

Modeling dynamic target deformation in camera calibration

Annika Hagemann
Robert Bosch GmbH
Computer Vision Research Lab
annika.hagemann@de.bosch.com

Moritz Knorr
Robert Bosch GmbH
Computer Vision Research Lab

Christoph Stiller
Institute of Measurement & Control Systems
Karlsruhe Institute of Technology

Abstract

Most approaches to camera calibration rely on calibration targets of well-known geometry. During data acquisition, calibration target and camera system are typically moved w.r.t. each other, to allow image coverage and perspective versatility. We show that moving the target can lead to small temporary deformations of the target, which can introduce significant errors into the calibration result. While static inaccuracies of calibration targets have been addressed in previous works, to our knowledge, none of the existing approaches can capture time-varying, dynamic deformations. To achieve high-accuracy calibrations despite moving the target, we propose a way to explicitly model dynamic target deformations in camera calibration. This is achieved by using a low-dimensional deformation model with only few parameters per image, which can be optimized jointly with target poses and intrinsics. We demonstrate the effectiveness of modeling dynamic deformations using different calibration targets and show its significance in a structure-from-motion application.

1. Introduction

Target-based camera calibration still constitutes the standard approach to estimating camera intrinsics and extrinsics in 3D computer vision. It relies on the known 3D geometry of precisely printed or measured objects, and can thereby achieve highly accurate estimates for camera characteristics (e.g. [32, 6, 21, 16]). Although target-less methods exist (e.g. [13, 17, 11]), they have not superseded the classical target-based approaches yet, as the known shape and appearance of calibration targets constitutes a relevant practical advantage.

To obtain an informative set of images for target-based camera calibration, the camera system and target are typi-

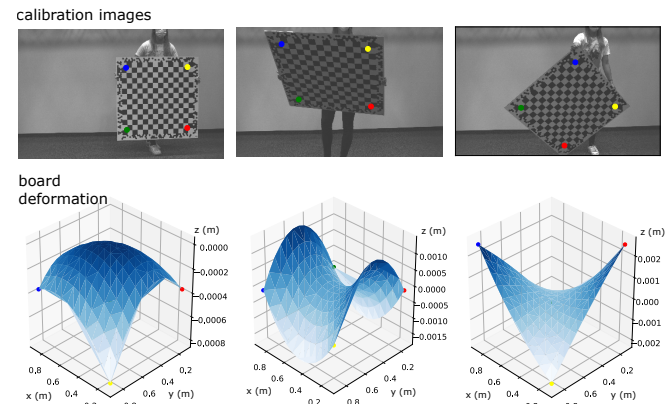


Figure 1. Carrying a calibration target leads to small deformations that can impact the calibration accuracy. The bottom line shows the deformation in each image, estimated using the proposed method. The z-axis is pointing towards the back of the board and the colored dots indicate the corresponding target corners.

cally moved w.r.t. each other [16, 19, 18]. This allows capturing images under different angles and distances, leading to a well-constrained calibration problem. Especially when the camera system is rigidly mounted (e.g. on a vehicle, robot or wall), practitioners oftentimes have no choice but to move the target. In fact, many publicly available computer vision datasets show a person carrying the calibration target [28, 4, 30], and computer vision toolboxes are specifically designed for this workflow [20, 15, 14].

Although carrying calibration targets is common practice, it has, to the best of our knowledge, never been addressed whether, and to which extent, the resulting mechanical forces impact the target geometry and consequently the calibration. While static deformations of a target, e.g. due to printing errors or plastic deformations, have been addressed in various works [10, 26, 1, 27, 24, 25, 8, 31], dynamic deformations remained unnoticed.

We show that moving or carrying a typical calibration target can lead to small elastic deformations of the material, causing the target to deviate from the assumed geometry (see Fig. 1). Although the magnitude is typically on the order of micrometers to millimeters, we show that these deformations can significantly impact the accuracy of the calibration.

To avoid the introduction of these errors into the calibration result, we propose a way to explicitly model the dynamic deformation of calibration targets within the calibration. In addition to allowing a static correction of the target geometry, we propose to include a deformation model of low parameter dimensionality, whose parameters are estimated for each individual image. We demonstrate that this augmentation does not only decrease the calibration error, but also leads to a practically relevant improvement, using the example of egomotion estimation in a Structure-from-Motion (SfM) experiment.

2. Related work

Target-based camera calibration is a well-established technique in computer vision and photogrammetry. It relies on the known 3D geometry of precisely printed or measured objects. By associating the known 3D points with the observed image points, camera characteristics are inferred.

Classical photogrammetry uses complex and precisely measured 3D targets that are imaged from a pre-defined set of perspectives [12]. To avoid this comparatively laborious process, practical computer vision and robotics typically rely on simpler targets and algorithms that do not require a precisely known set of perspectives [29, 32]. To this end, Zhang et al. [32] proposed to use planar calibration targets and developed an easy-to-use algorithm for the calibration. Besides the fact that planar targets are comparatively easy to produce, they enable a simple generation of initial values for the target poses using the direct linear transform (DLT) [7]. As a consequence, the use of planar calibration targets became one of the most widespread approaches to camera calibration [16, 14, 20, 15, 25]. However, one critical assumption of Zhang’s method, as well as for most other target-based approaches, is the accurate knowledge of 3D target points and the rigidity of the target. This assumption is violated in the case of dynamically deforming targets.

