

# CalQNet – Detection of Calibration Quality for Life-Long Stereo Camera Setups

Jiapeng Zhong\*, Zheyu Ye\*, Andrei Cramariuc, Florian Tschopp, Jen Jen Chung  
Roland Siegwart, Cesar Cadena

**Abstract**—Many mobile robotic platforms rely on an accurate knowledge of the extrinsic calibration parameters, especially systems performing visual stereo matching. Although a number of accurate stereo camera calibration methods have been developed, which provide good initial “factory” calibrations, the determined parameters can lose their validity over time as the sensors are exposed to environmental conditions and external effects. Thus, on autonomous platforms on-board diagnostic methods for an early detection of the need to repeat calibration procedures have the potential to prevent critical failures of crucial systems, such as state estimation or obstacle detection. In this work, we present a novel data-driven method to estimate the quality of extrinsic calibration and detect discrepancies between the original calibration and the current system state for stereo camera systems. The framework consists of a novel dataset generation pipeline to train CalQNet, a deep convolutional neural network. CalQNet can estimate the extrinsic calibration quality using a new metric that approximates the degree of miscalibration in stereo setups. We show the framework’s ability to predict the divergence of a state-of-the-art stereo-visual odometry system following a degraded calibration in two real-world experiments.

## I. INTRODUCTION

In the emerging industries of autonomous robotic systems, such as micro aerial vehicles and self-driving cars, sensor calibration plays an essential role and is especially important for stereo visual sensor setups [1]. In order to estimate an accurate calibration, a sophisticated procedure [2] is usually performed beforehand, and the obtained parameters are kept for the life-cycle of the product. The most accurate and commonly used calibration methods involve the use of a known target board and require the user to perform expert calibration procedures [3], [4], [5]. However, the originally determined calibration parameters can lose their validity over time, due to environmental transients and external disturbances, such as mechanical shocks. Such changes will lead to unavoidable biases and errors in the processes that utilize the sensor data, for example state estimation or obstacle detection algorithms which can lead to critical failures of vital systems of the vehicle. Yet, as more and more autonomous platforms and applications emerge, the capability of users to perform sophisticated calibration task to update the calibration parameters should not be an expected requirement. Online target-less calibration methods also exist, but are computationally very expensive and less precise [6], [7], [8], [9]. Thus, fault detection [10] for warning the user about the need for

recalibration in a timely manner, can greatly alleviate the dangers posed by degraded calibrations [11], regardless of which method for recalibration would be used.

The calibration of stereo cameras involves determining the *i*) intrinsics and *ii*) extrinsics, which describes the relationship between the 3D world to the 2D image plane, and the transformation between the cameras, respectively. In this work, we propose a data-driven method for detecting miscalibrations in the extrinsics of stereo camera setups, to bypass the need to recalibrate in order to check the validity of the current parameters. We propose a method in which for each specific stereo setup, a network is tailor-trained to predict when a recalibration is necessary for the current setup.

Using a learning-based approach poses its own set of challenges, among which firstly is the lack of a publicly available large-scale dataset, that would feature life-long setups of stereo visual sensors where the calibration parameters degrade over time. To tackle this issue we adopt a similar semi-synthesized data generation approach as proposed in [12]. The second challenge is that the degree of miscalibration in the extrinsics of a stereo camera is difficult to quantify. A classic approach would be to calculate a degree of miscalibration by measuring the average distance of features to corresponding epipolar lines [13]. However, this procedure relies on feature matching for each stereo pair, which performs poorly at high disturbance. We address these challenges and provide the following key contributions:

- A novel metric, *weighted overall disturbance effect* (WODE), quantifying extrinsic miscalibration of stereo camera setups.
- A data generation pipeline to train a convolutional neural network (CNN) called CalQNet, that predicts from a single stereo frame whether the platform is miscalibrated.
- Experiments on calibration quality prediction performance and use cases in two real world experiments using ORB-SLAM2 [14].

## II. RELATED WORK

One of the most common ways to evaluate the monocular and stereo camera calibration is using reprojection errors of known keypoints in the scene, e.g. by using calibration targets or fiducial markers [15]. In such controlled environments and under expert handling, accurate sensor calibrations can be obtained using a variety of methods [3], [4], [5], [16].

\*Authors contributed equally to this work.

