

# A Robust Extrinsic Calibration Framework for Vehicles with Unscaled Sensors

Celyn Walters, Oscar Mendez, Simon Hadfield, Richard Bowden<sup>1</sup>

*Abstract*—Accurate extrinsic sensor calibration is essential for both autonomous vehicles and robots. Traditionally this is an involved process requiring calibration targets, known fiducial markers and is generally performed in a lab. Moreover, even a small change in the sensor layout requires recalibration. With the anticipated arrival of consumer autonomous vehicles, there is demand for a system which can do this automatically, after deployment and without specialist human expertise.

To solve these limitations, we propose a flexible framework which can estimate extrinsic parameters without an explicit calibration stage, even for sensors with unknown scale. Our first contribution builds upon standard hand-eye calibration by jointly recovering scale. Our second contribution is that our system is made robust to imperfect and degenerate sensor data, by collecting independent sets of poses and automatically selecting those which are most ideal.

We show that our approach’s robustness is essential for the target scenario. Unlike previous approaches, ours runs in real time and constantly estimates the extrinsic transform. For both an ideal experimental setup and a real use case, comparison against these approaches shows that we outperform the state-of-the-art. Furthermore, we demonstrate that the recovered scale may be applied to the full trajectory, circumventing the need for scale estimation via sensor fusion.

## I. INTRODUCTION

Autonomous vehicles will be one of the most lifestyle-changing advancements to result from today’s Robotics and Computer Vision research. They have the potential to solve congestion, reduce fatalities and environmental impact. In order for these agents to operate safely in the real world, they need to interpret their surroundings and robustly localize with a high degree of accuracy. To this end, modern robotic platforms employ a variety of sensor types, such as stereo/monocular cameras, Inertial Measurement Units (IMUs), Global Positioning Sensors (GPSs) and wheel odometry combined via sensor fusion.

While using multiple sensors provides more robust perception, it is extremely important that the sensors are properly calibrated. Intrinsic calibration ensures that each individual sensor’s output is accurate, and for many commercial sensors, this calibration is performed in the factory. However, the extrinsic calibration between sensors must be carried out once they are mounted in their final positions. Generally, this action is avoided by mounting the sensors to a structure with precisely known dimensions. However, this puts restrictions on the sensor layout and ease of manufacture. Where this is not possible, the traditional method is to make use of external reference frames [1], [2]. Such a procedure is

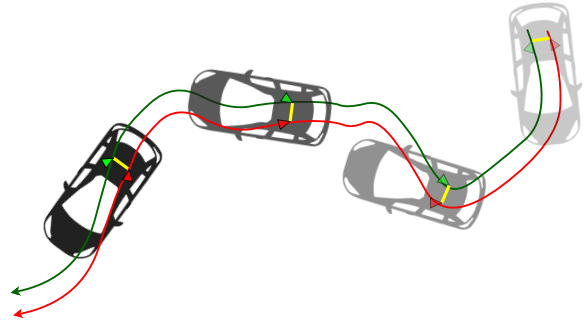


Fig. 1. Green and red sensor trajectories are used to recover the fixed extrinsic transform between them, shown as yellow.

relatively simple to carry out for a test platform, but platforms designed for continuous operation are susceptible to changes in calibration due to vibration and slip in mountings. This necessitates manual re-calibration, which is inconvenient for larger platforms and vehicles. Furthermore, a typical consumer will not have access to a calibration environment nor have the expertise to reliably carry out a complicated procedure.

An attractive alternative is calibration directly from sensor motion. Hand-eye calibration techniques [3], [4], [5] estimate the transform between two rigidly-attached sensors using their poses, as in fig. 1. It was originally formulated as a way to recover the transformation between a robotic arm and a mounted camera. It removes the need for overlapping views and may be broadly applied to visual and nonvisual sensors. However, there are several limitations to this approach that prevent it from being used as an automatic calibration tool in the field:

- 1) **Assumes known scale** — A prerequisite for hand-eye calibration is that the measurements are to a consistent scale, usually metric. This is not the case for monocular cameras, and the common solution is to rely on visible external reference frames or to rescale the sensor using another metric sensor such as GPS.
- 2) **Corruption by degenerate motion** — For a complete description of the relative transformation, hand-eye calibration requires stimulation of all degrees of freedom. Without appropriate motion, parts of the extrinsic transform will be unobservable. It is impractical for a consumer to be required to carry out calibration manoeuvres, but the necessary movements may occur during prolonged natural use. The truly unobservable parts, such as impossible vehicle motion, are by nature not important for sensor fusion-based odometry.

<sup>1</sup>C. Walters, S. Hadfield, O. Mendez and R. Bowden are with CVSSP at the University of Surrey, UK c.walters@surrey.ac.uk

**3) Vulnerable to drift** — Individual sensors exhibit compounded error in the form of drift. While additional data can increase the accuracy of the extrinsic calculation, acquiring data over a longer period exacerbates drift.