Lavest et al. [10] first noticed that small errors in the target geometry can impact the calibration accuracy. They proposed a method for 3D targets that does not only estimate camera intrinsics, but also the 3D geometry of the calibration target. A similar idea was developed for planar calibration targets [26]. Based on the observation that common printers lead to inaccuracies in the printed targets, a 2D-correction was proposed, which could be optimized during calibration [26]. Subsequent work [1] proposed to optimize not only a 2D correction, but the full scene structure (3D

points) during calibration and reported a significant reduction in residual errors.

Strobl and Hirzinger [27] further built upon the idea and proposed a calibration method in which the 3D coordinates of corners are estimated within the bundle adjustment. One of their main contributions was a parametrization that allowed the joint estimation of 3D coordinates and target poses without ambiguities. Their approach showed a significantly improved calibration accuracy when using an inaccurate calibration target (wrinkled paper) [27]. Subsequent works [8, 31] proposed additional strategies for calibrations with statically imperfect calibration targets.

While the existing approaches allow for an accurate calibration using an imperfect calibration target, the approaches have focused solely on *static deviations*. This means that only a single set of corrected 3D coordinates is estimated for the entire calibration dataset. To our knowledge, varying, or *dynamic deformations*, that can arise during the calibration process, have not been addressed to date.

3. Method

3.1. Camera projection modeling

The projection of a 3D world point onto the 2D image can be expressed by a function $p : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ that maps a 3D point $\mathbf{x} = (x, y, z)^T$ from a world coordinate system to a point $\bar{\mathbf{u}} = (\bar{u}, \bar{v})^T$ in the image coordinate system¹. This projection can be decomposed into a coordinate transformation from the world coordinate system to the camera coordinate system $\mathbf{x} \rightarrow \mathbf{x}_c = \mathbf{R}\mathbf{x} + \mathbf{t}$ and the subsequent projection from the camera coordinate system to the image coordinate system $p_C : \mathbf{x}_c \rightarrow \bar{\mathbf{u}}$:

$$\bar{\mathbf{u}} = p(\mathbf{x}, \boldsymbol{\theta}, \mathbf{\Pi}) = p_C(\mathbf{x}_c, \boldsymbol{\theta}) = p_C(\mathbf{R}\mathbf{x} + \mathbf{t}, \boldsymbol{\theta}), \quad (1)$$

where $\boldsymbol{\theta}$ are the *intrinsic* camera parameters and $\mathbf{\Pi}$ are the *extrinsic* parameters describing the rotation \mathbf{R} and translation \mathbf{t} of the camera w.r.t. the world coordinate system.

3.2. Calibration

Without loss of generality, we assume a single chessboard-style calibration target and a single camera in the following. We focus on the monocular case, as it constitutes the harder problem for deformation modeling. For stereo and multi-camera calibration, where the target is visible in more than one camera, the problem becomes more constraint. We argue that, if deformation modeling is successful for monocular calibration, it can be easily extended to the multi-camera case.

The calibration target contains a set of N_C easily detectable 3D points (e.g. chessboard corners $\mathcal{C} =$

¹We stick to the same calibration formulation and nomenclature as in [6].

$\{\text{corner}_i\}_{i=1}^{N_C}$ for checkerboard targets). The 3D coordinates of each target point i in the world (target) coordinate system are assumed to be known as $\mathbf{x}_i = (x_i, y_i, z_i)^T$. For simplicity and without loss of generality we assume the planar checkerboard surface to coincide with the world's x-y plane, such that $\mathbf{x}_i = (n_i d, m_i d, 0)^T$ where n_i and m_i are the index of corner i and d is the grid width of the checkerboard.

The calibration dataset is a set of $N_{\mathcal{F}}$ images $\mathcal{F} = \{\text{frame}_j\}_{j=1}^{N_{\mathcal{F}}}$, where the calibration target is visible from different perspectives. The image coordinates of chessboard-corners in each image are determined by a corner detection algorithm [25], giving the observations $\mathbf{u}_{ij} = (u_{ij}, v_{ij})^T$ for each corner i in each image j . As corner detection is not perfectly accurate, the observed coordinates \mathbf{u}_{ij} will deviate from the true image points $\bar{\mathbf{u}}_{ij}$ by a certain error ϵ_d , called observational noise. This error is typically assumed to be independent identically distributed (i.i.d.) $\epsilon_d \sim \mathcal{N}(\mathbf{0}, \sigma_d^2 \mathbf{I})$ where σ_d^2 is called the detector variance.

Following [32], camera intrinsics θ and target poses $\{\mathbf{\Pi}_j\}_{j=1}^{N_{\mathcal{F}}}$ are estimated by minimizing a calibration cost function, defined by the sum of squares of reprojection errors

$$\epsilon_{\text{res}}^2 = \sum_{j \in \mathcal{F}} \sum_{i \in \mathcal{C}} \|\mathbf{u}_{ij} - \mathbf{p}(\mathbf{x}_i, \theta, \mathbf{\Pi}_j)\|^2. \quad (2)$$

Having generated initial values for the target poses using the DLT, the calibration cost is optimized using a non-linear least-squares algorithm, giving parameter estimates $(\hat{\theta}, \{\hat{\mathbf{\Pi}}_j\}_{j=1}^{N_{\mathcal{F}}}) = \text{argmin}(\epsilon_{\text{res}}^2)$.

For simplicity, all formulas are written for the case where the target is fully visible in the image, although the methods are not restricted to this case. Furthermore, we present formulas for non-robust optimization. To reduce the impact of potential outliers, robustification (e.g. a cauchy kernel) can be included.

