Online Marker-free Extrinsic Camera Calibration using Person Keypoint Detections

Bastian Pätzold[®], Simon Bultmann[®], and Sven Behnke[®]

University of Bonn, Institute for Computer Science VI, Autonomous Intelligent Systems, Friedrich-Hirzebruch-Allee 8, 53115 Bonn, Germany {paetzold,bultmann,behnke}@ais.uni-bonn.de https://www.ais.uni-bonn.de

Abstract. Calibration of multi-camera systems, i.e. determining the relative poses between the cameras, is a prerequisite for many tasks in computer vision and robotics. Camera calibration is typically achieved using offline methods that use checkerboard calibration targets. These methods, however, often are cumbersome and lengthy, considering that a new calibration is required each time any camera pose changes. In this work, we propose a novel, marker-free online method for the extrinsic calibration of multiple smart edge sensors, relying solely on 2D human keypoint detections that are computed locally on the sensor boards from RGB camera images. Our method assumes the intrinsic camera parameters to be known and requires priming with a rough initial estimate of the camera poses. The person keypoint detections from multiple views are received at a central backend where they are synchronized, filtered, and assigned to person hypotheses. We use these person hypotheses to repeatedly solve optimization problems in the form of factor graphs. Given suitable observations of one or multiple persons traversing the scene, the estimated camera poses converge towards a coherent extrinsic calibration within a few minutes. We evaluate our approach in real-world settings and show that the calibration with our method achieves lower reprojection errors compared to a reference calibration generated by an offline method using a traditional calibration target.

Keywords: camera calibration \cdot human pose estimation \cdot factor graphs

1 Introduction

Sensor calibration is an essential prerequisite for most intelligent systems, as they combine data from numerous sensors to perceive a scene. To successfully interpret and fuse the measurements of multiple sensors, they need to be transformed into a common coordinate frame. This requires the sensor poses in a common reference frame—their extrinsic calibration. An imprecise calibration can lead to degradation in performance and can possibly cause critical safety issues.

This work was funded by grant BE 2556/16-2 (Research Unit FOR 2535 Anticipating Human Behavior) of the German Research Foundation (DFG).



Fig. 1. Extrinsic camera calibration using person keypoint detections (a) computed locally on different smart edge sensors: (b) based on a Google EdgeTPU Dev Board [12] and (c) based on an Nvidia Jetson Xavier NX Developer Kit [26]. 2D keypoints from multiple views are synchronized, filtered, and assigned to 3D person hypotheses (d) on a central backend. Observations are accumulated over time and a factor graph optimization is solved to obtain the optimal camera poses (coordinate systems in (d), with the blues axis being the view direction).

For multiple reasons, the task of camera calibration is an inherently difficult one to solve [25]: First, the calibration parameters change over time, by normal usage, e.g. due to vibration, thermal expansion, or moving parts. Therefore, it is not sufficient to calibrate the parameters only once during the construction of the system. Instead, calibration must be performed repeatedly throughout its lifetime. Second, the calibration parameters cannot be measured directly with sufficient precision; they must be inferred from the data captured by the considered cameras. Typically, the calibration is performed by actively deploying a calibration target of known correspondences in front of the cameras, e.g. a checkerboard pattern [34]. However, this requires expertise and might be perceived as cumbersome and lengthy when it has to be applied repeatedly for a large multi-camera system. Further challenges for inferring calibration parameters from image data involve accounting for noisy measurements and collecting a sufficient amount of data points spread over the entirety of the image planes.

In this work, we develop a novel, marker-free method for calibrating the extrinsic parameters of a network of static smart edge sensors, where each sensor runs inference for 2D human pose estimation. In particular, we infer the relative poses between the cameras of the deployed smart edge sensor boards in real time using the person keypoint detections transmitted by the sensors [3], as illustrated in Fig. 1. 2D keypoints from multiple views are synchronized, filtered, and assigned to 3D person hypotheses on a central backend, where observations

are accumulated over time and a factor graph optimization [8] is solved to obtain the optimal camera poses. The method can handle multiple persons in the scene, as well as arbitrary occlusions, e.g. from tables or pillars. We assume the intrinsic parameters of the cameras to be known and a rough initial estimate of the extrinsic calibration to be available, which can easily be obtained, e.g. from a floor plan or by tape measure.

Our proposed method alleviates many of the issues mentioned above: No specific calibration target is required; it suffices for one or several persons to walk through the scene. The method handles data association between multiple observed persons and their unknown dimensions. We propose an efficient online algorithm that optimizes the camera poses on-the-fly, giving direct feedback on success and when enough data has been captured. The calibration procedure thus can be repeated easily to account for parameter change over time, without expert knowledge. Furthermore, person keypoints can be detected from a significantly larger range of viewing angles (e.g. front, back, or side-view) than the pattern of a classical checkerboard calibration target, which is well detected only from a frontal view. This facilitates the collection of a sufficient amount of corresponding data points visible in multiple cameras that well constrain the factor graph optimization, further reducing the time required for calibration.

