https://doi.org/10.2352/ISSN.2470-1173.2021.18.3DIA-101

3D Multi-Camera Calibration with a Fractal Encoded Multi-Shape Target

Dominique E. Meyer [1]***, Pranav Verma*** [1] ***and Falko Kuester*** [1]
[1]Dronelab, University of California San Diego; La Jolla, CA
{dom, pverma, fkuester}@ucsd.edu

Abstract

Robust multi-camera calibration is a fundamental task for all multi-view camera systems, leveraging discreet camera model fitting from sparse target observations. Stereo systems, photogrammetry and light-field arrays have all demonstrated the need for geometrically consistent calibrations to achieve higher-levels of sub-pixel localization accuracy for improved depth estimation. This work presents a calibration target that leverages multi-directional features to achieve improved dense calibrations of camera systems. We begin by presenting a 2D target that uses an encoded feature set, each with 12 bits of uniqueness for flexible patterning and easy identification. These features combine orthogonal sets of straight and circular binary edges, along with Gaussian peaks. Our proposed feature extraction algorithm uses steerable filters for edge localization, and an ellipsoidal peak fitting for the circle center estimation. Feature uniqueness is used for associativity across views, which is combined into a 3D pose graph for nonlinear optimization. Existing camera models are leveraged for intrinsic and extrinsic estimates, demonstrating a reduction in mean re-projection error of for stereo calibration from 0.2 pixels to 0.01 pixels when using a traditional checkerboard and the proposed target respectively.

Introduction

Geometric camera calibration allows a camera model to be fit to map the geometric transformation of light through the optical elements of a camera system in addition to relative extrinsics in the case of multi-camera systems. Such calibrations enable the mapping of observed world features to image space, often needed in the case of multi-view and geometric computer vision applications. Despite the emergence of auto-calibration methods which allow systems to solve for calibration parameters at the time of deployment, explicit calibration is required for many industrial camera systems with limited computational power, yet real-time estimations needs.

Camera calibration is divided into two major steps, including intrinsic and extrinsic camera parameter estimation. Intrinsic calibrations are widely used across applications requiring physical lens and non-perfect sensor-lens placements to be corrected. Lens distortions, de-centering and focal lengths are the primary variables of such intrinsic calibrations. For extrinsic parameters of a system, it is merely a relative pose $R|t$. Any vision problem involving the mapping or inverse mapping of 3D world features into the image plane is subject to these physical effects that need to be accounted for.

The proposed calibration target addresses a set of limitations associated with current calibration methods. Firstly, current targets lack dense features that can be uniquely identified across multiple-views, limiting the image-space feature distribution and hence suffering from larger calibration errors around the edges and corners of cameras. Furthermore, a predominance of strategies leverage local corner features that are subject to localization errors larger than if larger features are used. This work therefore proposes a dense and unique calibration target alongside a feature extraction algorithm, and formulates the optimization problem to obtain camera parameters for improved calibration.

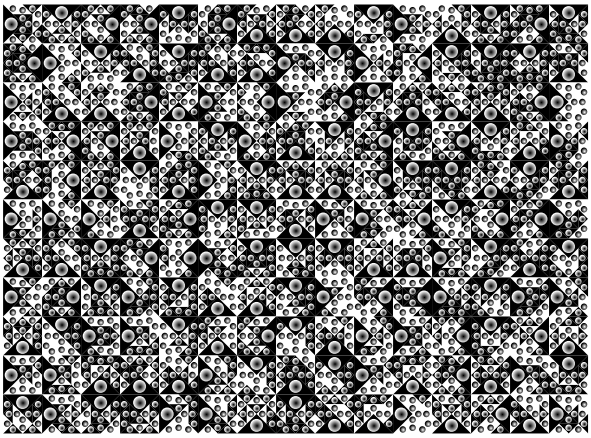

Figure 1: Proposed encoded calibration target.

