2019.
- [10] D. Zhang, L. Ma, Z. Gong, W. Tan, J. Zelek, and J. Li, “An overlap-free calibration method for LiDAR-camera platforms based on environmental perception,” *Trans. Instr. and Meas.*, vol. 72, 2023.
- [11] D. Tu, B. Wang, H. Cui, Y. Liu, and S. Shen, “Multi-camera-LiDAR auto-calibration by joint structure-from-motion,” in *Int. Conf. Intelligent Robots and Systems*, 2022, pp. 2242–2249.
- [12] Z. Zhang, “Flexible camera calibration by viewing a plane from unknown orientations,” in *Int. Conf. Comp. Vis.*, vol. 1, 1999, pp. 666–673.
- [13] M. Vel’as, M. Španěl, Z. Materna, and A. Herout, “Calibration of RGB camera with Velodyne LiDAR,” in *Int. Conf. in Central Europe on Comp. Graphics, Visualization and Comp. Vis.*, 2014, pp. 135–144.
- [14] J. Ou, P. Huang, J. Zhou, Y. Zhao, and L. Lin, “Automatic extrinsic calibration of 3D LIDAR and multi-cameras based on graph optimization,” *Sensors*, vol. 22, no. 6, p. 2221, 2022.
- [15] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, “g<sup>2</sup>o: A general framework for graph optimization,” in *Int. Conf. Rob. and Autom.*, 2011, pp. 3607–3613.
- [16] C. Yuan, X. Liu, X. Hong, and F. Zhang, “Pixel-level extrinsic self calibration of high resolution lidar and camera in targetless environments,” *Rob. and Autom. Letters*, vol. 6, no. 4, pp. 7517–7524, 2021.
- [17] J. Yin, F. Yan, Y. Liu, and Y. Zhuang, “Automatic and targetless LiDAR-camera extrinsic calibration using edge alignment,” *Sensors*, vol. 23, no. 17, pp. 19 871–19 880, 2023.
- [18] K. Koide, S. Oishi, M. Yokozuka, and A. Banno, “General, single-shot, target-less, and automatic LiDAR-camera extrinsic calibration toolbox,” in *Int. Conf. Rob. and Autom.*, 2023, pp. 11 301–11 307.
- [19] T. Caselitz, B. Steder, M. Ruhnke, and W. Burgard, “Monocular camera localization in 3D LiDAR maps,” in *Int. Conf. Intelligent Robots and Systems*, 2016, pp. 1926–1931.
- [20] J. Zhu, J. Xue, and P. Zhang, “CalibDepth: Unifying depth map representation for iterative LiDAR-camera online calibration,” in *Int. Conf. Rob. and Autom.*, 2023, pp. 726–733.






Fig. 10. Qualitative results of MDPCalib on the KITTI dataset and our in-house robotic platforms. Best viewed on screen and zoomed in.

- [21] J. Borer, J. Tschirner, F. Ölsner, and S. Milz, “From chaos to calibration: A geometric mutual information approach to target-free camera LiDAR extrinsic calibration,” in *Winter Conf. Applications of Comp. Vis.*, 2024, pp. 8394–8403.
- [22] Z. Taylor and J. Nieto, “Motion-based calibration of multimodal sensor arrays,” in *Int. Conf. Rob. and Autom.*, 2015, pp. 4843–4850.
- [23] N. Schneider, F. Piewak, C. Stiller, and U. Franke, “RegNet: Multimodal sensor registration using deep neural networks,” in *Intelligent Vehicles Symposium*, 2017, pp. 1803–1810.
- [24] D. Cattaneo and A. Valada, “CMRNext: Camera to LiDAR matching in the wild for localization and extrinsic calibration,” *arXiv preprint arXiv:2402.00129*, 2024.
- [25] X. Lv, B. Wang, Z. Dou, D. Ye, and S. Wang, “LCCNet: LiDAR and camera self-calibration using cost volume network,” in *Conf. Comp. Vis. and Pattern Recogni.*, 2021, pp. 2894–2901.
- [26] T. Liu and D. Tao, “On the robustness and generalization of cauchy regression,” in *Int. Conf. Inf. Sci. and Techn.*, 2014, pp. 100–105.
- [27] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multimap SLAM,” *Trans. on Rob.*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [28] N. Vödisch, D. Cattaneo, W. Burgard, and A. Valada, “CoVIO: Online continual learning for visual-inertial odometry,” in *Conf. Comp. Vis. and Pattern Recogni. Workshops*, 2023.
- [29] H. Zhan, C. S. Weerasekera, J. W. Bian, and I. Reid, “Visual odometry revisited: What should be learnt?” in *Int. Conf. Rob. and Autom.*, 2020, pp. 4203–4210.
- [30] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, “FAST-LIO2: Fast direct lidar-inertial odometry,” *Trans. on Rob.*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [31] Y. Shiu and S. Ahmad, “Calibration of wrist-mounted robotic sensors by solving homogeneous transform equations of the form  $AX=XB$ ,” *Trans. on Rob. and Autom.*, vol. 5, no. 1, pp. 16–29, 1989.
- [32] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The KITTI dataset,” *Int. Journal of Rob. Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [33] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, and J. Hays, “Argoverse 2: Next generation datasets for self-driving perception and forecasting,” in *Neu. Inf. Proc. Sys. Track on Datasets and Benchmarks*, 2021.
- [34] D. Cattaneo, M. Vaghi, A. L. Ballardini, S. Fontana, D. G. Sorrenti, and W. Burgard, “CMRNet: Camera to LiDAR-map registration,” in *Int. Conf. Intel. Transp. Systems*, 2019, pp. 1283–1289.
- [35] K. Yuan, Z. Guo, and Z. J. Wang, “RGGNet: Tolerance aware LiDAR-camera online calibration with geometric deep learning and generative model,” *Rob. and Autom. Letters*, vol. 5, no. 4, pp. 6956–6963, 2020.