All the authors of this paper are affiliated with the Autonomous Systems Lab, ETH Zürich, {jzhong, zheyue, crandrei, ftschopp, chungj, rsiegwart, cesarc}@ethz.ch.

In addition to methods based on known target observations, target-less calibration methods have also been investigated, also due to their potential to be used online during the life-cycle of the sensor. Such calibration methods [6], [7], [8], [9] can be used during online operation and in unknown environments, since they rely on a joint optimization of tracked feature locations and camera parameters. Similarly, more robustness can be added by including external sensors as well into the problem, as proposed in [17], [18], where an inertial measurement unit is used in addition to a camera. However, all the aforementioned online calibration methods are computationally expensive, produce less accurate calibrations and potentially require additional external sensors. Due to these disadvantages, these methods are impractical for the purpose of confirming the validity of the current calibration parameters in the regular use of stereo camera systems.

Other online methods that do not directly deal with the calibration, but perform sensor fault detection, rely on correlating information from redundant sensors [19], [20], [21]. In addition to the computational requirements these methods are more costly since they require multiple sensors to integrate and work independently. Fault detection can also be indirectly done by performing sanity-checks on outputs of algorithms that rely on good calibrations [22]. In these cases it becomes ambiguous whether the reason for the fault was bad calibration, the type of motion or the current environment. Furthermore, for many systems, performing tasks with potentially failing algorithms can pose a severe safety threat. Many optimization based odometry methods such as ORB-SLAM2 [14] have been shown to be able to function with inaccurate calibrations, but at a significant degradation in performance.

In [12] the authors propose a learning-based algorithm for detecting when the intrinsic parameters of monocular cameras have diverged from the original calibration. This proves the capability and applicability of CNNs to detect mismatches between the observed environment and the undistorted image. However, the same method can not be directly applied to the sensor extrinsics problem, as the process of image undistortion is fundamentally different from stereo image rectification.

### III. METHODOLOGY

In this section, our contributions are presented in detail. The dataset generation pipeline is explained in Section III-A. In Section III-B, our novel metric for extrinsic miscalibration is defined. Finally, the neural network CalQNet and its training are detailed in Section III-C.

#### A. Dataset Generation

In a real-life scenario pre-calibrated parameters are usually stored and remain untouched, and what changes during calibration drift is the physical sensor setup. Thus, the most straightforward procedure to create a dataset to represent such a scenario, is to manually vary the extrinsics between sensors from the nominal one and take measurements with

the disturbed setup. However, in this process re-calibration would be required each time the setup is perturbed in order to understand the deviation. This is time-consuming and impractical, especially since the procedure also needs to be repeated for each individual sensor setup. This repetition is necessary because we do not propose a generic solution for detecting when any sensor suffers a bad calibration, but instead we show that it is possible to fine-tune CalQNet to detect when a specific sensor and its current calibration are no longer valid. Thus, we adopt the strategy of semi-synthesization, introduced in [12]. This includes modifying the calibration parameters by adding randomly generated disturbance and re-rectifying the stereo pairs accordingly. This is equivalent to the visual effect on a setup, where the “true calibration” is the disturbed one and the amount of miscalibration is the disturbance added.

In our stereo dataset generation pipeline, we assume each camera has been correctly calibrated, i.e., the intrinsics and distortion coefficients are known and remain valid throughout the process of stereo rectification, and only the extrinsics are considered variable. The true extrinsic transformation (translation, rotation) is denoted as  $\Theta = \{x, y, z, \alpha, \beta, \gamma\}$ . For each pair of raw camera images  $I_{left}$  and  $I_{right}$ , a rectification map  $M' = f(\Theta)$ <sup>1</sup> can be calculated to obtain the rectified images  $I'_{left}$  and  $I'_{right}$  [15]. Due to the projective transformation and tangential undistortion process, the resulting rectified images have different shapes from the original images. Nevertheless, the use of a CNN requires a unified input size. Therefore, in order to account for all these restrictions, we define a validity mask  $\mathcal{R}$ , which corresponds to the largest rectangular region in the rectified images containing only valid pixels, while retaining the aspect ratio of the original images. The cropped areas are further resized to the size of the raw images, filling the empty pixels with bilinear interpolation, to get the final sample images  $\hat{I}_{left}$  and  $\hat{I}_{right}$ . An example of the stereo rectification process and validity mask is shown in Figure 1.