3.3. Modeling static board deformation

Static board deformations can be modeled by introducing an additional set of parameters $\{\Delta \mathbf{x}_i\}_{i \in \mathcal{C}}$ that describe the deviation of the 3D coordinate of the corner from the assumed planar grid [27]. Importantly, 7 degrees of freedom must remain fixed to prevent ambiguity with the target poses². In [27] this is achieved by setting the correction of two corners to zero, $\Delta \mathbf{x}_1 = (0, 0, 0)$, $\Delta \mathbf{x}_2 = (0, 0, 0)$ and furthermore fixing the last coordinate of a third corner $\Delta \mathbf{x}_3 = (\Delta x_3, \Delta y_3, 0)$. Here the choice of the fixed corners is free, as long as they are non-collinear [27]. The cost function with global deformation is then given by

$$\epsilon_{\text{res}}^2 = \sum_{j \in \mathcal{F}} \sum_{i \in \mathcal{C}} \|\mathbf{u}_{ij} - \mathbf{p}(\mathbf{x}_i, \theta, \mathbf{\Pi}_j, \Delta \mathbf{x}_i)\|^2. \quad (3)$$

²Six degrees of freedom corresponding to the target pose, plus one to fix the absolute scale [27].

3.4. Modeling dynamic board deformation

Modeling dynamic board deformations requires a time-dependent correction $\Delta \mathbf{x}_i(t)$ over time t . As the calibration dataset consists of a discrete set of images, time-dependence can be implemented by allowing for a different correction $\Delta \mathbf{x}_{ij}$ in every image j . For the calibration problem to still be well-conditioned, the correction cannot be modeled freely for every 3D point, as in equation (3). Instead, a low-dimensional parameterization with few parameters per image is needed.

We make use of the fact that physical deformations lead to spatially correlated deviations. If deviations $\Delta \mathbf{x}_{ij}$ of different corners i are correlated, they can be approximated with a function that only depends on the original corner coordinates \mathbf{x}_i and an additional set of parameters β_j that can differ in each image j :

$$\Delta \mathbf{x}_{ij} = \mathbf{f}(\mathbf{x}_i, \beta_j) \quad (4)$$

The parameters $\{\beta_j\}_{j=1}^{N_{\mathcal{F}}}$ can be estimated jointly with intrinsics θ and target poses $\{\mathbf{\Pi}_j\}_{j=1}^{N_{\mathcal{F}}}$ during calibration. The calibration cost function is then given by

$$\epsilon_{\text{res}}^2 = \sum_{j \in \mathcal{F}} \sum_{i \in \mathcal{C}} \|\mathbf{u}_{ij} - \mathbf{p}(\mathbf{x}_i, \theta, \mathbf{\Pi}_j, \beta_j)\|^2. \quad (5)$$

While this general formulation can be applied to any target, the specific choice of the function $\mathbf{f}(\mathbf{x}_i, \beta_j)$ may depend on the target's geometry and material. For planar checkerboard targets with comparatively rigid/stiff material (e.g. aluminum polyethylene composite), we propose a simple paraboloid to model the major deformations in the z-coordinate of the target

$$\mathbf{f}(\mathbf{x}_i, \beta_j) = (0, 0, a_j x_i^2 + b_j y_i^2 + c_j x_i y_i), \quad (6)$$

with $a_j, b_j, c_j \in \mathbb{R}$. Here x_i and y_i are the corner coordinates in the board coordinate system, the origin of which we assume to be in the center of the board. The formulation can be extended with additional terms, as long as the terms do not lead to an ambiguity with the target poses. A constant term (0th order), for instance, describes a translation in z-direction which would result in an ambiguity with the target translation. A linear term (first order), could result in an ambiguity with the target rotation.

We chose this simple model to demonstrate that no complex modeling of mechanical forces and deformations is needed to obtain a significant improvement in calibration accuracy. Clearly, more sophisticated models that preserve area and angles on the board, or that account for mechanical properties of the material could also be used. However, our experiments show that even this simple model leads to a significant reduction of calibration errors, which suggests that the model is capable of capturing the major deformations of a typical planar target.