We evaluate the proposed approach in real-world settings and show that the calibration obtained with our method achieves better results in terms of reprojection errors in the targeted application domain of 3D multi-person pose estimation, compared to a reference calibration generated by an offline method using a traditional calibration target. We make our implementation publicly available¹.

2 Related Work

Camera Calibration. Traditional methods for camera calibration are based on using artificial image features, so-called *fiducials*. Their common idea is to deploy a calibration target with known correspondences in the overlapping field of view (FoV) of the considered cameras. Zhang [34] utilizes a checkerboard pattern on a planar surface to perform intrinsic calibration of single cameras. The *kalibr* toolkit [28] uses a planar grid of *AprilTags* [27] to perform offline extrinsic and intrinsic calibration of multiple cameras, which allows to fully resolve the target's orientation towards the cameras and is robust against occlusions. We apply this method to obtain a reference calibration for evaluating our work.

Reinke et al. [29] propose an offline method for finding the relative poses between a set of (two) cameras and the base frame of a quadruped robot. They use a fiducial marker mounted on the end-effector of a limb as the calibration target. The camera poses are resolved using a *factor graph* [8], modeling kinematic constraints between the marker frame and the base frame together with the visual constraints. We take up the idea of using factor graphs to model the calibration constraints in our work.

¹ https://github.com/AIS-Bonn/ExtrCamCalib_PersonKeypoints

To cope with the issues of traditional approaches, methods for camera calibration have been proposed that do not extract fiducial features from calibration targets but use naturally occurring features instead. Komorowski et al. [18] extract SIFT features [23] and find correspondences between multiple views using RANSAC [11]. They use segmentation to remove dynamic objects and validate their approach on stereo vision datasets. Their method is targeted towards one or a few small-baseline stereo cameras and offline processing of a small batch of images. Bhardwaj et al. [1] calibrate traffic cameras by extracting vehicle instances via deep neural networks (DNNs) and matching them to a database of popular car models. The extracted features and known dimensions of the car models are then used to formulate a PnP problem [11]. They assume a planar ground surface in the vicinity of the cars and process results offline.

A variety of methods considering surveillance scenarios use pedestrians as calibration targets. Lv et al. [24] track head & feet detections of a single pedestrian walking on a planar surface during the leg-crossing phases to perform offline extrinsic calibration of a single camera based on the geometric properties of vanishing points. Following a tracking approach, they resolve the corresponding intrinsic parameters based on Zhang [34]. Hödlmoser et al. [16] use a similar approach as [24], but expand the method to calibrate a camera network from pairwise relative calibrations. The absolute scale of the camera network is resolved by manually specifying the height of the walking person. Liu et al. [21] require a moderately crowded scene to perform online intrinsic and extrinsic calibration of a single camera by assuming strong prior knowledge regarding the height distribution of the observed pedestrians. The approach is based on computing vanishing points using RANSAC. In [22] they expand their method by introducing a joint calibration for a network of cameras based on the Direct Linear Transform [14]. Henning et al. [15] jointly optimize the trajectory of a monocular camera and a human body mesh by formulating an optimization problem in the form of a factor graph. They apply a human motion model to constrain sequential body postures and to resolve scale.

Guan et al. [13,31] detect head & feet keypoints for each observable pedestrian and perform pairwise triangulation assuming an average height for all visible persons in the image pair. They then compute the calibration offline, using RANSAC, followed by a gradient descent-based refinement scheme. Their resulting calibration is only defined up to an unknown scale factor, which must be resolved manually. The method assumes the center lines between all pedestrians to be parallel, in other words, all persons are assumed to stand upright during the calibration, whereas other poses, e.g. sitting persons, are not supported. Our method, in contrast, extracts up to 17 keypoints per person [20] using convolutional neural networks (CNNs) and assumes neither the dimensions or height of the persons, nor their pose or orientation towards the cameras to be known. As for the unknown dimensions of the persons, the scale of our calibration is also ambiguous up to a single scale factor. To address this issue, we force the scale of the initial estimate of the extrinsic calibration to be maintained throughout the calibration procedure. Human Pose Estimation. Human pose estimation refers to the task of detecting anatomical keypoints of persons on images. Early works use manually designed feature extractors like HOG-descriptors [6] or pictorial structures [10,9]. In recent years, approaches using CNNs have become popular and yield impressive results. Two well-known state-of-the-art and publicly available methods for human pose estimation are *OpenPose* [5] and *AlphaPose* [19]. Both methods estimate the poses of multiple persons simultaneously and in real time.