We use the extrinsic calibration parameters  $\Theta^*$  provided by the dataset as ground truth, and the resulting image pairs and rectification maps as  $\hat{I}^*$  and  $\hat{M}^*$ , respectively. To obtain samples of miscalibration, a random vector  $d \in \mathbb{R}^6$  is sampled uniformly between  $[-1.5d_{thr}, 1.5d_{thr}]$ , where  $d_{thr}$  is a manually set threshold so that the application is still expected to function. We train by sampling a higher interval so that CalQNet can learn to predict miscalibration cases where the target application is no longer expected to work. In our case,  $d_{thr}$  is chosen according to the performance of ORB-SLAM2, which will be discussed in detail in Section IV. This process allows us to generate an arbitrarily large dataset for training a CNN with a relatively small number of raw images and one valid calibration.

<sup>1</sup>This mapping also depends on the camera intrinsics and distortions. However, in our notation this effect is excluded by our assumption that they remain fixed.



Fig. 1. Example of a rectified stereo pair with true calibration parameters from the KITTI dataset [23]. **Top:** Validity mask  $\mathcal{R}$  (green) corresponding to the largest rectangle containing only valid pixels with same aspect ratio as the original image. **Bottom:** Valid region after cropping and resizing to original image size as final input into the network




Fig. 2. Examples of stereo rectifications with the same amount of disturbance. For different DoFs, the disturbance affects image rectification differently. **Left:** disturbance of 0.1 m in  $x$ ,  $y$ , and  $z$ , respectively. **Right:** disturbance of 0.1 rad in  $\alpha$ ,  $\beta$ , and  $\gamma$ , respectively.

### B. Metrics for Quantifying Extrinsic Miscalibration

As described in Section III-A, stereo rectification with accurate intrinsics is a transformation parameterized by the extrinsic calibration parameters. Due to the continuous nature of these parameters, one can apply arbitrarily small disturbance to the calibration. Furthermore, the same amount of disturbance applied to different directions has a different effect on the mapping, as shown in Figure 2.

Therefore, we propose adding weight terms to quantify the different effects on the transformation. This is achieved by capturing the influence of each degree of freedom  $i$  in  $\Theta$ . During the stereo rectification, image planes of stereo cameras are rotated with corresponding projective transformation matrices  $H_{left}(d)$  and  $H_{right}(d)$  to make both images coplanar, which are dependent of the extrinsic parameters. Thus, for each parameter  $i \in SE(3)$ , we calculated a weight  $w_i$  based on how much the original plane is rotated differently from the true calibration, i.e.,

$$w_i(d_i) = \left\| \text{Rodrigues}(H_{left,i}^{-1}(d_i) \cdot \hat{H}_{left}^*) \right\|_2 + \left\| \text{Rodrigues}(H_{right,i}^{-1}(d_i) \cdot \hat{H}_{right}^*) \right\|_2, \quad (1)$$

where  $\hat{H}_{left}^*$  and  $\hat{H}_{right}^*$  are the projection matrices calculated using the true transformation vector. The inverse of the matrices calculated with disturbed parameters are multiplied

with the ones from true parameters and then converted to rotation vectors. As the calculation of such projective transformations is dependent on both the translation and rotation between the two cameras, this process quantitatively captures the relative difference between the effects on the resulting rectified image plane for the same amount of disturbance in different DoF. Finally, the overall effect (WODE),  $\delta$ , is calculated as the weighted sum of disturbances in all directions

$$\delta = \sum_{i \in SE(3)} \|w_i(d_i) \cdot d_i\|. \quad (2)$$

For training,  $\delta$  is further normalized in order to achieve  $\bar{\delta} = 1$  for the fully disturbed case, i.e. disturbing every DoF with  $d_{thr}$

$$\bar{\delta} = \frac{\delta}{\delta_{thr}} = \frac{\sum_{i \in DoF(6)} \|w_i(d_i) \cdot d_i\|}{\sum_{i \in DoF(6)} \|w_{i,thr} \cdot d_{thr}\|}, \quad (3)$$

where  $\delta_{thr}$  is the normalization factor and  $w_{i,thr}$  is the weight calculated with  $d = d_{thr}$  in Equation (1) on the corresponding DoF. We choose to normalize the distribution of WODE for training to avoid numerical instability in the network, when  $d_{thr}$  is chosen to be very small. Furthermore, it makes the output more interpretable with respect to the expected functioning of the target application.