Prior Work

Prior work covers a broad spectrum of approaches to fit a camera model from a calibration target, ranging from 2D planar targets to 3D rigs, and features of varying type. The most common approaches adopted across open-source vision libraries such as OpenCV [1] and Matlab use a calibration approach based on a singular planar checkerboard and a traditional corner feature detection algorithm [10]. These detection methods achieve sub-pixel accuracy by solving for the local corner features within a window that best represents convergence of tangential gradient lines [3], or the minimum eigenvalue of the local gradient co-variance matrix [7]. Such checkerboard patterns can however only be detected when all inner corners are visible, making acquisition of the pattern features difficult across the full Field of View (FoV) of a single camera, and as such even more challenging for multi-camera systems. Figure 2 demonstrates the coverage issues which arises from limited visibility. In the case of multi-view systems, it is often desirable to enforce calibration constraints dependent on the system baseline direction. Checkerboard edges only contain

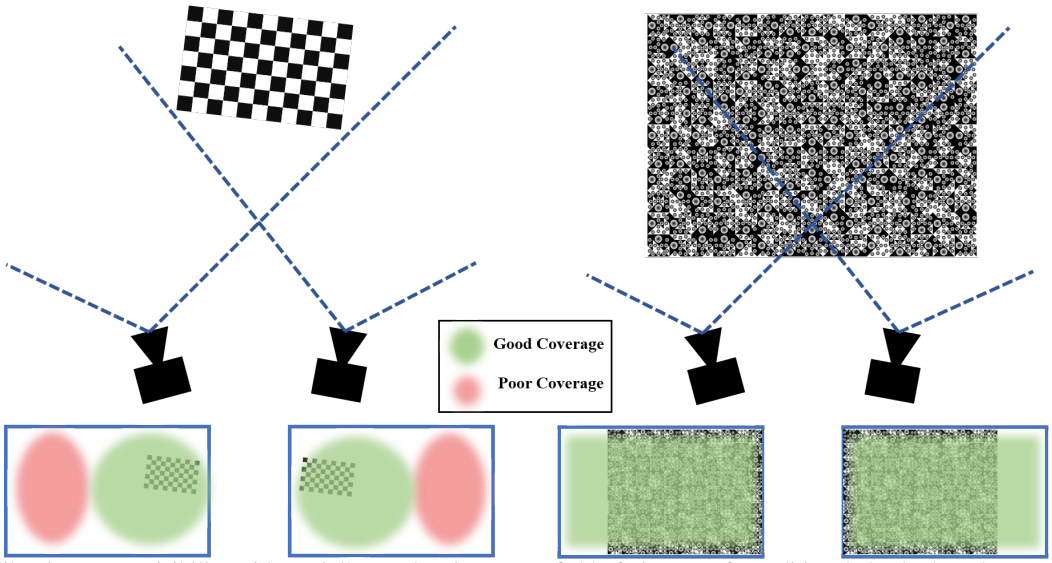

Figure 2: Calibration target visibility with partially overlapping stereo field of views. Left- traditional checkerboard target requiring full observations for detection. Right- proposed target with unique local feature patches allowing for partial visibility.

horizontal and vertical edges, which in the case of perpendicular arrays (trinocular systems in an L-configuration, or lightfield systems with both vertical and horizontal distribution of cameras) are insufficient to enforce geometric constraints. Augmenting calibration targets with diagonal edges and circular features, enables such constraints to be further increased.

Circular features have proven to improve localization accuracy due to their fitting from a larger neighborhood [6]. While the improved feature localization can improve geometric calibrations, overall calibration quality is still subject to having full FoV coverage of features to ensure the calibration accounts for corner aberrations in a set of observations. In the case of circular patterns without unique encoding, these still require all features to be visible, again introducing the visibility limitation shown in Figure 2.

Alternatively, a fractal target encoding local square patches of decreasing size has been used [8], with the work emphasizing the importance of sampling density and uniqueness to enable feature samples across the full camera FoV to achieve improved calibration. Despite the coverage improvement from this target, the corner feature localization is still limited to the accuracy of corner detection methods.