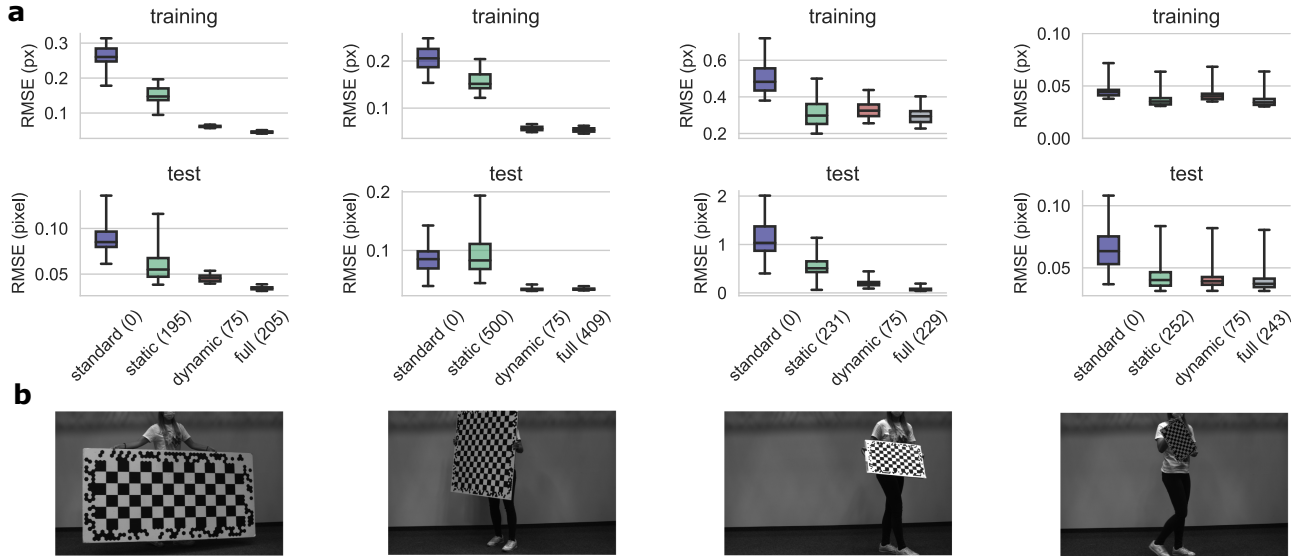


Figure 2. Comparison of the different calibration methods for different types of calibration targets. **a** Training error and test error of the different calibration methods (*standard*, *static*, *dynamic* and *full*) on datasets with different targets. The number in brackets indicates the number of deformation parameters. **b** Example images of the targets T1 (1 m × 2 m × 4 cm, styrofoam composite), T2 (1 m × 1 m × 0.6 cm, aluminum polyethylene composite), T3 (60 cm × 42 cm × 0.7 cm, paper glued to cardboard) and T4 (37.5 cm × 25 cm × 0.25 cm, aluminum polyethylene composite). We used an Allied Vision camera MG-235B with XENOPLAN 1.4/17 lens. The results and images correspond column-wise. The results for the reference dataset, where the target is lying statically on the ground, are shown in the Supplementary Fig. S1.

3.5. Combining a static in-plane with a dynamic out-of-plane deformation

In practice, a superposition of static and dynamic deformations is likely. A static deformation can be caused by an imprecise print, while dynamic deformations can arise from gravity induced sag while recording calibration images. To cover both, we propose to superimpose a static in-plane deformation in the x - y -plane with dynamic out-of-plane deformations in the target's z -direction. When including a static in-plane deformation $(\Delta x_i, \Delta y_i)$, we have to constrain four degrees of freedom³, which can be achieved by setting $(\Delta x_1, \Delta y_1) = (0, 0)$ (fixes the target position) and $(\Delta x_2, \Delta y_2) = (0, 0)$ (fixes the target scale and orientation within the $z = 0$ plane). The cost function is then given by

$$\epsilon_{\text{res}}^2 = \sum_{j \in \mathcal{F}} \sum_{i \in \mathcal{C}} \|\mathbf{u}_{ij} - \mathbf{p}(\mathbf{x}_i, \boldsymbol{\theta}, \boldsymbol{\Pi}_j, \boldsymbol{\beta}_j, \Delta x_i, \Delta y_i)\|^2. \quad (7)$$

3.6. Extension to multi-camera calibration

Modeling board deformations is even better constraint in the multi-camera case, as compared to the monocular case. Therefore, the methods that we proposed for monocular calibration are directly applicable to stereo- and multi-camera calibration. For a set of $N_{\mathcal{K}}$ cameras $\mathcal{K} = \{\text{camera}_k\}_{k=1}^{N_{\mathcal{K}}}$,

³Two degrees of freedom for the target translation in x - and y -direction, one for the rotation within the $z = 0$ plane and one to fix the absolute scale.

the cost function for modeling dynamic deformations is given by

$$\epsilon_{\text{res}}^2 = \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{F}} \sum_{i \in \mathcal{C}} \|\mathbf{u}_{ijk} - \mathbf{p}(\mathbf{x}_i, \boldsymbol{\theta}_k, \boldsymbol{\pi}_k, \boldsymbol{\Pi}_j, \boldsymbol{\beta}_j)\|^2, \quad (8)$$

where $\boldsymbol{\pi}_k$ with $k = 1, \dots, N_{\mathcal{K}}$ denotes the relative poses of all cameras w.r.t. the first camera, and $\boldsymbol{\Pi}_j$ is the pose of the calibration target w.r.t. the first camera in image j .

4. Experimental setup

To investigate the effect of target deformations on the calibration, we collected multiple calibration datasets and compared four different calibration methods:

- **standard**: the calibration target is assumed to be planar with fixed board coordinates $\mathbf{x}_i = (n_i d, m_i d, 0)^T$. This is the standard approach in camera calibration (equation (2)).
- **static**: the target is assumed to statically deviate from the assumed geometry. One global correction is estimated for the full dataset (equation (3)).
- **dynamic (ours)**: the target is assumed to deform dynamically (different in each image), where the deformation is described by a paraboloid (equation (5)).

- **full (ours)**: a static in-plane deformation is combined with a dynamic out-of-plane deformation (equation (7)).

4.1. Calibration using different targets

We used four calibration targets of different size and material (see Fig. 2b). All targets are coded to enable a clear association of corners [25]. Targets T1, T2 and T4 are high-quality, precisely manufactured targets, while target T3 is self-made by printing the target on two sheets of paper and fixing them on cardboard. This should mimic a quick and comparatively inaccurate attempt. To investigate the effect of dynamic board deformations, we collected datasets in which the targets were carried by a person and moved w.r.t. the camera (see Fig. 2b). The full calibration dataset for each target contained more than 500 images. Out of each of these datasets, we drew 50 random subsets containing 25 images each, on which we ran a calibration with each method. This allowed a statistical analysis of the different calibration methods.