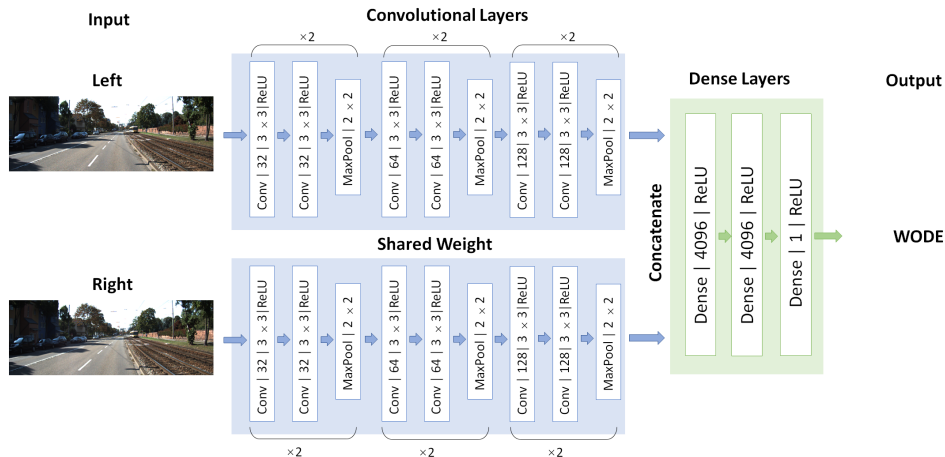


Fig. 3. Network architecture of CalQNet to predict WODE. The input is a pair of rectified stereo images which are first processed in a Siamese network and then concatenated through fully connected layers to predict the calibration quality.

### C. Network for Detecting Stereo Miscalibration

The architecture of the WODE prediction network CalQNet for stereo cameras is presented in Figure 3. The base network is adopted from [12], and transformed into a *Siamese* architecture [24], where two branches of CNNs share the same structure and weights. Images pairs are rectified and fed separately through the convolutional backbone and the outputs are concatenated together as an input to a dense network for final predictions of the calibration validity. During training, mean squared error loss is applied to the error between the networks WODE prediction and the ground truth WODE. The loss is optimized using the ADAM optimizer [25] with a reductive learning rate (halved every 20 epochs). Dropout [26] is applied to avoid over-fitting and the network parameters are initialized with Xavier’s initialization method [27]. The network consists of 25M parameters and can be inferred on a NVIDIA GeForce GTX TITAN X GPU with KITTI images ( $1392 \times 512$  pixels) in 18.5 ms.

## IV. EXPERIMENTS AND EVALUATIONS

In this section, we evaluate our trained model and metric on two publicly available datasets. The results of evaluating the trained models are discussed in Section IV-A. The evaluation of WODE as a metric is detailed in Section IV-B. Finally, in Section IV-C we show how WODE and CalQNet can be used in conjunction with an odometry system.

### A. Miscalibration Detection with a Network

To evaluate the performance of CalQNet in predicting WODE, we show results on the widely used KITTI [23] and EuRoC [28] datasets, which feature various stereo camera setups. For KITTI we only use the sequences from September 26, 2011 and perform a random split into 80% for training data and the rest for testing. A similar split is performed for all the sequences in EuRoC. We add uniformly sampled translational and rotational disturbance to the ground truth calibration as detailed in Section III-A. The thresholds for sampling are  $d_{thr, KITTI} = 0.05$  and  $d_{thr, EuRoC} = 0.025$ ,

expressed in meters or radians. The threshold values were determined empirically based on the range of values within which an odometry system might still work. The two datasets have different threshold due to the very different stereo setups with which they were recorded. The choice of  $d$  results in maximal WODE values of around  $\bar{\delta}_{max} \sim 2$ .

The performance of CalQNet is first evaluated on image pairs which were used for training, but with new random disturbances (seen scene). Secondly, new images from the test set are evaluated with random disturbances (new scene). As shown in Figures 4 and 5 CalQNet is able to capture the general trend for the sensor setup, even from scenes that the network has never encountered before, though performing not as well as on the seen scenes.