In this paper, we present a single flexible solution which can overcome all of these issues. **(1)** is addressed by simultaneously solving for the extrinsic transform and its corresponding scale by optimizing a similarity matrix. This allows monocular sensors to be calibrated in the same way as any metric sensor. This means that algorithms which operate on unscaled sensors such as Visual Odometry (VO) may be combined with other data sources in a metric sensor fusion framework. To address **(2)**, the calibration runs in real-time without explicit user intervention. The agent does not need to perform specialized movements, as the framework updates each parameter when it is observable, providing robustness against periods of insufficient movement or lack of visual features. For **(3)**, our framework mitigates the effects of drift caused by poor visibility and noise, which commonly interfere with VO & IMU processing. We exploit the fact that data is often locally accurate, and combine several overlapping sections of the trajectory which contain less drift. In addition, the calibration discards the sections containing discontinuities that may be generated, for example, by GPS corrections, or by a loop closure in a visual odometry system.

Since direct measurements are not used, the calibration is indifferent to the source of data, allowing remarkable flexibility. For instance, a VO/Simultaneous Localisation and Mapping (SLAM) algorithm may be freely chosen.

This paper demonstrates how our framework achieves accurate extrinsic measurements even when presented with imperfect input data. Section II surveys other extrinsic calibration approaches and the foundations for this work. Section III-A describes how preliminary estimates for transformation and scale are generated. In section III-B, the estimates are selected and processed to account for measurement inaccuracies and degenerate motion. Finally, the performance is evaluated in sections IV-A to IV-D against competing approaches.

## II. RELATED WORK

A simple approach to general extrinsic calibration is described by Zhao *et al.* [2], in which markers attached to camera housings are simultaneously viewed by an external ‘support’ camera. This way, even non-overlapping cameras can be simply calibrated without degenerate cases. Although flexible, the method depends on exact measurements between each sensor and its attached marker which may be difficult to obtain in practice.

Traditional approaches to extrinsic calibration exploit sensor-specific capabilities to maximize accuracy. Zhang and Pless use a calibration grid to retrieve the extrinsic transform between a camera and a 2D laser range finder [1]. Velas *et al.* replace the calibration grid with a novel 3D marker [6]. For intrinsic and extrinsic calibration of IMUs and RGB cameras, the Kalibr calibration toolbox by Furgale *et al.* may be used [7], [8]. All of these rely on a dedicated calibration procedure, which may not be practical for a large vehicle that

cannot undergo certain motions. A more automated method is used in our framework, in which the calibration data is obtained automatically during normal operation.

An approach which is less dependent on specific visible cues is to instead use natural features. Ling and Shen achieve this for finding the offset between each sensor in a stereo camera [9]. For monocular cameras, Ataer-Cansizoglu *et al.* exploit a previously generated SLAM model, where 2D-3D correspondences are formed from the images to find the global camera positions [10]. More closely related to our work is the approach of Maye *et al.*, in which the self-calibration is online and robust to small motion [11]. Ours differs as it does not directly use landmarks nor assume these are available.

Most algorithms which directly process visual features are restricted to those which use known geometry or where there are overlapping views. Our approach, which builds on hand-eye calibration, requires only relative sensor pose information. Solutions for hand-eye calibration were published as early as 1989 by Shiu and Ahmad [5]. Early methods estimate the rotation and translation independently. Traditionally, rigid transformation is decomposed in this way to follow the rotation-then-translation formalization. This leads to a common perception that their relationship is separate. In [3], Chen argues that decoupling the two adversely affects the generality and efficacy of the algorithm. They approach the problem with screw motion to unify them as a single parameter. Daniilidis also jointly parameterizes rotation and translation, this time using dual quaternions [4].

The problem of scale is often overlooked, as it can usually be obtained through a separate process. Although Schmidt *et al.* extend [4] by encoding scale as the norm of the dual quaternion [12], the rotation, scale, and translation are obtained separately in evaluation. The relative error also seems high and the scale error is not provided. Heller *et al.* use second-order cone programming to recover the scaled translation component. Again, the rotation is separately calculated, and convergence time is in the order of minutes. In contrast to the above approaches, we opt to estimate both the extrinsic transform and corresponding scale in the same operation in real-time.

Vehicles with restricted motion modes cannot stimulate all the necessary degrees of freedom required for a full hand-eye calibration. For example, for wheeled vehicles which can only travel with planar motion and evidence of rotation is limited to a single axis, a sensor’s height is unobservable. Ruland *et al.* and Heng *et al.* both apply hand-eye calibration for cameras in a vehicular context, although they ignore the estimation of relative heights, or rely on a separate technique [13], [14]. Heng *et al.* utilizes the dual-quaternion approach from [4] and makes use of the visual features between views as a refinement in an offline process. In our work, all parameters come from the same optimization and are updated if and when they are recoverable.