# Automatic Target-Less Camera-LiDAR Calibration From Motion and Deep Point Correspondences

## - Supplementary Material -

Kürsat Petek<sup>1\*</sup>, Niclas Vödisch<sup>1\*</sup>, Johannes Meyer<sup>1</sup>, Daniele Cattaneo<sup>1</sup>, Abhinav Valada<sup>1</sup>, Wolfram Burgard<sup>2</sup>

In this supplementary material, we report results on the Argoverse 2 [33] dataset, release calibration parameters for the KITTI dataset [32], and provide more detailed numbers for our ablation studies.

### S-I. RESULTS ON ARGOVERSE 2

In this section, we provide results on the Argoverse 2 [33] dataset. In contrast to the KITTI odometry benchmark, Argoverse 2 contains many short sequences lasting only 15-30s, which is too short for our method to be applied. Furthermore, since the ground truth extrinsic calibration is not consistent between all sequences, we collect sequences that share the same set of ground truth parameters and combine them in a larger sequence with small gaps in between that can be detected by our method via the lost-track feature of ORB-SLAM3 [27]. In particular, in this experiment, we utilize sequence 05fb81ab-5e46-3f63-a59f-82fc66d5a477 as the reference sequence. Note that CMRNext [24] has not been trained on the Argoverse 2 [33] dataset, therefore, allowing us to use sequences from the `train` split for evaluation. In Tab. S-I, we report the calibration error for both stereo cameras of CMRNext [24] and our method MDPCalib.

### S-II. KITTI CALIBRATION PARAMETERS

To foster future research towards camera-LiDAR sensor fusion, we release the calibration parameters that we obtained with our method for sequence 00 of the KITTI dataset [32]. In detail, we provide the rotation matrices and translation vectors for projecting the LiDAR point cloud into the image.

*Left RGB Camera:*

$$\mathbf{R} = \begin{bmatrix} -1.2619\text{e-}4 & -9.9997\text{e-}1 & -8.2230\text{e-}3 \\ -7.8537\text{e-}3 & 8.2238\text{e-}3 & -9.9994\text{e-}1 \\ 9.9997\text{e-}1 & -6.1602\text{e-}5 & -7.8545\text{e-}3 \end{bmatrix}$$

$$\mathbf{t} = [5.1090\text{e-}2 \quad -5.5873\text{e-}2 \quad -2.9575\text{e-}1]^T$$

*Right RGB Camera:*

$$\mathbf{R} = \begin{bmatrix} -1.8202\text{e-}3 & -9.9996\text{e-}1 & -8.2416\text{e-}3 \\ -8.2823\text{e-}3 & 8.2564\text{e-}3 & -9.9993\text{e-}1 \\ 9.9996\text{e-}1 & -1.7518\text{e-}3 & -8.2970\text{e-}3 \end{bmatrix}$$

$$\mathbf{t} = [-4.5166\text{e-}1 \quad -4.8448\text{e-}2 \quad -2.8787\text{e-}1]^T$$

TABLE S-I  
CALIBRATION ERROR ON THE ARGOVERSE 2 DATASET

Method	Magnitude		Stereo left camera			Rotation [°]		
			Translation [cm]					
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw
CMRNext [24]	17.31	0.35	10.52	10.82	8.47	<b>0.01</b>	0.30	0.18
MDPCalib ( <i>ours</i> )	<b>9.48</b>	<b>0.18</b>	<b>7.86</b>	<b>5.03</b>	<b>1.70</b>	0.04	<b>0.15</b>	<b>0.09</b>

Method	Magnitude		Stereo right camera			Rotation [°]		
			Translation [cm]					
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw
CMRNext [24]	31.27	0.36	29.15	9.92	5.44	0.14	0.29	0.15
MDPCalib ( <i>ours</i> )	<b>22.90</b>	<b>0.21</b>	<b>20.88</b>	<b>8.86</b>	<b>3.17</b>	<b>0.01</b>	<b>0.17</b>	<b>0.12</b>

We use sequence 05fb81ab-5e46-3f63-a59f-82fc66d5a477 and sequences with the same set of ground truth calibration parameters. Bold values denote the best score per metric.