Calibration Target

This paper proposes a calibration target (Figure 1) designed to address limitations of current calibration patterns and methods. Specifically it features uniquely encoded patches that allow partial visibility and enable complete sampling across camera FoVs. The target leverages diverse straight edges and circular features to improve the localization accuracy of feature fitting and enables multi-view constraint enforcement.

Each target is divided into a set of arbitrary number of user-defined patches with a spatial sizing chosen to match camera resolution and FoV for reliabale detection. Target dimensions should be selected such that the size of the smallest circle features observed by a given camera system results in reliable feature detection. The square patches as shown in Figure 3 consist of large triangle-circle pairs that anchor the patch and enforce rotational in-variance. The rest of the patch is composited with smaller triangle sets with circles of opposite gradient. The encoding for each patch is achieved using this set of twelve smaller triangles and circles where the gradient direction defines the binary value. This approach supports up to $2^{12} = 4096$ unique patches, containing thirteen circular features each, resulting in up to 53,248 unique circle features.

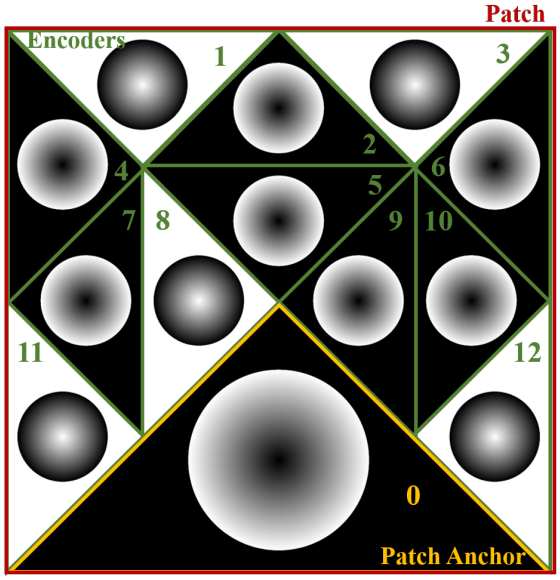

Figure 3: Annotated sections of the calibration target patches.

Additional features can be added in a fractal pattern by subdividing each small triangle by 4, enabling the spacing of 6 additional features per small triangle as can be seen in Figure 4. Since complete patches need to be detected to identify the patch identifier, and with the desired objective to have maximum image coordinate coverage of features, it is preferable to have more

smaller patches over increased fractal levels of circular features.

The square patches contain diagonals introduced through the triangular features. These angled edges enable 8 different local gradient orientations (black to white transitions). Additionally the circle edges have 2π angular range forming close contour ellipses when observed. The circular patches contain a continuous spherical gradients that assume a Gaussian distribution. This gradient defines the local orientation of each circle that is used to derive the directions of the peaks and may in the future serve as areas way for a secondary peak fitting method [9].

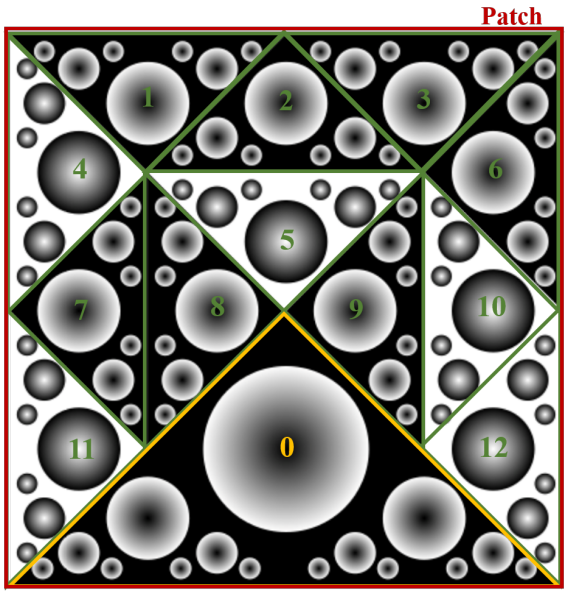

Figure 4: Example of a unique patch for a 3-level fractal target.