Zhu *et al.*’s process extends hand-eye calibration by exploiting generic planes using Structure from Motion (SfM) [15]. This allows calibration of non-overlapping views from multiple cameras. Since our method does not assume

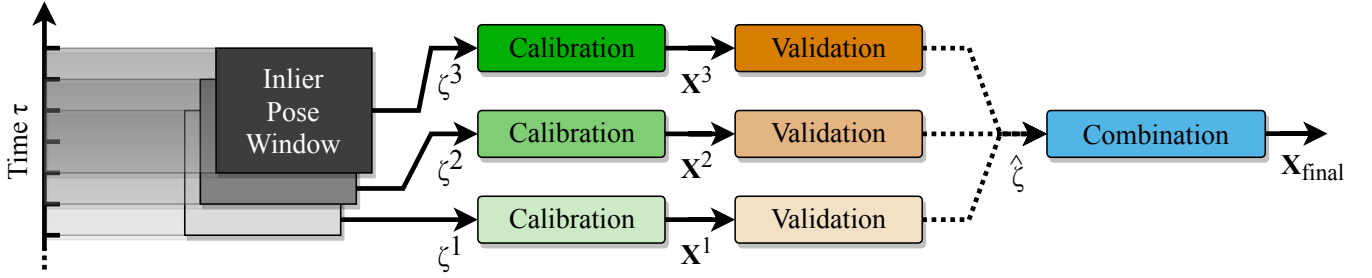


Fig. 2. Systems diagram of calibration framework.

visual information is available, cameras may be treated as black boxes which provide unscaled pose information. This is advantageous in the case of cameras which process VO/SLAM algorithms on-board, and allows simpler data connections for a more modular setup.

### III. METHODOLOGY

State-of-the-art approaches cater for specific use-cases, and often rely on particular sensors. Instead, we propose an extrinsic calibration framework, laid out in fig. 2, which allows fully modular configuration of sensors and their pose-estimation algorithms. Its operation may be summarized as follows: First, sets of time-aligned poses from two arbitrary sensors are collected (black blocks, fig. 2). Estimates for the extrinsic calibration between the sensors are obtained using each pose set (green blocks), and described in section III-A. The sets are filtered based on the type of motion and the quality of the estimate determined (orange blocks), and is described in detail in section III-B. Finally, the inlier estimates are continuously combined (blue block) for a stable calibration result which is refined over time.

#### A. Extrinsic Calibration

There are multiple formulations for hand-eye calibration but the most widely-used solves the following equation:

$$\mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{B}, \quad (1)$$

where  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{X}$  are homogeneous transformations in the form

$$\mathbf{A}, \mathbf{B} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ 0^T & 1 \end{bmatrix}, \quad \mathbf{X} = s \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ 0^T & 1/s \end{bmatrix}. \quad (2)$$

$\mathbf{A}$  and  $\mathbf{B}$  represent the pose for the first and second sensors respectively. By rearranging eq. (1) into  $\mathbf{A} = \mathbf{X}\mathbf{B}\mathbf{X}^{-1}$ , it can be seen that the pose from sensor  $\mathbf{B}$  transformed by the correct transformation  $\mathbf{X}$  will be identical to the pose from sensor  $\mathbf{A}$ . Assuming that  $\mathbf{A}$  and  $\mathbf{B}$  are rigidly attached, the same transform  $\mathbf{X}$  will solve for a series of  $\mathbf{A}$  and  $\mathbf{B}$  poses. Therefore,  $\mathbf{X}$  represents the constant unknown transformation that maps the reference frame of  $\mathbf{A}$  to that of  $\mathbf{B}$ . This concept is shown in fig. 3.

For calibration, a series of time-synchronized poses  $\zeta_A$  and  $\zeta_B$  are recorded for each sensor relative to  $\mathbf{O}_A$  and  $\mathbf{O}_B$  respectively, where

$$\begin{aligned} \zeta_A &= \{\mathbf{A}_0, \mathbf{A}_1 \cdots \mathbf{A}_\tau\} \\ \zeta_B &= \{\mathbf{B}_0, \mathbf{B}_1 \cdots \mathbf{B}_\tau\}, \end{aligned} \quad (3)$$

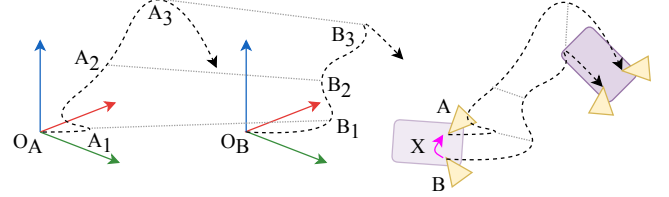


Fig. 3. Concept of hand-eye calibration for two cameras  $A$  and  $B$ . Poses  $\mathbf{A}_{1-3}$  and  $\mathbf{B}_{1-3}$  are recorded in an unknown frame relative to their original position  $\mathbf{O}_A$  and  $\mathbf{O}_B$ . When correctly placed, their trajectories have constant relative transformation  $\mathbf{X}$ .