2D keypoint detections from multiple, calibrated camera views can be fused to obtain 3D human poses. We consider a network of smart edge sensors, introduced by Bultmann et al. [3,4], performing 3D human pose estimation in real time. Each smart edge sensor performs 2D human pose estimation, processing the image data locally on the sensor boards and transmitting only the obtained keypoint data. A central backend fuses the data received by the sensors to perform 3D human pose estimation via direct triangulation, and a semantic feedback loop is implemented to improve the 2D pose estimation on the sensor boards by incorporating global context information.

We adopt the smart edge sensor network [3] for our work, using the 2D person keypoints detected by the sensors as calibration features and aim to improve and facilitate the camera calibration required for this application scenario.

3 Method

Our method uses the image streams of a multi-camera system S_N with N > 1projective cameras C_i , $i \in [0 ... N - 1]$, to extract and maximize knowledge about the relative poses between all cameras in real-time, i.e. finding the translation $\mathbf{t}_{ij} \in \mathbb{R}^3$ and rotation $\mathbf{R}_{ij} \in SO(3) \subset \mathbb{R}^{3\times 3}$ between all camera pairs ij, where the pose of the optical center of C_i is defined by

$$\mathbf{C}_{i} = \begin{pmatrix} \mathbf{R}_{i} \ \mathbf{t}_{i} \\ \mathbf{0} \ 1 \end{pmatrix} \in \mathcal{SE}(3) .$$
(1)

We call this the *extrinsic calibration* of the multi-camera system S_N . Without loss of generality, we chose the first camera C_0 to be the origin of the global reference frame and set $\mathbf{C}_0 = \mathbf{I}_{4\times 4}$. In its local coordinate system, the view direction of each camera is the z-axis.

Fig. 2 gives an overview of our proposed pipeline. Each camera stream is fed into a person keypoint detector on the connected inference accelerator [3]. We refer to the unity of camera and detector as a *smart edge sensor*. The keypoint detections are transmitted to a central backend where they are timesynchronized and processed further. The clocks of sensors and backend are software-synchronized via NTP and each keypoint detection message includes a timestamp representing the capture time of the corresponding image.

The *preprocessing* stage removes redundant and noisy detections after which *data association* is performed, where correspondences between person detections from multiple views are established. Corresponding person detections are fused to form a person hypothesis and attached to a queue, which serves to decouple



Fig. 2. Proposed pipeline for extrinsic camera calibration using smart edge sensors and person keypoint detections. Images are analyzed locally on the sensor boards. Keypoint detections are transmitted to the backend where multiple views are fused to construct and solve optimization problems using factor graphs. A queue decouples the preprocessing and optimization stages.

the preprocessing stage from the rest of the pipeline. The optimization stage continuously reads from this queue, selects several person hypotheses, and uses them to construct and solve an optimization problem in the form of a factor graph [8]. The *refinement* stage updates the current estimate of the extrinsic calibration by smoothing the intermediate results generated by the optimization and compensates for scaling drift w.r.t. the initialization. As prerequisites for our method, we assume the intrinsic parameters of the cameras to be known and a rough initial estimate of the extrinsic calibration to be available, e.g. by tape measure or from a floor plan. The FoVs of all cameras must overlap in such a way that S_N forms a connected graph.

3.1 Preprocessing

The backend receives N person keypoint detection streams and synchronizes and preprocesses them such that they can be used for optimization. Each keypoint detection \mathcal{D}_{i}^{p} is associated to a person instance p and defined as

$$\mathcal{D}_{i}^{p} = \left\{ \left(u, v \right)^{\mathsf{T}}, c, \boldsymbol{\Sigma} \right\}, \tag{2}$$

where $(u, v)^{\mathsf{T}}$ are the image coordinates, $c \in [0, 1]$ is the confidence, $\Sigma \in \mathbb{R}^{2 \times 2}$ is the covariance of the detection, and j is the joint index. The covariance Σ is determined from the heatmaps used for keypoint estimation [3]. First, incoming streams are synchronized into sets of time-corresponding detection messages of size N, which we will refer to as *framesets* in the following. Preprocessing then rejects false, noisy, redundant, and low-quality detections, passing through only detections that are considered accurate and suitable for contributing to improving the extrinsic calibration. For this, we check different conditions for each frameset, which address the number of detections per sensor, the timestamps associated to each sensor, or the confidence value of each detection. In particular, we reject all framesets where the maximum span or standard deviation of timestamps exceeds a threshold and consider only joint detections with a minimum confidence of



Fig. 3. Data Association: 3D back-projection rays embedded in the global coordinate system for the joint detections of one person (black), the corresponding reduction to line segments after applying depth estimation (green), and the center of mass of the corresponding person hypothesis (black). 3D human pose estimation according to [3] shown for illustration purposes only.