To ensure that the features are uniformly distributed with both white and black peaks and edges of all directions, the patch identifiers are randomly sampled without replacement. Furthermore each patch is randomly rotated *pi* radians to ensure a random distribution of large circles across the target. At the time of detection, this orientation is recovered from the large anchor circle and set of locations of smaller circles.

Detection and Calibration

The circular features in the proposed calibration target and patch encoding scheme require a detection pipeline which leverages high sub-pixel localization accuracy to achieve reliable calibrations. We begin by an edge detection, followed by feature linkage to separate straight lines from ellipses. Once all features have been fit, the patches are recognized from the feature locations and identifiers to provide a dictionary of feature-coordinate pairs used by the model fitting within the calibration. The full pipeline is depicted in Figure 5 and will be discussed within this section.

Feature Detection

The importance of feature localization accuracy on calibration quality imposes the requirement for good edge fitting, in turn used for the ellipse fitting and centroid finding. Sub-pixel edge algorithms can use first and second order image gradients, with the first requiring additional optimization to find local ridges and the latter explicitly being defined as the roots of second order gradients. We use the steerable filtering approach introduced by [5] with user-tunable Gaussian kernel sizes to locally steer the second order derivatives to the locally dominant orientation. This provides an image gradient and an orientation image which is then used for solving for the real roots of the bi-linear surface yielding continuous edges. Edges are linked to provide edge point groups which belong to continuous feature contours as shown in Figure 6. With the help of the feature fitting described in the next step, these are then separated into ellipses and lines, rejecting any remaining outliers that do not meet fitting requirements and a minimum contour length.

The target uses circles as its primary features, which when subject to a projective transformation assume elliptical geometries [6]. While locally non-linear lens distortions may affect that assumption, it holds true as long as projected circles are small with respect to the overall lens distortions. The ellipse fitting is done using a direct least squares approach proposed in [2]. All contributing points within a linked contour are used for the fitting, and a secondary pass is used to merge disjoint ellipse segments for an optimized fitting. All fit ellipses have a major and minor axis diameter, eccentricity, with a center point and a fitting quality metric (mean point to ellipse distance from edge points). These metrics are used to validate only reliable ellipse fits, and the di-

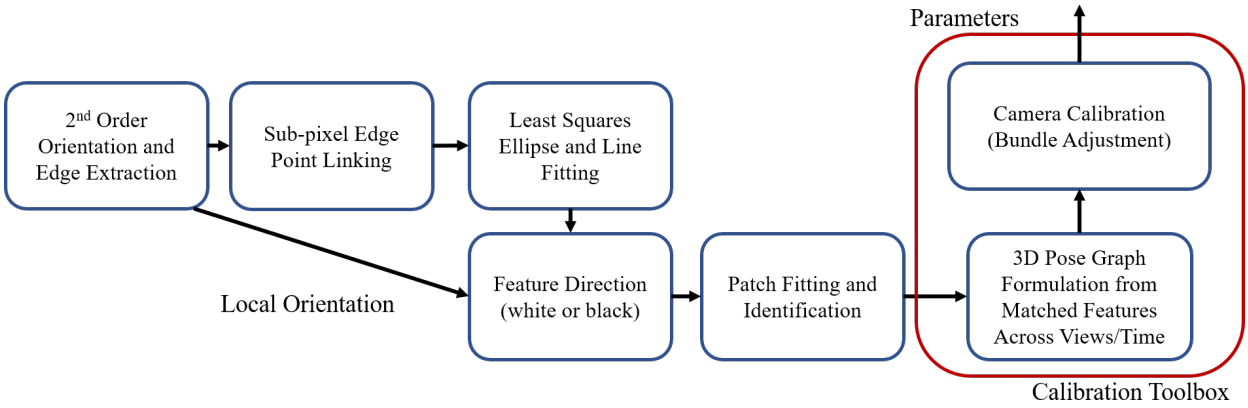

Figure 5: Target detection pipeline.