and  $\mathbf{O}$  is defined as the initial position of each set of poses  $\zeta$ , such that

$$\mathbf{O}_A \mathbf{A}_0^{-1} = \mathbf{I}, \quad (4)$$

where  $\mathbf{I}$  is the identity transform matrix. The reason that this is necessary is that the sets of poses  $\zeta_A$  and  $\zeta_B$  are independent and are defined in their own arbitrary reference frame. It is not important what sensor this data originates from, which gives rise to the flexibility of the approach. Poses from cameras may be obtained through any VO or SLAM algorithm, which are abundantly available [16], [17], [18]. For a given autonomous/robotic platform, it is likely that such an algorithm already needs to be run for sensor fusion and mapping purposes. This being the case, there is no additional computational cost in terms of obtaining these poses. Similarly, there are odometry solutions available for lidar scanners, and poses may also be obtained from IMUs, GPS or wheel odometry.

Hand-eye calibration is performed using  $\zeta_A$  and  $\zeta_B$ , resulting in an estimate for the extrinsic transformation. To optimize for scale in the same calibration step, it is necessary to be able to incorporate a scale parameter into each input measurement  $\mathbf{A}_\tau$  and  $\mathbf{B}_\tau$ . For this reason, we use a similarity matrix. This necessitates the use of a nonlinear optimization, using the Frobenius norm in the cost function

$$h(\mathbf{X}|\zeta_A, \zeta_B) = \sum_{\tau} \|\mathbf{A}_\tau \mathbf{X} - \mathbf{X} \mathbf{B}_\tau\|. \quad (5)$$

$\mathbf{A}_\tau$  and  $\mathbf{B}_\tau$  are the synchronized poses of the first and second sensor respectively, and  $\mathbf{X}$  is the extrinsic estimate. The problem is minimized using the Levenberg-Marquardt algorithm initialised with identity transform or, if one exists, the previously known calibration. If both sensors are metric,

the result of scale parameter  $s$  from eq. (2) is usually near 1.0. Fixing the value reduces the degrees of freedom and allows for higher accuracy.

### B. Robust Estimation Pipeline

In a real-world scenario, data is subject to noise and drift. The minimization does not make adjustments to the input poses as this would require sensor-specific techniques such as Bundle Adjustment (BA) for cameras and bias estimation for IMUs. Optimizing over longer periods makes it more likely that all 6 degrees of freedom are exercised, enabling a fully observable solution. On the other hand, this also incorporates more error due to drift.

The following describes our methods which bring robustness to our scaling hand-eye calibration, allowing it to run in a deployed system without supervision. To allow isolation of recorded data which is subject to drift or discontinuity, intermediate calibrations are repeatedly carried out over a sliding window, shown by the vertical axis in fig. 2. This gives frequent estimates  $\mathbf{X}^i$  of the extrinsic transformation,

$$\mathbf{X}^i = \arg \min_{\mathbf{X}} (h(\mathbf{X} | \zeta_A^i, \zeta_B^i)), \quad (6)$$

where  $h$  is our scaling hand-eye calibration and  $i$  is the iteration.

For a complete solution including all axes of translation, the input data needs to exercise rotational motion in each of the X, Y, and Z axes. For example, in a movement with no rotation, all points rigidly attached to an agent produce identical trajectories relative to their starting positions. In this case, there are infinitely many values of  $X$  which have an identical minimizer cost. Many vehicles do not rotate freely and spend the majority of the time moving in ways which do not allow a full extrinsic calibration. By selecting windows  $\zeta^i$  based on their eigenvalues, we can ensure that they contain enough motion. For nonholonomic vehicles, the ratio between eigenvalues computed over position indicates rotation. This allows our system to automatically determine which parts of the extrinsic transformation can reliably be estimated and disable calibration when stationary.

Some windows inevitably contain errors or degenerate motion, leading to incorrect extrinsic estimates. Early rejection is applied to remove the most severe failures by thresholding their respective optimizer costs. When problematic input results in a low cost, the produced translation or orientation is an obvious outlier and is removed using RANSAC. Overlapping windows allows a greater number of calibrations to be obtained, giving stability to RANSAC. Although pose pairs may be incorporated into multiple overlapping windows, each is relative to a different reference frame and therefore estimates may be treated independently. The mean translation and rotation over all inlier windows consolidates and refines the accumulated evidence over longer periods.