In addition, CalQNet appears to perform worse in both cases when WODE is extremely high. This performance drop could be caused by the training disturbances which are generated randomly from a uniform distribution between  $\pm 1.5d_{thr}$ , result in a non-uniformly distributed  $\bar{\delta}$  with larger values being less frequent. The second possible reason is that the cropping and resizing introduced in Section III-A result in a worse image quality with larger disturbances. Moreover, motion distortion and blur, although present in the data captured in a moving vehicle, are not explicitly addressed by the analysis, and therefore the results also account for them. Although the network structure of CalQNet is relatively simple, and potential performance could be gained by using more modern architectures, the experiments demonstrate the sensitivity of CalQNet to artifacts caused by wrong calibrations.

### B. WODE and Calibration Parameters

The most common way to evaluate stereo rectification related algorithms is by utilizing epipolar geometry. Thus, it is of interest to compare our metric with such an approach. We calculate the root mean squared distances  $e_{epi}$  of matched SIFT features [29] in a stereo image to their






Fig. 4. Network performance on KITTI dataset ( $\delta_{thr} = 0.0263$ ). **Left:** Seen scenes, **Right:** New scenes.




Fig. 5. Network performance on EuRoC dataset ( $\delta_{thr} = 0.0243$ ). **Left:** Seen scenes, **Right:** New scenes.

corresponding epipolar lines. The influences on WODE and  $e_{epi}$  from disturbances in different degrees of freedom are shown in Figures 6 and 7. Calculations are based on the KITTI 2011-09-26 calibration session and EuRoC MH\_01\_easy sequences.

Both metrics grow as the extrinsic parameters diverge from the provided calibration. This means that for both metrics larger scores represent larger disturbances. However, it can be observed in Figures 6 and 7 that  $e_{epi}$  has more discontinuities and is less smooth than the analytically calculated WODE. This is mainly due to the fact that the number of features successfully detected and matched decreases when disturbances increase and even fails to detect any features after a certain point, as shown in Figure 8.

Furthermore, disturbances in the same category (translation, rotation) have similar influence on WODE, with the exception of the direction of the baseline where the influence on WODE is minimal. This further helps WODE to generalize in capturing the calibration quality when all DoFs are disturbed, as they have a more balanced effect on the total miscalibration metric.

### C. WODE and Odometry Algorithm Performance

One way to quantify the calibration quality, from a deployment perspective in a real system, is by evaluating how much the error in calibration influences the task performance, e.g. errors in SLAM or odometry. Hence, we conducted experiments with ORB-SLAM2 [14] on both the KITTI and EuRoC datasets, in order to better evaluate the newly proposed metric WODE. Images are rectified with randomly disturbed extrinsic parameters in all 6 DoF. Trajectories extracted by the odometry algorithm are compared with the provided ground truth using the trajectory evaluation




Fig. 6. Evaluation of the influence of disturbances in different DoF based on an image from KITTI. **Left:**  $e_{epi}$ ; **Right:** WODE.




Fig. 7. Evaluation of the influence of disturbances in different DoF based on an image from EuRoC. **Left:**  $e_{epi}$ , discontinuities result from SIFT detection/matching failures; **Right:** WODE.




Fig. 8. Number of successfully matched SIFT features w.r.t. disturbances on different DoF on **Left:** an image from KITTI; **Right:** an image from EuRoC.

toolbox [17]. Figure 9 and Figure 10 show the relationship between both  $e_{epi}$  and WODE w.r.t. the root mean squared error (RMSE) of the trajectory estimate from ORB-SLAM2 on KITTI and EuRoC, respectively, as well as the success rate of ORB-SLAM2. *Successful* is defined as ORB-SLAM2 finishing odometry on a given trajectory regardless of the drift. Thus, in a *non-successful* or *diverged* run, ORB-SLAM2 loses track and can not continue its operation.