0.6. We further require the hip and shoulder detections of each person to be valid, which is necessary for robust data association. After filtering, we use the distortion coefficients of each camera to undistort the coordinates of all valid detections using the OpenCV library [2].

3.2 Data Association

In the data association step, we find correspondences between detections from different sensors based on the current estimate of the multi-view geometry of the camera network, which can still be inaccurate. First, we back-project each 2D detection \mathcal{D} into 3D, obtaining a ray $\mathbf{p}_{\mathcal{D}}$ with undetermined depth originating at the optical center of the respective camera. Next, we reduce each ray $\mathbf{p}_{\mathcal{D}}$ to a line segment $l_{\mathcal{D}}$ by estimating the interval $[z_{\min}, z_{\max}]$ in which the depth z of each detection \mathcal{D} lies, as illustrated in Fig. 3. For the depth interval estimation, we assume a minimum and maximum torso height and width for the detected persons, derived by a specified minimum and maximum person height to be expected during calibration. The four torso keypoints (shoulders and hips) empirically are the most stable and least occluded ones, and the physical distances between them can be assumed constant for a person due to the human anatomy, independent of the pose. The respective measured distance, however, depends on the persons? orientation towards the cameras. Here, we assume worst-case orientations, leading to larger depth intervals. In summary, we do not require persons to always stand upright but support arbitrary poses and orientations towards the cameras instead. Specifying a short person height interval leads to a more constrained search space during data association, accommodating for an inaccurate initial estimate of the extrinsic calibration or a crowded scene. A wider interval, however, yields equal results in common scenarios, while supporting small and tall persons as calibration targets alike. Depth estimation is the only component in the pipeline, where human anatomy is exploited.

To find the correspondences between person detections from multiple views, we deploy an iterative greedy search method similar to the approach of Tanke et al. [30], using the distances between the previously estimated line segments as data association cost. We define the distance of two line segments l_1 , l_2 as the Closest Point-distance described by Wirtz et al. [33]:

$$d_{\text{closestpoint}}(\mathbf{l}_1, \mathbf{l}_2) = \min\left(d_{\perp}\left(\mathbf{l}_1, \mathbf{l}_2\right), d_{\perp}\left(\mathbf{l}_2, \mathbf{l}_1\right)\right) \,. \tag{3}$$

To further improve the robustness of the approach, as the extrinsic calibration is not precisely known, we iterate over all person detections, sorted in ascending depth order, utilizing the depth estimation $(z_{\min} + z_{\max})/2$ of each person detection. This exploits the fact that near person detections have a relatively short interval $[z_{\min}, z_{\max}]$ and, thus, a more constrained localization in 3D space. Lastly, we compute the *center of mass* for each person hypothesis, which serves to give a rough localization of each person hypothesis in 3D space, by averaging the center points of the line segments of all assigned torso keypoints.

The data association stage outputs a list of person hypotheses that have been observed by at least two different cameras, which will be used to construct a factor graph optimization problem in the following.

3.3 Factor Graph Optimization

The optimization stage processes the person hypotheses obtained through data association to extract knowledge about the extrinsic calibration of the utilized multi-camera system. To this end, we construct a factor graph [7] encoding projective constraints, based on a selection of person hypotheses, as well as prior knowledge on the camera poses w.r.t. the initial estimate of the extrinsic calibration or the results of previous optimization cycles.

The optimal selection of person hypotheses used in an optimization cycle is determined by a selection algorithm from all available hypotheses, ensuring an optimal spatial and temporal distribution of the observations, to obtain a well-constrained optimization problem while also maintaining a reasonable degree of entropy between selections over consecutive optimization cycles. For this, we generate a random permutation of the indices of all available person hypotheses, which is biased towards selecting newer hypotheses first. Selecting newer hypotheses with higher probability is advantageous, as their data association and center of mass are estimated more reliably, given that the extrinsic calibration improves over time. Additionally, we ensure a minimum spacing between all selected person hypotheses, w.r.t. to their center of mass, by only including the next person hypothesis within the permutation if its distance towards all previously selected person hypotheses is above a spacing threshold s = 0.2 m.