In contrast to other calibration approaches, ours can take advantage of existing knowledge. It is often possible to physically measure the distance between sensors. The function

in eq. (6) may be extended with a regularizer term,

$$\mathbf{X}_{\text{final}} = \arg \min_{\mathbf{X}} (h(\mathbf{X} | \hat{\zeta}_A, \hat{\zeta}_B) + \alpha \omega), \quad (7)$$

where  $\alpha$  is a weight and  $\omega$  is the difference between the Euclidean distance of the extrinsic estimate and the scalar measurement. This results in a significantly tighter estimate cluster and may even overcome unobservable axes. Since the transformation is not expected to dramatically change, the previous calibration can also be used as a prior guess. Even so, accumulation of evidence allows the system to recover from situations such as sensor mountings shifting during operation.

In summary, this section introduced the concept of hand-eye calibration and how it applies generally to the task of recovering relative extrinsic transforms. It also highlighted several limitations which reduce the effectiveness of traditional hand-eye calibration in real scenarios. These include calibrating with flawed input data, and compromises between gathering sufficient motion and minimizing drift. Methods by which our framework addresses these limitations were explained. The following section shows that these solutions are effective in comparison to previously published work.

## IV. RESULTS

This section demonstrates that the capabilities of our framework allow it to achieve competitive accuracy with other hand-eye calibration methods when exposed to ideal movement and known scale. In the absence of these, we show that the approach significantly outperforms other techniques. The performance of the framework is evaluated in three scenarios: firstly, for a hand-held platform to provide nondegenerate movement; secondly, for a ground-based robotic platform using monocular vision restricted to planar motion; and finally, using a well-known benchmark with multiple sensor types to test performance in real-world conditions.

### A. Effect of Baseline

A test was performed to examine the effect of baseline length between sensors has on hand-eye calibration position accuracy. Two matching stereo cameras were fixed to a rigid bar at different spacings, with translation constrained to a single axis and no rotation. Camera motion was extracted using stereo ORB-SLAM2 [17] in a feature-rich environment with smooth motions and each test was repeated five times. The even spread around the correct distance at each separation in fig. 4 shows that wider baselines do not significantly affect the accuracy, and it is not valid to evaluate calibration error as a percentage of the ground truth. Hand-eye calibration has stronger dependence on the accuracy of input poses.

### B. Twin Stereo Cameras

To evaluate real-world performance in the best case, the data needs to be in metric scale and perform motion which exercises all six degrees of freedom to make sure all parameters are observable. The input data should be as error-free as possible. This demands slow, smooth movement with well-lit, static surroundings with plenty of features. To fulfil these needs, a platform was assembled for dataset capture.

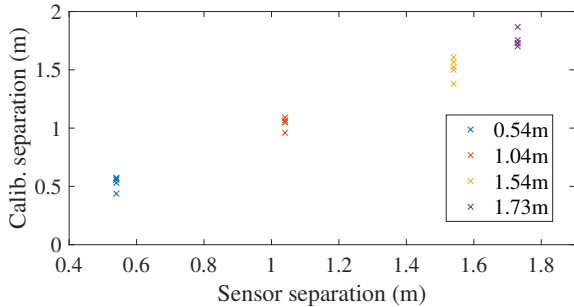


Fig. 4. Position of calibration results at different baselines.

TABLE I  
TRANSLATION AND ROTATION CALIBRATION ERROR, AND EVALUATION OF RESULTING TRAJECTORIES USING THE CALIBRATION

Calibration	T (m)	R (°)	ATE (m)	RPE (m, °)
Schmidt [4] (full)	0.2118	2.68	0.188	0.075 2.07
Ours (full)	0.0620	3.21	0.159	0.049 2.15
Schmidt (robust)	0.0270	1.57	0.162	0.039 1.71
Ours (robust)	<b>0.0105</b>	1.49	<b>0.147</b>	0.039 <b>1.64</b>
Ours (scale)	0.0276	<b>1.41</b>	0.160	<b>0.038</b> 1.65

Two matching stereo cameras (Stereolabs ZED cameras) were mounted to a rigid base, shown in fig. 5.

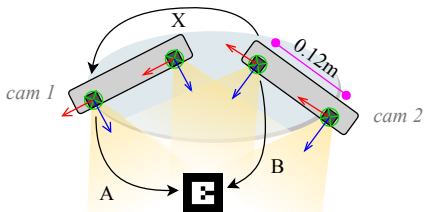


Fig. 5. Overview of stereo camera layout. Relative 3D rotation is arbitrary.

The ground truth transforms  $A$  and  $B$  between each camera and the ARTag were obtained using a tracking library [19]. These were used to calculate the ground truth relative transform  $X$  between cameras. The overlapping views and ARTag are only for ground truth measurement; they do not benefit the calibration. To isolate VO error, poses from  $cam 1$  were obtained using stereo VO [20], and poses from  $cam 2$  were the same transformed by the ground truth. The base was moved to excite all degrees of freedom over 90 seconds. Predictably, all hand-eye calibration approaches give very close to zero error. However, using independent VO for  $cam 2$ 's poses instead shows a significant loss of accuracy and the necessity for robustness.