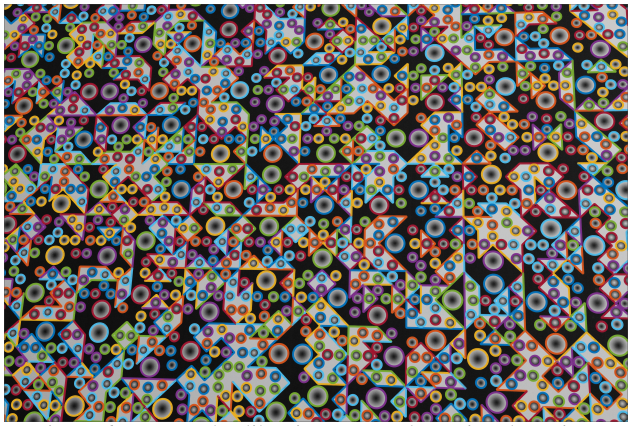

Figure 6: Proposed calibration target edge point clustering.

ameter is used to separate large from small ellipses with an Otsu threshold biased by the ratio of large to small circles.

The edge gradient directions of the ellipses are used to extract whether the ellipse is white or black on an oppositely shaded triangle. This direction is identified from the ellipse curvature normal direction and the locally dominant orientation from the image filtering. If they are in agreement, both vectors point towards the ellipse center implying a white peak, and if they have opposite directions, they imply a black peak.

Target Extraction

To implicitly match features across image views for the calibration, a direct identifier is associated to each feature, stemming from the patch and identifier they lie in. These patches are extracted from the local distribution of ellipse centers, which are referenced from the large ellipses. Each patch is extracted only if all ellipses within the patch are confirmed valid. The features are linked to the 3D world points from the target plane by associating the feature image coordinates and world coordinates via their identifiers. A complete extraction of patches is demonstrated in Figure 7a.

Model Fitting & Calibration

We leverage the located and matched correspondences to formulate the optimization problem as a bundle adjustment problem: given the 3D points and the observed 2D locations in each view, simultaneously optimize for the camera parameters (intrinsics and distortion parameters) and the pose at each view such that the reprojection error is minimized. This bundle adjustment problem can be solved using a variety of non-linear optimization strategies, with the Levenberg-Marquardt optimization being commonly chosen to combine the convergence benefits of both Gauss-Newton and gradient descent [11]. For this problem, we utilize OpenCV's implementation of the Levenberg-Marquardt based calibration to optimize over intrinsics and extrinsics.

Based on empirical evidence, it has been observed that it is beneficial to perform the calibration process in two stages: In the first stage, for each camera separately, we estimate and optimize