As a reference calibration, we collected a dataset of 650 images with the calibration target T2 lying statically on the ground, while moving the camera. In this setup, dynamic deformations can be ruled out and only a static correction of the board may be needed. We therefore used the *static* method and all 650 images to obtain a reference calibration. To verify that the calibration did not change during recording, we collected a second reference dataset after performing all recordings (see Fig. 3).

In all calibrations, we used a pinhole model with radial distortion

$$\begin{aligned} u &= f_x \frac{x_c}{z_c} (1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + \text{ppx}, \\ v &= f_y \frac{y_c}{z_c} (1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + \text{ppy}, \end{aligned} \quad (9)$$

where $r = \sqrt{\left(\frac{x_c}{z_c}\right)^2 + \left(\frac{y_c}{z_c}\right)^2}$, f_x and f_y are the focal lengths, (ppx, ppy) is the principal point and k_1, k_2, k_3 are the distortion parameters.

4.2. Computing the training and test error

To assess the accuracy of calibration results obtained with different methods and targets, we made use of different metrics. The most straight-forward approach is the inspection of calibration residuals and the computation of the root mean squared error (RMSE). However, this error corresponds to a *training* error, as it is computed on the dataset on which the parameters were optimized. Therefore, it does not capture a possible overfitting, which could occur when introducing additional parameters. To capture overfitting, we additionally computed a *test* error on an independent test dataset. For that purpose, we used the reference calibration dataset, where the calibration target lied statically

on the ground. To compute the test error for a given set of estimated intrinsics $\hat{\theta}$, we then ran a "reduced" calibration on the reference dataset. In this reduced calibration, the intrinsics $\hat{\theta}$ were fixed, and only the target poses were optimized⁴ (see also [18]). The residual errors of this reduced calibration are a test error for the accuracy of estimated intrinsics.

4.3. Computing the mapping error

An alternative approach to assess the accuracy of the different intrinsic calibration results is the direct comparison with the ground-truth intrinsic parameters. As there is no ground-truth available in real-world experiments, we used the reference calibration as an approximation for the ground-truth.

To compare the different sets of intrinsics, we used *mapping error* in image space [3, 5, 6]. This metric propagates the differences in the individual parameters to a difference in image space. It is obtained by projecting a regular grid of image points from the image to 3D points using one set of intrinsics, and then projecting these 3D points back to the image using the other set of intrinsics. The mapping error is then defined by the RMSE of original image points compared to the projected image points.

4.4. Structure-from-Motion

To investigate the practical impact of dynamic board deformations, we use a Structure-from-Motion experiment. We collected a dataset with the camera moving around a building and returning to the starting point. To make sure that start and endpoint are exactly identical, we added the first image again to the end of the dataset.

We used AliceVision meshroom [2], an open-source 3D reconstruction software, to infer the camera trajectory and 3D structure of the building. Importantly, we only allowed for pairwise sequential feature matching, and deactivated the automatic loop detection. Thereby, calibration errors can accumulate along the trajectory.

The accuracy of the trajectory can now be assessed by the loop closure error. As in [16], we assess the loop closure error in image space: all 3D points that were reconstructed based on features extracted in the *first* image are projected to the image using the estimated pose of the *last* image. The RMSE of the original feature points compared to the projected 3D points then serves as a measure for the loop-closure error. The main advantage of this assessment is that it is unaffected by the scale ambiguity of monocular SfM.