Table I lists the distance and shortest rotation from the ground truth for the following methods: 1) The dual quaternion method of [4] over the entire trajectory, 2) our Frobenius norm parameterization as shown in eq. (5), 3) [4] using our robust pipeline, 4) our parameterization using the robust pipeline, and 5) our robust method with scale optimization. The root mean squared error of Absolute Trajectory Error (ATE) and Relative Pose Error (RPE) [21] is calculated using the trajectory from  $cam 1$  and  $cam 2$  after transforming by the

inverse extrinsic transform result. Drift and small loop closures (which are widespread in robotic vehicles) corrupt the global trajectories despite efforts to minimise them. As with other hand-eye calibration techniques, the dual quaternion approach has been shown to be highly accurate [4], [12], [14], but only in situations with significantly less measurement noise. Table I shows our approach is more suitable in this area, and also that redundantly estimating scale does not detract from the overall accuracy.

### C. Turtlebot

A second dataset was required to demonstrate our approach's benefits to automatic scale estimation and robust self-calibration in the context of robotics. We used the dataset from Mendez *et al.* [22], consisting of a ground-based robotic platform, Turtlebot, travelling with periods of rotation and of linear motion. A forward-facing Microsoft Kinect depth sensor was mounted with a known relative transform, to be retrieved by our calibration. The ground truth trajectory was generated using Labbé and Michaud's RGB-D SLAM algorithm [23], and registered to the floorplan. With yaw rotation only, the relative sensor heights are unobservable. However, as the sensor outputs are relative to the plane of motion, the height difference is inconsequential for applications such as sensor fusion. The pose sources were the wheel odometry fused with gyroscope, and monocular VO (SVO 2.0 [16]) using the Kinect Red-Green-Blue (RGB) image. Our calibration does not use the depth information. Monocular VO typically derives its scale from the mean depth of tracked points and therefore depends on the dynamic scene view. Without incorporating global optimization such as bundle adjustment, the VO exhibits scale variation over a sequence and every time it reinitializes. In contrast, our calibration derives scale from relative transformation between sensor positions which is static.

A qualitative evaluation of the scale recovery can be made with fig. 6. The initial scale of the VO (blue) has been manually matched to metric. Its separate trajectory on the right of the figure is caused by positive scale drift, and loss of visual tracking causes the sharp corner near the end. By applying our continuous scale estimation (red), the path closely resembles the ground truth (green) and makes it through the doorway. The final section is somewhat recovered, though it is heavily dependent on the tracking quality. Global drift is not an issue when being fused with an absolute positional reference, such as GPS.

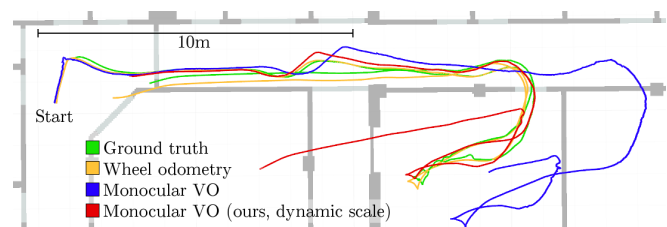


Fig. 6. Turtlebot path from different sources overlaid onto the floorplan. The dynamically scaled version of monocular VO brings it much closer to the ground truth.

TABLE II

TRANSLATION ERROR WHEN CALIBRATING MONOCULAR VO TO WHEEL ODOMETRY + GYRO, FOR TWO DIFFERENT SEQUENCE LENGTHS

Calibration	T (m)	R (°)	ATE (m)	RPE (m, °)
Schmidt [4]	0.3126	1.28	0.3175	0.6535 27.24
Schmidt (scaled)	0.1116	1.73	0.0987	0.6507 25.91
Ours	<b>0.0139</b>	1.31	<b>0.0755</b>	<b>0.6391</b> 26.95
Ours (6Dof)	0.0239	<b>1.04</b>	0.0924	0.6396 <b>24.34</b>

The unfiltered intermediate and combined calibrations are plotted over time in fig. 7. The noise in the unfiltered estimates clearly shows the need for intelligent data selection, and the diminishing error of the final transform is a result of the accumulation of evidence. The scale increase at 60 seconds corresponds with poor VO on the final turn, seen in fig. 6.