For each optimization cycle t, we construct a factor graph \mathcal{G}_t by using a selection of person hypotheses \mathcal{H}_t . A factor graph is a bipartite graph consisting of variable nodes and factor nodes [8]. Variable nodes represent the unknown random variables of the optimization problem, i.e. the 3D joint positions of the person hypotheses in \mathcal{H}_t (landmark nodes) and the considered camera poses of the multi-camera system \mathcal{S}_N (camera nodes). Factor nodes constrain the variable nodes by encoding the available knowledge about the underlying distribution



Fig. 4. Factor graph with camera variable nodes for the camera poses C_i and landmark variable nodes for the 3D person joint positions \mathcal{L}_j . Camera and landmark nodes can be connected via binary projection factors to constrain the reprojection error of a person keypoint detection. Each landmark node must be connected to at least two projection factors for allowing triangulation. All camera nodes are connected to a unary prior factor that encodes the initial uncertainty of the camera pose.

of the considered random variables. In particular, this refers to the obtained observations contained in \mathcal{H}_t as well as the resulting camera poses from previous optimization cycles. Each factor node uses a Gaussian noise model that reflects the confidence in the constraint it represents. The constructed factor graph is illustrated in Fig. 4. We equip every camera node \mathcal{C}_i^t with a unary prior factor, encoding prior knowledge about the camera pose and its uncertainty, and use binary projection factors connecting camera nodes to landmark nodes to encode observation constraints based on person keypoint detections. Projection factors calculate the reprojection error for a 2D detection w.r.t. the corresponding camera pose and landmark position using the known intrinsic parameters \mathbf{K}_i . Camera nodes are initialized with the current estimate for the extrinsic calibration (for t = 0, we use the initial estimate and for t > 0, we result of the previous time step). Landmark nodes are initialized by triangulation of the 2D observations using the latest camera geometry estimate [14]. Note, that we perform triangulation in every optimization cycle, even when using a person hypothesis that was already utilized in a previous optimization cycle. Hence, the triangulation results are updated based on the current estimate of the extrinsic calibration. We solve each factor graph \mathcal{G}_t for the most likely camera poses by applying a Levenberg-Marquardt optimization scheme, provided by the GTSAM framework [7]. A successful optimization yields a new candidate for the extrinsic calibration of \mathcal{S}_N . We forward this candidate to the refinement stage where the current estimate of the extrinsic calibration will be updated based on this candidate. The updated estimate for the extrinsic calibration will then be used for constructing and initializing the factor graph in the next optimization cycle.

3.4 Camera Pose Refinement

After each successful optimization, we obtain new candidates \hat{C}_i^t for the extrinsic calibration of a subset of cameras $i \subseteq S_N$ that were constrained by the factor graph \mathcal{G}_t . We smooth between the previous state and the new measurement using a Kalman filter to obtain the current estimate of the extrinsic calibration \mathcal{C}_i^t . As each optimization cycle contains only a limited number of observations in the factor

graph to enable real-time operation, smoothing prevents overconfidence towards a specific set of observations and improves the convergence behavior. We update the previous estimate with the result of the factor graph optimization using the marginalized uncertainty from the optimized factor graph as measurement noise. Between optimization cycles, we add a constant process noise to each predicted camera pose uncertainty, enabling convergence over longer time horizons. Finally, we prevent scaling drift of the updated extrinsic calibration by applying the scaling factor that minimizes the distance towards the initial estimate of the extrinsic calibration according to Umeyama's method [32].

4 Evaluation

We evaluate the proposed method in challenging real-world settings in our lab, a large room with an area of $\sim 240 \text{ m}^2$ and a height of 3.2 m. As the room is partly a robotics workshop and partly a desk-based workspace, it is densely filled with different objects and furniture, which can cause false detections and occlusion. The cameras are distributed throughout the room as illustrated in Fig. 5 at similar heights of around 2.6 m. We deploy 20 smart edge sensors, 16 of which are based on the Google EdgeTPU, and 4 are based on the Nvidia Jetson Xavier NX (cf. Fig. 1). Both sensor types provide person keypoint detections in identical format and will be treated in the same way during the experiments.

For evaluation, we apply our pipeline to recordings of one and two persons (1.96 m & 1.70 m) crossing the room and generating detections in all cameras over ~180 s. We repeat the experiment 10 times with different initializations and compare our results towards a reference calibration obtained by applying the kalibr toolkit [28]. We apply an initial error of 0.25 m and 10° in a random direction w.r.t. the reference calibration for all cameras C_i for i > 0 and use default parameters provided in the linked repository. We empirically verified that errors of this order of magnitude are easily attainable via manual initialization utilizing a floor plan, height measurements, and RGB images from all cameras.



Fig. 5. Sketched floor plan with camera poses for the evaluation experiments.

Table 1. Statistics of the position and orientation error towards the reference calibration averaged over 10 repetitions of the experiment.