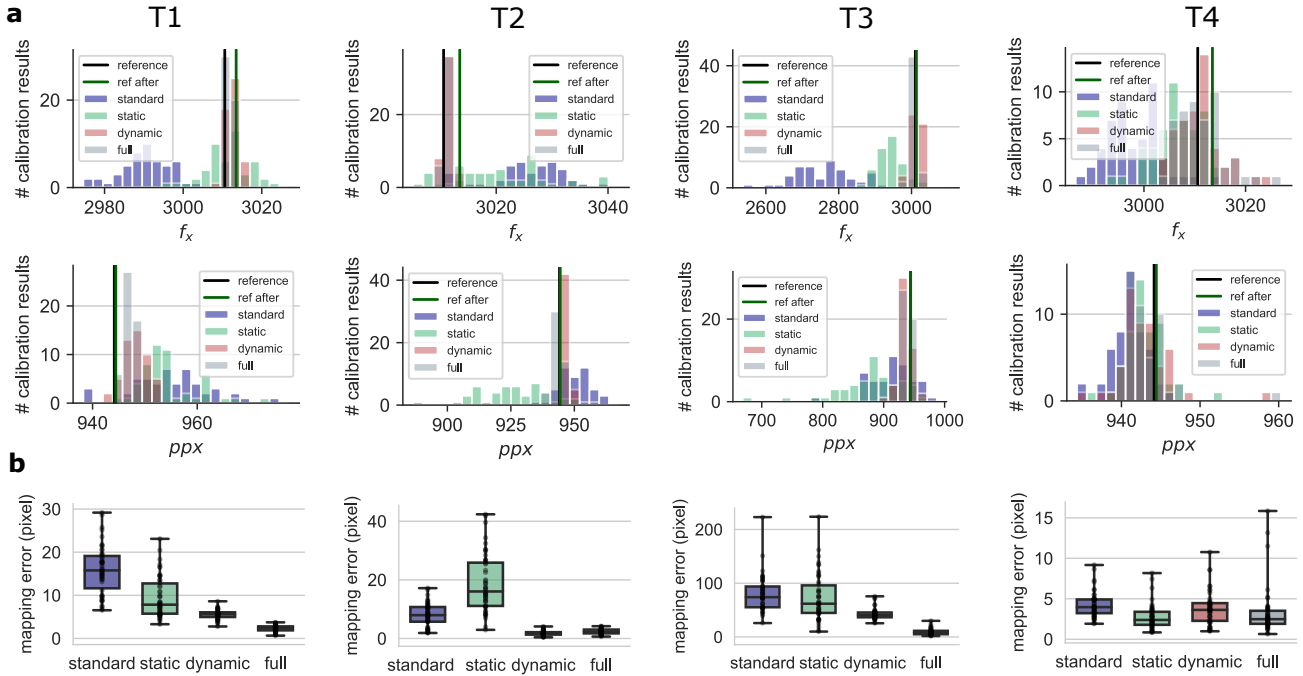


Figure 3. **a** Distribution of estimated intrinsic parameters using the different targets (T1, T2, T3, T4) and calibration methods (standard, static, flexible, full). The focal length f_x and the principal point ppx are depicted. For all other parameters see Supplementary Figs. S2 and S3. The vertical lines show the reference calibrations before and after all experiments. **b** Error in estimated intrinsic parameters w.r.t. the reference calibration, measured in terms of the mapping error in image space. High values indicate a large deviation from the reference calibration.

5. Results

5.1. Reprojection errors on training and test dataset

To analyze the impact of board deformations, we ran 50 calibrations with each of the four calibration methods *standard*, *static*, *dynamic* and *full*, for each of the calibration targets shown in Fig. 2b. The reprojection errors are shown in Fig. 2a. For all targets, the RMSE was reduced by modeling board deformation. Targets T1 and T2 show similar results with $\text{RMSE}_{\text{standard}} > \text{RMSE}_{\text{static}} > \text{RMSE}_{\text{dynamic}} > \text{RMSE}_{\text{full}}$. Interestingly, this order does not reflect the number of additional parameters: although the *static* method contains more additional parameters than the *dynamic* method, it leads to a smaller reduction in the RMSE. Target T3 shows the same trend, but higher RMSE values, as it is a self-made, less accurate target. Target T4, on the other hand, shows comparatively low RMSE values because (i) it is small and light such that less deformations are to be expected and (ii) due to its size it only covers small image regions, so that errors could partially be compensated by adjusting the pose estimate. For this target, deformation modeling only had a small effect on the RMSE. In sum-

⁴To correct the static imperfection of the target, we included a fixed, static correction $\{\Delta \mathbf{x}_i\}_{i \in C}$ that had previously been estimated by the reference calibration.

mary, modeling dynamic board deformation led to a clear reduction in reprojection errors in three out of the four targets.

The reprojection error is always a *training error*. As the number of parameters increases when modeling board deformations, a reduction in the RMSE is to be expected regardless of whether they lead to an improvement of the calibration. In particular, adding a large number of parameters comes with the risk of overfitting to the given dataset. We therefore additionally assessed the *test error*, using the reference dataset with the target lying statically on the ground.

Across most calibrations, the test error was highly correlated with the training error (see Fig. 2a). This suggests that the reduction in training errors was caused by an improved calibration rather than overfitting. For the *static* method, however, the test error for target T2 was increased. This suggests that the introduction of a static deformation led to less accurate estimates of intrinsic parameters. One possible explanation is that target T2 showed dynamic deformations which the *static* method could not model. Estimating a fixed correction when deformations are in fact dynamic resulted in wrong intrinsics. In summary, Fig. 2 shows that for the three targets with larger physical extent, modeling dynamic deformations (*dynamic* or *full*) led to a significant reduction in the test error.

5.2. Intrinsic parameters and mapping error

We compared the intrinsic parameters obtained with each of the four calibration methods *standard*, *static*, *dynamic* and *full*. Fig. 3a shows the distributions of the estimated focal length f_x and the principal point pp_x for all targets and methods. Additionally, we annotated the reference calibrations, obtained before and after the different recordings. All parameters show statistical fluctuations, i.e. the results vary depending on the image sample. However, the variance tended to be higher for the methods *standard* and *static*, as compared to *dynamic* and *full*. In fact, using dynamic deformation modeling, the parameters remained comparatively close to the reference calibration across all image samples.

To express the deviation in the individual parameters as a single number, we computed the mapping error w.r.t. the reference calibration [5, 3, 6]. Fig. 4b shows the resulting values for all targets and methods. Consistent with the test error (Fig. 2), the mapping error for targets T1, T2 and T3 reduced significantly when dynamic deformations were modeled (*dynamic* or *full*). Only for target T4, the choice of the calibration method did not lead to a significant difference in calibration results.

6. Structure-from-Motion

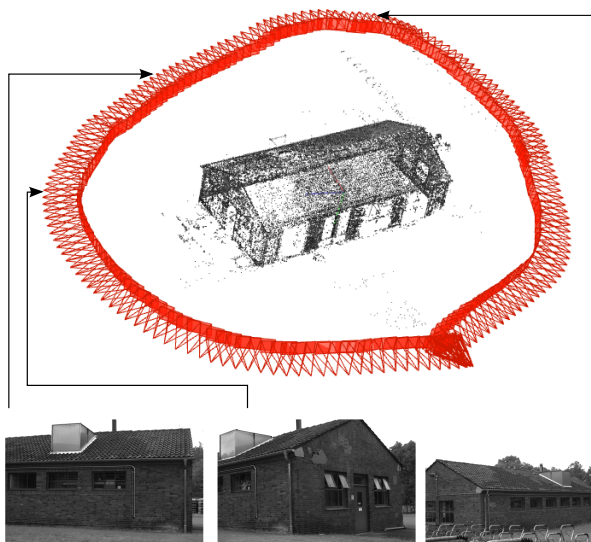


Figure 4. Structure-from-Motion experiment. The camera was moved around a building and then returned to the starting point. The bottom row shows exemplary images.