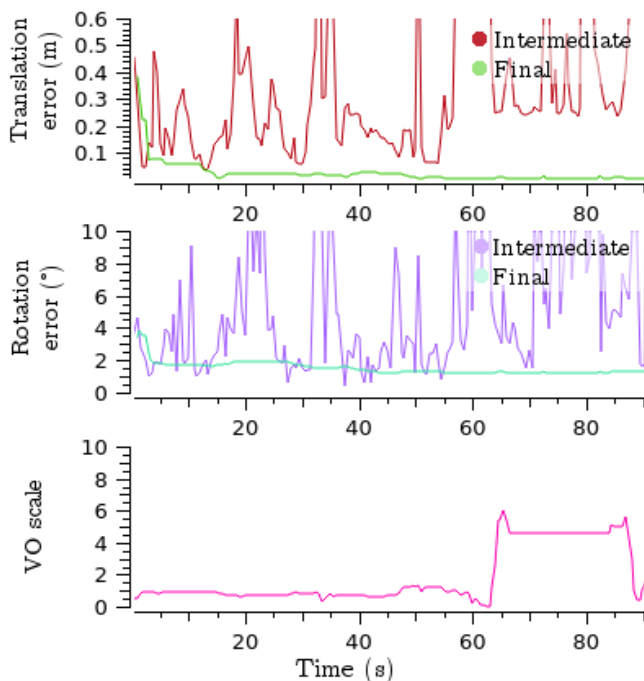


Fig. 7. Error, scale for unfiltered estimates and updating calibration.

The final calibration errors can be found in table II. ‘Schmidt (scaled)’ used a 40 second section with minimal scale drift and was manually scaled before calibrating. Unsurprisingly, incorrectly scaled monocular values cause standard hand-eye calibration to fail. The final row uses the regularizer with weight  $\alpha = 0.1$  to work in SE3 space (includes the ‘unobservable’ vertical component). It also yields slightly better rotation estimates but requires an initial guess. Here our framework overcomes both the scale issue and VO shortcomings without directly using sensor data.

#### D. Dataset Benchmarks

As the focus of our work is robustness during normal operation in real scenarios, a third experiment was performed to gauge performance on representative data: The well-established KITTI dataset [24]. ‘Odometry’ sequences contain

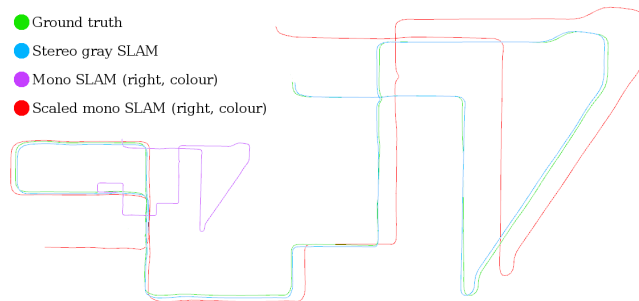


Fig. 8. SLAM trajectories from odometry sequence 8. The scaled version of monocular SLAM locally matches stereo SLAM and the ground truth.

TABLE III

CALIBRATION ERROR BETWEEN DIFFERENT SENSORS WITH RESPECT TO THE GROUND TRUTH FOR KITTI ODOMETRY SEQUENCE 8

<i>Poses in metric scale</i>		
Sensors calibrated	T (m)	R (°)
GPS/IMU → lidar	0.0311	0.23
Stereo SLAM, gray → color	0.0190	0.31
Lidar → stereo SLAM color	0.0277	0.13
Lidar → stereo SLAM color (3DoF)	0.0877	0.77
<i>Solving for scale</i>		
GPS/IMU → mono color	0.0518	0.19
Stereo SLAM gray → mono color	0.0629	0.25
Stereo SLAM gray → mono color (3DoF)	0.0633	0.10

stereo images from two cameras (grayscale and RGB), 3D Velodyne point cloud and fused GPS/IMU. The sensor extrinsics were accurately calibrated by the dataset authors on each day of recording.

We obtained the camera trajectories using stereo ORB-SLAM2 [17], and the lidar poses were generated using an open-source version of Zhang and Singh’s Lidar Odometry and Mapping (LOAM) [18]. The SLAM maps are not seen by our calibration. The vertical translation component is evaluated to be unobservable due to rotations only on the ground plane. As with the Turtlebot example however, the vertical offsets would not affect their reported motion not therefore are not required in sensor fusion. Introduction of a weak regularizer as outlined in eq. (7) allows recovery of the vertical displacement, despite being defined only by a scalar distance.

Table III shows the results of our calibration for a selection of sensor pairs (with and without scale estimation) compared with the dataset calibrations. Again, the rows noted ‘3DoF’ make use of the regularizer ( $\alpha = 0.1$ ) and include the usually unrecoverable vertical component. Scale in monocular SLAM drifts over time, but our calibration can correct for it when the motion allows. A visualisation of the scale recovery is shown in fig. 8 (purple to red). The misalignment is partially a result of comparing separate non-deterministic SLAM instances. Scale estimates are based on accumulated data and may also be slightly delayed, but it shows subsections are locally accurate.

## V. CONCLUSION

In conclusion, this work has demonstrated that robust hand-eye calibration is both convenient and can be more effective than other extrinsic retrieval techniques. We demonstrated how providing robustness to a constrained technique allows it to gather evidence over time and overcome sections of degenerate or noisy data. This is in addition to metric scale recovery which enables fusion of monocular sensors. It is dependent on the quality of the sensor odometry, so best results come from state-of-the-art algorithms in feature-rich environments.