	1 Person				2 Persons			
Error	Avg.	Std.	Min.	Max.	Avg.	Std.	Min.	Max.
Position	$0.053\mathrm{m}$	$0.030\mathrm{m}$	$0.011\mathrm{m}$	$0.119\mathrm{m}$	$0.052\mathrm{m}$	$0.030\mathrm{m}$	$0.017\mathrm{m}$	$0.122\mathrm{m}$
Orientation	0.390°	0.184°	0.120°	0.891°	0.436°	0.177°	0.154°	0.818°



Fig. 6. Evolution of mean and min–max span of (a) position and (b) orientation error towards the reference calibration. Convergence is faster when observing multiple persons.

Tab. 1 shows the statistics of the final position and orientation error distributions towards the reference calibration averaged over all repetitions of the experiments with one or two persons present in the scene, respectively. The position error is obtained by rigid alignment of the calibration result towards the reference according to Umeyama's method [32] without rescaling. The orientation error is computed as the angle between two orientations via the shortest arc [17]. We do not observe a significant difference in the final result between calibrating with one or two persons. However, convergence is faster in the two-person case, as all cameras provide detections earlier in the procedure.

Fig. 6 shows the evolution of the error over time for one exemplary repetition of the resp. experiment with one or two persons. The majority of the convergence takes place in the first ~ 50 optimization cycles or ~ 35 s and after ~ 100 optimizations, the camera poses and errors remain stable in the two-person experiment. With only a single person, convergence is slower. Observations from all cameras are obtained after ~ 110 optimization cycles and it takes ~ 150 iterations for the poses to remain stable.

Tab. 2 shows the final position error for different initialization errors with two persons. Convergence remains stable for initial errors up to 35 cm and 15° but becomes less reliable for larger errors. In particular, the likelihood of the camera poses being stuck in a local minimum consistent with queued person hypotheses containing false data association increases with larger initialization errors, as the accuracy of the data association relies on the geometry of the provided initialization.

Additionally, we compare reprojection errors, measured using two different evaluation pipelines, using the calibration obtained from our experiments with

Initial Error	$0.10\mathrm{m},5^\circ$	$0.25\mathrm{m},10^\circ$	$0.35\mathrm{m},15^\circ$	$0.50\mathrm{m},20^\circ$	$0.75m,20^\circ$
Final Error	0.050 m	0.052 m	$0.067\mathrm{m}$	0.118 m	0.214 m
Std.	$0.003{ m m}$	0.030 m	$0.011\mathrm{m}$	$0.073{ m m}$	$0.164\mathrm{m}$

Table 2. Position error towards the reference calibration for different initial errors.



Fig. 7. Comparison of the reprojection error per camera between our method and the reference calibration using (a) keypoint- and (b) marker-based evaluation pipelines.

two persons. The first evaluation processes keypoint detections for 3D human pose estimation [3] and matches the domain in which our calibration was obtained. Here, we use a distinct recording unseen during the calibration. The second pipeline uses a sequence of multi-view images of the AprilTag grid used to obtain the reference calibration [28] and, thus, matches its data domain. In general, the keypoint-based evaluation is biased towards our keypoint-based calibration method, while the AprilTag evaluation is biased towards the reference calibration.

Fig. 7 shows the reprojection error per camera for both evaluation pipelines and Tab. 3 reports the averaged reprojection error. For the keypoint-based evaluation, we observe that our calibration achieves lower reprojection errors for all but two cameras. For the marker-based evaluation, our calibration achieves similar reprojection errors as for the keypoint-based pipeline, while the reprojection errors of the reference calibration are significantly lower. Our flexible, marker-free method achieves lower reprojection errors for the envisaged application of 3D multi-person pose estimation and still achieves a coherent result when evaluating with a traditional calibration target. The difference in accuracy for the second evaluation is mainly due to our method being marker-free using features from persons of unknown dimensions for calibration, while the reference method knows the exact scale of the calibration (and evaluation) target. Also, the noise in the joint detections may be larger than for the tag detections.

The averaged reprojection error per joint group for the keypoint-based evaluation is shown in Tab. 4. Our method achieves lower reprojection errors in all categories. The reprojection error is larger for faster-moving joints like ankles and wrists, while it is smaller for more stable joints. This can be explained by limitations in the synchronization within framesets.

It is worth noting that the measured reprojection error does not exclusively originate from the provided extrinsic calibration, but also from other factors, e.g. the intrinsic camera calibration, or the approach for detection, data association, and triangulation.

 Table 3. Comparison of the average reprojection error of our method and the reference calibration for keypoint- and AprilTag-based evaluations, averaged over 10 repetitions.

Calibration	Keypoints	AprilTag Grid
Reference	$4.57\mathrm{px}$	$1.95\mathrm{px}$
Our Method	4.01 px	$5.00\mathrm{px}$

Table 4. Comparison of the reprojection errors per joint group between our method and the reference calibration averaged over 10 repetitions of the experiment.