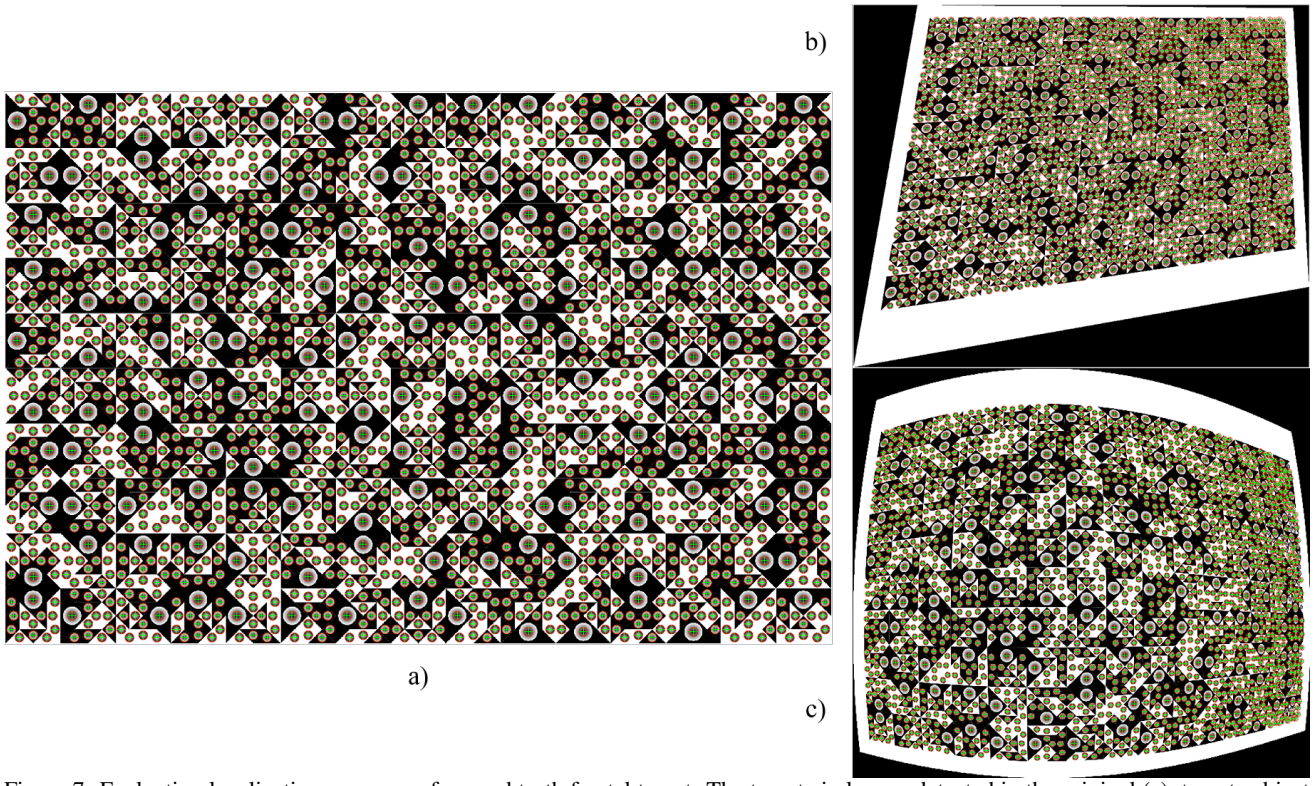

Figure 7: Evaluating localization accuracy of ground truth fractal target. The target circles are detected in the original (a), target subject to projective distortion (b) and target subject to radial distortion (c). Detected features are shown as green crosses in red circles, overlayed on the target.

the intrinsic parameters and pose at each view, minimizing the reprojection error computed by projecting the world points to image space using the intrinsics and pose estimates. Then, in the second stage, we use these intrinsics estimates to optimize for the relative extrinsics between the left and right camera, by estimating and optimizing the pose of each view in both cameras so that the relative pose between left and right camera remains constant, and the reprojection error using these pose estimates and the fixed intrinsics is minimized. It has been observed that separating the intrinsic and extrinsic calibration optimizations ensures convergence to a better optima.

Results

The results are divided into two sets of experiments with separate objectives: validating the localization accuracy of elliptical features and evaluating the camera calibration performance for monocular and stereo camera calibrations. We summarize these as:

- Sub-pixel localization accuracy of spatial images features given synthetically rendered images with varying degrees of geometric distortions, to demonstrate target and algorithmic reliability at extracting correct feature locations.
- Demonstrable and real-world results for the intrinsic and pose estimation of a stereo camera moved around a checkerboard and fractal target. Evaluating the reprojection error residuals across the images using a) only a monocular approach for intrinsic calibration and b) a stereo calibration to solve the relative pose using the monocular estimated intrinsics.

Localization of Features in Synthetic Targets

For the first experiment, we generated a synthetic fractal target and its ground truth sub-pixel level feature locations. To compare the feature localization accuracy, we subjected the target, and the ground truth feature locations, to different kinds of distortions: projective homographies (to yield targets similar to Figure 7b), radial and tangential distortions (to yield targets similar to Figure 7c), and then ran our feature extractor on the warped targets.

For projective distortions, we create a projective transformation matrix of the form:

$$T = \begin{bmatrix} cos(\theta) & sin(\theta) & 0 \\ -sin(\theta) & cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

and varied θ between -15 and 15 degrees.