Commercial robots and autonomous vehicles have complex sensor configurations and must be made as a single unit if they are to be factory calibrated. A tailored calibration procedure can be expected to yield a higher degree of accuracy. However, our approach displays the potential for exceptional flexibility, where existing vehicles of different brands and models could be retro-fitted with sensors without the need for a specialized procedure.

## REFERENCES

- [1] Q. Zhang and R. Pless, "Extrinsic calibration of a camera and laser range finder (improves camera calibration)," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, vol. 3, 2004, pp. 2301–2306.
- [2] F. Zhao, T. Tamaki, T. Kuriita, B. Raytchev, and K. Kaneda, "Marker based simple non-overlapping camera calibration," in *IEEE Int. Conf. on Image Processing (ICIP)*, 2016, pp. 1180–1184.
- [3] H. Chen, "A screw motion approach to uniqueness analysis of head-eye geometry," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 1991, pp. 145–151.
- [4] K. Daniilidis, "Hand-eye calibration using dual quaternions," *Int. Journal of Robotics Research (IJRR)*, vol. 18, no. 3, pp. 286–298, 1999.
- [5] Y. C. Shiu and S. Ahmad, "Calibration of Wrist-Mounted Robotic Sensors by Solving Homogeneous Transform Equations of the Form  $AX = XB$ ," *IEEE Trans. on Robotics and Automation*, vol. 5, no. 1, pp. 16–29, 1989.
- [6] M. Velas, M. Spanel, Z. Materna, and A. Herout, "Calibration of RGB Camera With Velodyne LiDAR," in *WSCG Int. Conf. on Computer Graphics, Visualization and Computer Vision*, 2014, pp. 135–144.
- [7] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 1280–1286.
- [8] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple IMUs and of individual axes," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2016, pp. 4304–4311.
- [9] Y. Ling and S. Shen, "High-precision online markerless stereo extrinsic calibration," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016, pp. 1771–1778.
- [10] E. Ataer-Cansizoglu, Y. Taguchi, S. Ramalingam, and Y. Miki, "Calibration of non-overlapping cameras using an external SLAM system," in *Int. Conf. on 3D Vision (3DV)*, 2014, pp. 509–516.
- [11] J. Maye, P. Furgale, and R. Siegwart, "Self-supervised calibration for robotic systems," in *IEEE Intelligent Vehicles Symposium (IV)*, 2013, pp. 473–480.
- [12] J. Schmidt, F. Vogt, and H. Niemann, "Calibration-Free Hand-Eye Calibration: A Structure-from-Motion Approach," in *Pattern Recognition*, vol. 3663, 2005, pp. 67–74.
- [13] T. Ruland, H. Loose, T. Pajdla, and L. Krüger, "Hand-eye autocalibration of camera positions on vehicles," in *IEEE Int. Conf. on Intelligent Transportation Systems (ITSC)*, 2010, pp. 367–372.
- [14] L. Heng, B. Li, and M. Pollefeys, "CamOdoCal: Automatic intrinsic and extrinsic calibration of a rig with multiple generic cameras and odometry," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 1793–1800.
- [15] C. Zhu, Z. Zhou, Z. Xing, Y. Dong, Y. Ma, and J. Yu, "Robust plane-based calibration of multiple non-overlapping cameras," in *Int. Conf. on 3D Vision (3DV)*, 2016, pp. 658–666.
- [16] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semi-Direct Visual Odometry for Monocular and Multi-Camera Systems," *IEEE Trans. on Robotics*, vol. 33, no. 2, pp. 249–265, 2017.
- [17] R. Mur-Artal and J. D. Tardos, "ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras," *IEEE Trans. on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [18] J. Zhang and S. Singh, "LOAM: Lidar Odometry and Mapping in Real-time," *IEEE Trans. on Robotics*, vol. 32, no. July, pp. 141–148, 2015.
- [19] M. Rohs and A. Oulasvirta, "Target acquisition with camera phones when used as magic lenses," in *Conf. on Human Factors in Computing Systems*, 2008, pp. 1409–1418.
- [20] Stereolabs, "ZED Stereo Camera," <https://www.stereolabs.com/zed/>, accessed: 15/06/2018.
- [21] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 573–580.
- [22] O. Mendez, S. Hadfield, N. Pugeault, and R. Bowden, "SeDAR - Semantic Detection and Ranging: Humans can Localise without LiDAR, can Robots?" in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2018, pp. 1–8.
- [23] M. Labbé and F. Michaud, "Online global loop closure detection for large-scale multi-session graph-based SLAM," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2014, pp. 2661–2666.
- [24] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets Robotics: The KITTI Dataset," *Int. Journal of Robotics Research (IJRR)*, pp. 1231–1237, 2013.