Calibration	Head	Hips	Knees	Ankles	Shlds	Elbows	Wrists
Reference	4.02 px	$5.10\mathrm{px}$	4.83 px	5.88 px	3.61 px	$4.29\mathrm{px}$	$5.14\mathrm{px}$
Our Method	$3.55\mathrm{px}$	4.28 px	4.20 px	$5.27\mathrm{px}$	$3.21\mathrm{px}$	$3.75\mathrm{px}$	$4.55\mathrm{px}$

5 Conclusion

In this work, we developed a marker-free online method for extrinsic camera calibration in a scene observed by multiple smart edge sensors, relying solely on person keypoint detections. The keypoint detections are fused into 3D person hypotheses at a central backend by synchronization, filtering, and data association. Factor graph optimization problems are repeatedly solved to estimate the camera poses constrained by the observations. Knowledge about the camera poses obtained through the optimization of one factor graph is used during the construction of the next factor graph, enabling the accumulation of knowledge and the convergence of all cameras towards an accurate pose. Lastly, the convergence behavior is improved by a refinement scheme based on a Kalman filter.

Our method is designed to be robust against false or sparse sets of detections and occlusions, and is free of many typical assumptions of similar methods: It does not require a specific calibration target, can cope with and exploit the detections of multiple persons simultaneously, and handles arbitrary person poses. We evaluate the proposed method in a series of experiments and compare our calibration results to a reference calibration obtained by an offline calibration method based on traditional calibration targets. We show that our calibration results are more accurate than the reference calibration by reliably achieving lower reprojection errors in a 3D multi-person pose estimation pipeline used as application scenario. Not only provides our method a quick and easy-to-use calibration utility, but it also achieves state-of-the-art accuracy.

The limitations of our method are mainly related to scaling ambiguity and data association: The scale of the initial estimate of the extrinsic calibration is maintained throughout the calibration procedure. It inherently cannot resolve the scale in case the initial estimate is biased or inaccurate, as the dimensions of the persons used as calibration targets are unknown. Data association could be improved for inaccurate initial estimates of the extrinsic calibration by using visual re-id descriptors. We assume the intrisic camera calibration to be known. Our method could be extended by including intrinsic camera parameters in the optimization. Finally, our method could be extended by also using additional environment features for calibration.

References

- Bhardwaj, R., Tummala, G.K., Ramalingam, G., Ramjee, R., Sinha, P.: AutoCalib: Automatic traffic camera calibration at scale. ACM Transactions on Sensor Networks (TOSN) 14(3-4), 1–27 (2018). https://doi.org/https://doi.org/10.1145/3199667
- 2. Bradski, G.: The OpenCV Library. Dr. Dobb's Journal of Software Tools (2000)
- Bultmann, S., Behnke, S.: Real-time multi-view 3D human pose estimation using semantic feedback to smart edge sensors. In: Robotics: Science and Systems XVII (RSS) (2021). https://doi.org/10.15607/rss.2021.xvii.040
- Bultmann, S., Behnke, S.: 3D semantic scene perception using distributed smart edge sensors. In: IEEE International Conference on Intelligent Autonomous Systems (IAS) (2022). https://doi.org/10.48550/ARXIV.2205.01460
- Cao, Z., Hidalgo, G., Simon, T., Wei, S.E., Sheikh, Y.: OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 43(1), 172–186 (2021). https://doi.org/10.1109/TPAMI.2019.2929257
- Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR). vol. 1, p. 886–893 (2005). https://doi.org/10.1109/CVPR.2005.177
- Dellaert, F.: Factor graphs and GTSAM: A hands-on introduction. Tech. Rep. GT-RIM-CP&R-2012-002, Georgia Institute of Technology (2012), https://research. cc.gatech.edu/borg/sites/edu.borg/files/downloads/gtsam.pdf
- Dellaert, F., Kaess, M.: Factor graphs for robot perception. Foundations and Trends in Robotics (FNT) 6(1-2), 1–139 (2017). https://doi.org/10.1561/2300000043
- Felzenszwalb, P.F., Huttenlocher, D.P.: Pictorial structures for object recognition. International Journal of Computer Vision (IJCV) 61(1), 55–79 (2005). https://doi.org/10.1023/B:VISI.0000042934.15159.49
- Fischler, M., Elschlager, R.: The representation and matching of pictorial structures. IEEE Transactions on Computers C-22(1), 67–92 (1973). https://doi.org/10.1109/T-C.1973.223602
- Fischler, M.A., Bolles, R.C.: Random Sample Consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM 24(6), 381–395 (1981). https://doi.org/10.1145/358669.358692
- 12. Google: EdgeTPU dev board (2020), https://coral.ai/docs/dev-board/datasheet, accessed: 2022-03-25
- Guan, J., Deboeverie, F., Slembrouck, M., Van Haerenborgh, D., Van Cauwelaert, D., Veelaert, P., Philips, W.: Extrinsic calibration of camera networks based on pedestrians. Sensors 16(5) (2016). https://doi.org/10.3390/s16050654
- Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, USA, 2 edn. (2003)
- Henning, D., Laidlow, T., Leutenegger, S.: BodySLAM: Joint camera localisation, mapping, and human motion tracking. In: European Conference on Computer Vision (ECCV) (2022). https://doi.org/10.48550/ARXIV.2205.02301
- Hödlmoser, M., Kampel, M.: Multiple camera self-calibration and 3D reconstruction using pedestrians. In: International Symposium on Visual Computing (ISVC) (2010). https://doi.org/10.1007/978-3-642-17274-8_1
- Huynh, D.Q.: Metrics for 3D rotations: Comparison and analysis. Journal of Mathematical Imaging and Vision 35(2), 155–164 (2009)
- Komorowski, Jacekand Rokita, P.: Extrinsic camera calibration method and its performance evaluation. In: Computer Vision and Graphics. pp. 129–138 (2012). https://doi.org/10.1007/978-3-642-33564-8_16