For radial distortions, we considered the 2-variable model for radial distortion:

$$x_{distorted} = x(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) \quad (2)$$

$$y_{distorted} = y(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) \quad (3)$$

$$r^2 = x^2 + y^2 \quad (4)$$

Where x, y are the undistorted normalized image coordinates. For our experiment, we set $k_1 = k, k_2 = k^2, k_3 = 0$, where k varied logarithmically between $1e-5$, and 0.33.

For each of the tests, we compute the ground truth feature locations in the warped image by applying the same transformations to the ground truth feature locations, and computing the mean L2 distance between the detected feature locations and the ground truth locations.

We present both qualitative and quantitative comparisons to demonstrate the performance of the feature extractor. Figures 7b and 7c show that our model can accurately detect most of the features seen in the ground truth target image (Figure 7a, even under large distortions. Quantitative results are shown in Figures 8 and 9, where we compute the mean feature localization error vs distortion magnitude.

We observe that our feature extractor can localize features to accuracies greatly smaller than 0.1 pixels even under extreme image distortions while the traditional corner extractor for the checkerboard only reached accuracies in minimal distortions conditions of 0.15 pixels with a rapidly degrading localization when either projective and radial distortions were increased.

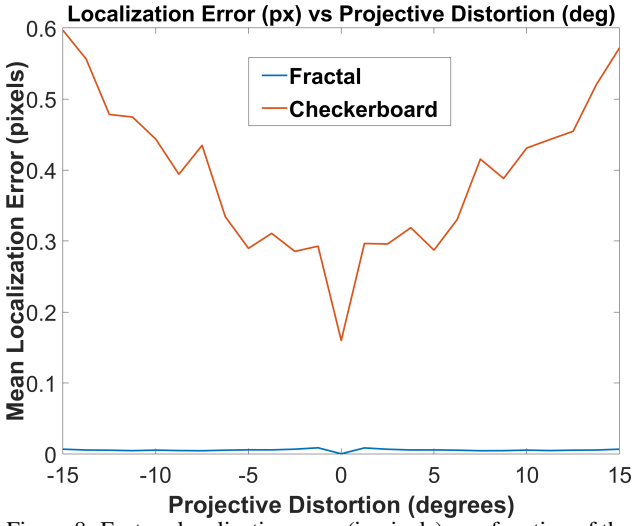

Figure 8: Feature localization error (in pixels) as a function of the image distortion, on our target vs min-eigen feature detector on checkerboard.

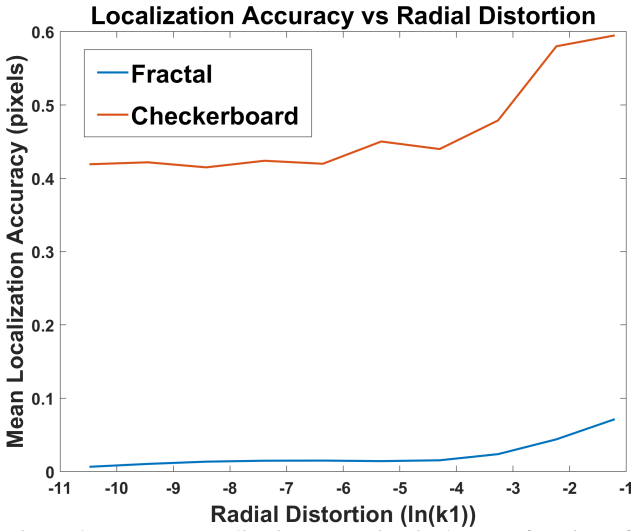

Figure 9: Feature Localization error (in pixels) as a function of the radial distortion, on our target vs min-eigen feature detector on checkerboard.