Based on Figures 4 and 5, which show the prediction accuracy of CalQNet for KITTI and EuRoC, we can calculate a standard deviation of  $\sigma_{KITTI} = 0.130$  and  $\sigma_{EuRoC} = 0.113$  for unseen scenes on KITTI and Euroc, respectively. Assuming a two  $\sigma$  confidence interval for predicting WODE, and an arbitrarily chosen operating point, for example an RMSE drift of no more than 75m for KITTI and 2m for EuRoC, respectively, we can now calculate from Figure 9 and Figure 10 how well we can predict miscalibration in these conditions. In this case the true positive rate will be the number of successful runs of ORB-SLAM2, where the drift is lower than the threshold and the WODE value is within the confidence interval around zero (meaning we predict the calibration remains unchanged). Similarly true negatives are all the runs that diverge or have a drift higher than the




Fig. 9. Performance of ORB-SLAM2 evaluated on KITTI 2011\_09\_30\_drive\_0034 (total trajectory length 920 m) w.r.t **Left:**  $e_{epi}$ ; **Right:** WODE.




Fig. 10. Performance of ORB-SLAM2 evaluated on EuRoC MH.01\_easy (total trajectory length 80 m) w.r.t **Left:**  $e_{epi}$ ; **Right:** WODE.

threshold and also have a WODE value that is further from zero than the confidence interval of prediction (meaning we predict the calibration has changed). Therefore, for KITTI we have an accuracy of 62% on predicting how likely we are to maintain our target performance given the current calibration. As an additional advantage, our method has a very high precision of 99%, meaning that we are very unlikely to trigger a recalibration unnecessarily, i.e. reporting a false positive. Similarly for EuRoC, for an arbitrary operating point of 2 m RMSE and the same confidence interval of two  $\sigma$ , we achieve an accuracy of 86% and a precision of 97%.

Not taking the failure cases into account, in Table I we calculate the Spearman’s rank correlation coefficient between ORB-SLAM2 drift and the  $e_{epi}$  and WODE metrics presented in Figure 9 and Figure 10. On KITTI, only WODE shows a significant correlation with the RMSE. On EuRoC, both  $e_{epi}$  and WODE have a significant moderate correlation with the RMSE, but the correlation with WODE is stronger. This shows that WODE is a better metric than  $e_{epi}$  to measure the expected performance of a stereo visual odometry algorithm. Overall, small WODE values correlate not only to a small trajectory error with ORB-SLAM2, but also to a higher success rate regarding the odometry diverging. The prediction accuracy and confidence could further be increased by using a sequence of predictions from different images to obtain a final decision.

## V. CONCLUSION

Our newly proposed metric *weighted overall disturbance effect* (WODE) successfully captures the degree of distortion from extrinsic miscalibration without depending on explicit

TABLE I

P-VALUE ( $p$ ) AND SPEARMAN’S RANK CORRELATION COEFFICIENT ( $r$ ) BETWEEN ORB-SLAM2 DRIFT (IGNORING FAILURE CASES) AND THE DIFFERENT METRICS ON THE KITTI AND EUROC DATASETS

	$p$		$r$	
	EuRoC	KITTI	EuRoC	KITTI
$e_{epi}$	0.00013	0.24	0.54	0.21
WODE	0.00002	0.01	0.59	0.44

traditional feature extraction and matching. We have exploited the semi-synthetic data generation introduced by [12] and successfully trained a simple neural network architecture to regress WODE from a single stereo image. In two real-world experiments covering both autonomous driving and micro aerial vehicle applications, we show that CalQNet is able to predict the divergence of ORB-SLAM2 [14] due to a bad calibration with high confidence and an accuracy between 62% and 86%. CalQNet can be deployed on-board a mobile robotic platform, where it could run at a very low frequency online to prevent potential degradation of a stereo camera system’s calibration, due to mechanical wear or other external factors. Currently, the network only performs the detection based on a single frame, which may cause inconsistency among consecutive frames. Further improvement is possible by introducing sequences and using the recurrent information.

## VI. ACKNOWLEDGMENT

This work was partially supported by Bosch Research, Siemens Mobility GmbH, Germany and the ETH Mobility Initiative under the project *PROMPT*.