Online Marker-free Extr. Camera Calib. using Person Keypoint Detections

15

- Li, J., Wang, C., Zhu, H., Mao, Y., Fang, H.S., Lu, C.: CrowdPose: Efficient crowded scenes pose estimation and a new benchmark. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10855–10864 (2019). https://doi.org/10.1109/CVPR.2019.01112
- Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft COCO: Common objects in context. In: European Conference on Computer Vision (ECCV). pp. 740–755 (2014). https://doi.org/10.1007/978-3-319-10602
- Liu, J., Collins, R.T., Liu, Y.: Surveillance camera autocalibration based on pedestrian height distributions. In: British Machine Vision Conference (BMVC). p. 144 (2011). https://doi.org/10.5244/C.25
- Liu, J., Collins, R.T., Liu, Y.: Robust autocalibration for a surveillance camera network. In: IEEE Workshop on Applications of Computer Vision (WACV). pp. 433–440 (2013). https://doi.org/10.1109/WACV.2013.6475051
- Lowe, D.: Object recognition from local scale-invariant features. In: IEEE International Conference on Computer Vision (ICCV). vol. 2, pp. 1150–1157 vol.2 (1999). https://doi.org/10.1109/ICCV.1999.790410
- Lv, F., Zhao, T., Nevatia, R.: Camera calibration from video of a walking human. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 28(9), 1513–1518 (2006). https://doi.org/10.1109/TPAMI.2006.178
- Maye, J., Furgale, P., Siegwart, R.: Self-supervised calibration for robotic systems. In: IEEE Intelligent Vehicles Symposium (IV). pp. 473–480 (2013). https://doi.org/10.1109/IVS.2013.6629513
- NVIDIA: Nvidia Jetson Xavier NX Developer Kit (2020), https://developer.nvidia. com/embedded/jetson-xavier-nx-devkit, accessed: 2022-03-25
- Olson, E.: AprilTag: A robust and flexible visual fiducial system. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 3400–3407 (2011). https://doi.org/10.1109/ICRA.2011.5979561
- Rehder, J., Nikolic, J., Schneider, T., Hinzmann, T., Siegwart, R.: Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes. In: IEEE International Conference on Robotics and Automation (ICRA). pp. 4304–4311 (2016). https://doi.org/10.1109/ICRA.2016.7487628
- Reinke, A., Camurri, M., Semini, C.: A factor graph approach to multi-camera extrinsic calibration on legged robots. In: IEEE International Conference on Robotic Computing (IRC). pp. 391–394 (2019). https://doi.org/10.1109/IRC.2019.00071
- Tanke, J., Gall, J.: Iterative greedy matching for 3D human pose tracking from multiple views. In: DAGM German Conference on Pattern Recognition (GCPR) (2019). https://doi.org/10.1007/978-3-030-33676-9_38
- Truong, A.M., Philips, W., Guan, J., Deligiannis, N., Abrahamyan, L.: Automatic extrinsic calibration of camera networks based on pedestrians. In: IEEE International Conference on Distributed Smart Cameras (ICDSC) (2019). https://doi.org/10.1145/3349801.3349802
- Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 13(04), 376–380 (1991). https://doi.org/10.1109/34.88573
- 33. Wirtz, S., Paulus, D.: Evaluation of established line segment distance functions. Pattern Recognition and Image Analysis 26, 354–359 (2016). https://doi.org/10.1134/S1054661816020267
- Zhang, Z.: A flexible new technique for camera calibration. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI) 22(11), 1330–1334 (2000). https://doi.org/10.1109/34.888718