To investigate the practical impact of dynamic board deformations, we used a Structure-from-Motion experiment. Using the same camera as before, we collected a dataset where the camera is moved in a loop around a building.

accurate calibration

inaccurate calibration

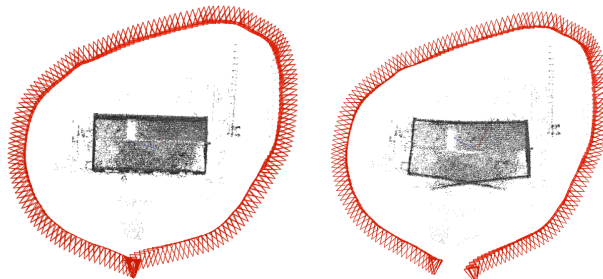


Figure 5. Top-view of two exemplary reconstruction results. Using an inaccurate calibration leads to significant errors in the 3D reconstruction, as well as in the estimated trajectory.

Fig. 4 shows the reconstructed scene and the camera trajectory, as well as some example images (for visualization we used COLMAP [22, 23] since we preferred the visualization style; quantitative analyses were performed with AliceVision meshroom [2]). The reconstruction shown in Fig. 4 is based on the reference calibration. If the assumed calibration is accurate, and sufficient features could be extracted and matched, the SfM algorithm is expected to return an accurate reconstruction. If, however, the calibration used within the SfM is highly inaccurate, less accurate reconstructions and trajectory estimates are expected.

Fig. 5 shows a top-view on the reconstruction using the reference calibration, and an inaccurate calibration (obtained from target T3 using the standard method). As loop closure detection is disabled, the errors can accumulate along the trajectory. For the inaccurate calibration, this results in a large pose gap between first and last image. Furthermore, the 3D structure of the building contains substantial errors. This shows that the accuracy of a calibration has a clear practical impact.

To assess the different calibration methods in this practical experiment, we ran the SfM experiment using the intrinsics obtained from the different calibration runs (using different targets and deformation models), and computed the loop closure error in image space. Table 1 shows the results. Across all targets, modeling dynamic board deformations led to significantly lower loop closure errors. This demonstrates that dynamic board deformations are not a minor effect, but that they can have a practically relevant impact in applications.

7. Discussion

We have shown that carrying a calibration target can lead to dynamic deformations that impact the accuracy of calibration results. As moving a calibration target in front of the camera constitutes one of the most widespread approaches to camera calibration, we consider it important for practitioners to be aware of this effect.

target	calibration method	loop closure RMSE (px)
T1	standard	16.1 ± 5.7
T1	static	5.1 ± 3.7
T1	dynamic (ours)	3.5 ± 1.8
T1	full (ours)	1.5 ± 0.7
T2	standard	9.2 ± 5.0
T2	static	11.8 ± 7.6
T2	dynamic (ours)	1.4 ± 0.6
T2	full (ours)	1.8 ± 0.6
T3	standard	248.7 ± 109.3
T3	static	49.8 ± 31.5
T3	dynamic (ours)	14.2 ± 9.9
T3	full (ours)	8.0 ± 5.6
T4	standard	9.7 ± 3.2
T4	static	5.1 ± 2.8
T4	dynamic (ours)	3.6 ± 2.4
T4	full (ours)	3.6 ± 2.4

Table 1. Reprojection error (root mean squared error, RMSE) in the Structure-from-Motion experiment. The error reflects the loop closure error of the trajectory in image space. The values are mean and standard deviation across 50 calibrations.

We proposed a way to explicitly model dynamic target deformations in the calibration. We thereby extended the existing work on static board deformation that only accounted for a single global correction across the entire dataset. Our results suggest that incorporating dynamic deformations leads to significantly more accurate and precise calibration results.

Clearly, the magnitude of the effect depends on the target’s size, shape, weight and material. Our results suggest that deformations are more relevant for large calibration targets, while the effect was smaller for the small, light target. Furthermore, self-made targets clearly show more deformations than professional targets. However, even for high-quality targets, the incorporation of target deformations led to more accurate calibration results.

In general, modeling static or dynamic board deformations introduces additional parameters into the calibration. Depending on the amount of data, this can come with a risk of overfitting to the data, as well as a risk of ambiguities (e.g. wrongfully attributing lens distortions to a deformed target). However, we did not observe such an effect in our experiments. Instead, the calibration results showed a lower variance than the standard calibration results. To ensure that deformation modeling did not lead to overfitting on a given dataset, one can either evaluate the calibration result on a separate test dataset, or evaluate the model uncertainty, e.g. using the expected mapping error [6]. Furthermore, model selection criteria, such as the geometric akaike infor-

mation criterion [9] and the bias ratio [6] can be applied to decide whether deformation modeling is needed for a specific dataset.

Importantly, our results suggest that a simple paraboloid is sufficient to achieve a significant improvement of calibration results. Due to this simplicity, the proposed model can be easily incorporated into existing calibration frameworks. In future, one could also account for temporal correlations in deformations, to further reduce the number of additional parameters. Furthermore, more elaborate deformation models could be tried, which explicitly account for the mechanical properties of the target.