## REFERENCES

- [1] C. Strecha, W. Von Hansen, L. Van Gool, P. Fua, and U. Thoennessen, “On benchmarking camera calibration and multi-view stereo for high resolution imagery,” in *2008 IEEE Conference on Computer Vision and Pattern Recognition*. Ieee, 2008, pp. 1–8.
- [2] Z. S. Roth, B. W. Mooring, and B. Ravani, “An Overview of Robot Calibration,” *IEEE Journal on Robotics and Automation*, vol. 3, no. 5, pp. 377–385, 1987.
- [3] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 11, pp. 1330–1334, 2000.
- [4] S. Abraham and W. Förstner, “Fish-eye-stereo calibration and epipolar rectification,” *ISPRS Journal of photogrammetry and remote sensing*, vol. 59, no. 5, pp. 278–288, 2005.
- [5] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, “Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4304–4311.
- [6] Y. Furukawa and J. Ponce, “Accurate camera calibration from multi-view stereo and bundle adjustment,” *International Journal of Computer Vision*, vol. 84, no. 3, pp. 257–268, 2009.
- [7] T. Dang, C. Hoffmann, and C. Stiller, “Continuous stereo self-calibration by camera parameter tracking,” *IEEE Transactions on Image Processing*, vol. 18, no. 7, pp. 1536–1550, 2009.

- [8] Z. Zhang, Q.-T. Luong, and O. Faugeras, "Motion of an uncalibrated stereo rig: self-calibration and metric reconstruction," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 1, pp. 103–113, 1996.
- [9] R. Horaud, G. Csorika, and D. Demirdjijan, "Stereo calibration from rigid motions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 12, pp. 1446–1452, 2000.
- [10] M. L. Visinsky, J. R. Cavallaro, and I. D. Walker, "Robotic fault detection and fault tolerance: A survey," *Reliability Engineering and System Safety*, vol. 46, no. 2, pp. 139–158, 1994.
- [11] K. Ni, N. Ramanathan, M. N. H. Chehade, L. Balzano, S. Nair, S. Zahedi, E. Kohler, G. Pottie, M. Hansen, and M. Srivastava, "Sensor network data fault types," *ACM Transactions on Sensor Networks*, vol. 5, no. 3, pp. 1–29, 2009.
- [12] A. Cramariuc, A. Petrov, R. Suri, M. Mittal, R. Siegwart, and C. Cadena, "Learning camera miscalibration detection," *ICRA*, 2020.
- [13] O. Faugeras and O. A. FAUGERAS, *Three-dimensional computer vision: a geometric viewpoint*. MIT press, 1993.
- [14] R. Mur-Artal and J. D. Tardos, "ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [15] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge: Cambridge University Press, 2004.
- [16] B. Atcheson, F. Heide, and W. Heidrich, "CALTag: High precision fiducial markers for camera calibration," *VMV 2010 - Vision, Modeling and Visualization*, pp. 41–48, 2010.
- [17] Z. Zhang and D. Scaramuzza, "A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry," *IEEE International Conference on Intelligent Robots and Systems*, pp. 7244–7251, 2018.
- [18] T. Schneider, M. Li, C. Cadena, J. Nieto, and R. Siegwart, "Observability-aware self-calibration of visual and inertial sensors for ego-motion estimation," *IEEE Sensors Journal*, vol. 19, no. 10, pp. 3846–3860, 2019.
- [19] J. P. Mendoza and R. Simmons, "Mobile robot fault detection based on redundant information statistics," *2012 International Conference on Intelligent Robots and Systems*, 2012. [Online]. Available: <http://www.cs.cmu.edu/~mmv/papers/12irosw-MendozaVeloSoSimmons.pdf>
- [20] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey, "Sensor fault detection and identification in a mobile robot," in *IEEE International Conference on Intelligent Robots and Systems*, vol. 3, no. October, 1998, pp. 1383–1387.
- [21] P. Sundvall and P. Jensfelt, "Fault detection for mobile robots using redundant positioning systems," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2006, no. May, pp. 3781–3786, 2006.
- [22] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. USA: John Wiley & Sons, Inc., 2002.
- [23] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset. The International Journal of Robotics Research," *The International Journal of Robotics Research*, no. October, pp. 1–6, 2013.
- [24] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. Lecun, C. Moore, E. SÄCKINGER, and R. Shah, "Signature verification using a "SIAMESE" time delay neural network," *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 07, no. 04, pp. 669–688, 8 1993.
- [25] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015, pp. 1–15.
- [26] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
- [27] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Journal of Machine Learning Research*, vol. 9, 2010, pp. 249–256.
- [28] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The EuRoC micro aerial vehicle datasets," *International Journal of Robotics Research*, vol. 35, no. 10, pp. 1157–1163, 2016.
- [29] D. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, 1999, pp. 1150–1157 vol.2.