### S-III. ABLATION STUDIES

In this section, we provide the numerical values that are used to generate the figures in the main manuscript. Throughout the tables, we highlight the parameters that correspond to our overall setting in gray. For the magnitude errors  $E_t$  and  $E_R$ , we report the mean and standard deviation of three runs. For the errors of the individual axes, we report only the mean.

\* Equal contribution.

<sup>1</sup> Department of Computer Science, University of Freiburg, Germany.

<sup>2</sup> Department of Eng., University of Technology Nuremberg, Germany.

TABLE S-II  
NUMBER OF POSES

Pose count	Magnitude		Translation [cm]			Rotation [°]			Time [s]
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw	
1000	2.94±0.13	0.14±0.00	0.66	2.78	0.49	0.03	0.05	0.13	3402
800	3.14±0.07	0.14±0.01	0.53	3.04	0.42	0.02	0.05	0.13	3011
600	3.42±0.10	0.15±0.00	0.08	3.39	0.36	0.02	0.05	0.14	1303
400	4.08±0.39	0.17±0.01	0.17	4.06	0.28	0.01	0.04	0.16	729
200	3.52±0.44	0.15±0.02	0.68	3.44	0.23	0.02	0.04	0.14	355

This table corresponds to Fig. 6. For the point correspondence constraints, we use 5% correspondences per image-point cloud pair from 100 pairs.

TABLE S-III  
RELATIVE NUMBER OF POINT CORRESPONDENCES

Corr. count	Magnitude		Translation [cm]			Rotation [°]			Time [s]
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw	
5%	2.94±0.13	0.14±0.00	0.66	2.78	0.49	0.03	0.05	0.13	3402
4%	3.17±0.03	0.14±0.00	0.39	3.12	0.34	0.01	0.04	0.13	2939
3%	3.44±0.31	0.15±0.02	0.50	3.36	0.43	0.02	0.05	0.14	2558
2%	3.44±1.06	0.15±0.04	0.19	3.42	0.34	0.01	0.04	0.08	1153
1%	3.64±0.51	0.15±0.02	0.18	3.59	0.31	0.02	0.05	0.14	808

This table corresponds to Fig. 7. For the motion constraints, we use 1,000 poses. We extract correspondences from 100 image-point cloud pairs.

TABLE S-IV  
NUMBER OF IMAGE-POINT CLOUD PAIRS

Pair count	Magnitude		Translation [cm]			Rotation [°]			Time [s]
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw	
100	2.94±0.13	0.14±0.00	0.66	2.78	0.49	0.03	0.05	0.13	3402
80	2.87±0.33	0.13±0.01	0.43	2.78	0.41	0.02	0.05	0.12	2760
60	2.99±0.52	0.13±0.02	0.10	2.96	0.32	0.02	0.05	0.12	2386
40	4.43±0.91	0.17±0.04	0.60	4.35	0.49	0.01	0.05	0.16	1745
20	3.80±1.48	0.16±0.04	0.69	3.60	0.12	0.01	0.05	0.15	902

This table corresponds to Fig. 8. For the motion constraints, we use 1,000 poses. Per image-point cloud pair, we select 5% of the correspondences.

TABLE S-V  
PAIR DIVERSITY VS. RELATIVE CORRESPONDENCES

Pair / Corr. count	Magnitude		Translation [cm]			Rotation [°]			Time [s]
	$E_t$ [cm]	$E_R$ [°]	x	y	z	roll	pitch	yaw	
100 / 5.0%	2.94±0.13	0.14±0.00	0.66	2.78	0.49	0.03	0.05	0.13	3402
80 / 6.25%	3.32±0.65	0.14±0.01	0.58	3.24	0.31	0.01	0.05	0.13	3054
60 / 8.33%	3.46±0.88	0.14±0.02	0.21	3.41	0.34	0.02	0.04	0.13	3696
40 / 12.5%	3.38±0.41	0.14±0.02	0.38	3.32	0.41	0.02	0.05	0.13	3904
20 / 25.0%	3.96±0.49	0.18±0.02	0.89	3.80	0.41	0.02	0.06	0.16	4451

This table corresponds to Fig. 9. For the motion constraints, we use 1,000 poses. We vary the pose count and the relative number of correspondences to obtain an approximately constant absolute number of correspondences.