References

- [1] Andrea Albarelli, Emanuele Rodolà, and Andrea Torsello. Robust camera calibration using inaccurate targets. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31:376–383, 2009.
- [2] AliceVision. Photogrammetric computer vision framework. meshroom. <https://github.com/alicevision/meshroom>. Retrieved on July 29th, 2021.
- [3] Johannes Beck and Christoph Stiller. Generalized B-spline Camera Model. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 2137–2142. IEEE, 2018.
- [4] Michael Burri, Janosch Nikolic, Pascal Gohl, Thomas Schneider, Joern Rehder, Sammy Omari, Markus W Achtelik, and Roland Siegwart. The euroc micro aerial vehicle datasets. *The International Journal of Robotics Research*, 2016.
- [5] Annika Hagemann, Moritz Knorr, Holger Janssen, and Christoph Stiller. Bias detection and prediction of mapping errors in camera calibration. In *German Conference on Pattern Recognition*. Springer, 2020. in press.
- [6] Annika Hagemann, Moritz Knorr, Holger Janssen, and Christoph Stiller. Inferring bias and uncertainty in camera calibration. *arXiv preprint arXiv:2107.13484*, 2021.
- [7] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. 2004. OCLC: 171123855.
- [8] Lei Huang, Qican Zhang, and Anand Asundi. Flexible camera calibration using not-measured imperfect target. *Applied optics*, 52(25):6278–6286, 2013.
- [9] Kenichi Kanatani. Geometric information criterion for model selection. *International Journal of Computer Vision*, 26(3):171–189, 1998.
- [10] Jean-Marc Lavest, Marc Viala, and Michel Dhome. Do we really need an accurate calibration pattern to achieve a reliable camera calibration? In *European Conference on Computer Vision*, pages 158–174. Springer, 1998.
- [11] Manuel Lopez, Roger Mari, Pau Gargallo, Yubin Kuang, Javier Gonzalez-Jimenez, and Gloria Haro. Deep single image camera calibration with radial distortion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11817–11825, 2019.
- [12] Thomas Luhmann, Stuart Robson, Stephen Kyle, and Jan Boehm. *Close-range photogrammetry and 3D imaging*. Walter de Gruyter, 2013.

- [13] Q-T Luong and Olivier D Faugeras. Self-calibration of a moving camera from point correspondences and fundamental matrices. *International Journal of computer vision*, 22(3):261–289, 1997.
- [14] MathWorks. Matlab Computer Vision Toolbox. <https://de.mathworks.com/help/vision/camera-calibration-and-3-d-vision.html>. Retrieved on January 8th, 2021.
- [15] OpenCV. OpenCV Tutorial Camera Calibrations. https://docs.opencv.org/master/d7/d21/tutorial_interactive_calibration.html. Retrieved on January 8th, 2021.
- [16] Songyou Peng and Peter Sturm. Calibration Wizard: A Guidance System for Camera Calibration Based on Modelling Geometric and Corner Uncertainty. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1497–1505, 2019.
- [17] Marc Pollefeys, Reinhard Koch, and Luc Van Gool. Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters. *International Journal of Computer Vision*, 32(1):7–25, 1999.
- [18] Andrew Richardson, Johannes Strom, and Edwin Olson. AprilCal: Assisted and repeatable camera calibration. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1814–1821. IEEE, 2013.
- [19] Pavel Rojtberg and Arjan Kuijper. Efficient pose selection for interactive camera calibration. In *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 31–36. IEEE, 2018.
- [20] ROS. ROS Tutorial MonocularCalibration. http://wiki.ros.org/camera_calibration/Tutorials/MonocularCalibration. Retrieved on January 8th, 2021.
- [21] Thomas Schoeps, Viktor Larsson, Marc Pollefeys, and Torsten Sattler. Why Having 10,000 Parameters in Your Camera Model is Better Than Twelve. *arXiv preprint arXiv:1912.02908*, 2019.
- [22] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016.
- [24] T. Strauss. *Kalibrierung von Multi-Kamera-Systemen*. KIT Scientific Publishing, 2015.
- [25] Tobias Strauss, Julius Ziegler, and Johannes Beck. Calibrating multiple cameras with non-overlapping views using coded checkerboard targets. In *17th international IEEE conference on intelligent transportation systems (ITSC)*, pages 2623–2628. IEEE, 2014.
- [26] Klaus H Strobl and Gerd Hirzinger. More accurate camera and hand-eye calibrations with unknown grid pattern dimensions. In *2008 IEEE International Conference on Robotics and Automation*, pages 1398–1405. IEEE, 2008.
- [27] Klaus H Strobl and Gerd Hirzinger. More accurate pinhole camera calibration with imperfect planar target. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1068–1075. IEEE, 2011.
- [28] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012.
- [29] Roger Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal on Robotics and Automation*, 3(4):323–344, 1987.
- [30] V. Usenko, N. Demmel, and D. Cremers. The double sphere camera model. In *Proc. of the Int. Conference on 3D Vision (3DV)*, September 2018.
- [31] Suzhi Xiao, Wei Tao, and Hui Zhao. Flexible and accurate camera calibration using imperfect planar target. In *International Conference on Optical and Photonics Engineering (icOPEN 2016)*, volume 10250, page 102501C. International Society for Optics and Photonics, 2017.
- [32] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334, 2000.