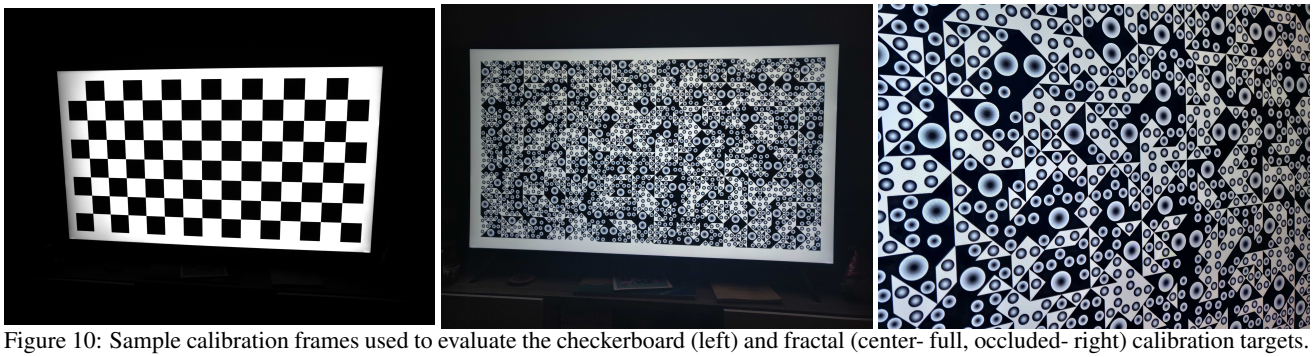

Figure 10: Sample calibration frames used to evaluate the checkerboard (left) and fractal (center- full, occluded- right) calibration targets.

Real-world Camera Calibration

The second set of experiments evaluated real-world performance of the proposed calibration routine. For data capture, we used dual TRI120S Lucid machine vision cameras (Sony IMX304 sensor) with fixed 12mm Computar lenses (V1226-MPZ), shot at an aperture of f/5.6. The acquisitions were separated into 3 datasets, each with 50 frames- one with the checkerboard (see Figure 10a), one with the fractal calibration target fully seen in all views (see Figure 10b), and one with only a partial/occluded views of the fractal target (see Figure 10c). The targets were displayed using a 65 inch 4k TV with an aspect ratio of 16:9.

For each of the target frames, we extracted the features with their identifiers and matched them across different views to establish feature point correspondences. The checkerboard ones leveraged the OpenCV pipeline with the *findChessboardCorners* and *cornerSubPix* functions for feature extraction (these use the algorithm proposed by [7]), while the fractal target ones used the proposed feature extractor from the previous section. Then, using OpenCV's *calibrateCamera* function, we estimated the intrinsic and extrinsic camera parameters for the cameras.

All evaluations were done with the True Pixel Error (TPE) metric, which compares the reprojected world points with the 2D image features given an estimated camera pose. This metric is used across multi-view bundle adjustment problems such as photogrammetry, and is considered a good measure of consistency for pose estimation and feature localization [4].

Monocular and Stereo Calibration

The monocular calibration considered the left and right cameras of the stereo captures independently, with the results presented in Table 1.

We pose the stereo camera calibration as a non-linear optimization of the camera intrinsics, extrinsics, and the relative poses between the left and right stereo cameras given the 3D world coordinates of the feature points and the 2D image coordinates of the features seen in each frame by both cameras. For an efficient and accurate implementation, we first leverage the intrinsic parameters for both the left and right cameras using OpenCV's *calibrateCamera* previously calculated in a monocular seeting, and utilize these estimates to optimize for the extrinsics using a Levenberg-Marquardt based non-linear optimization, using OpenCV's *stereoCalibrate* function.

Both monocular and stereo calibrations are evaluated using the checkerboard and fractal targets. The mean reprojection error is calculated for each 3D world point back projected into image coordinates from observed views. The statistical distributions of the mean reprojection errors is summarized in Table 1 and the errors are plotted over the image coordinates in Figure 11.

We validate the coverage of the respective targets from the various calibration views in Figure 11. The plots include the image coverage of points across the full dataset of 50 frames for both the checkerboard and fractal data. The shear difference in number of points indicates that the fractal target better samples the space, providing improved data for the calibration. The color magnitude shows the reprojection error for the estimated points. Despite the significantly improved coverage with the proposed fractal target, there remain small areas in the corners that lack points. This is likely due to the patches in those areas not being fully detected, eliminating them from being used for the calibration.

Discussion

The previously introduced re-projection error is an important metric to validate geometric camera calibration. During the calibration optimization, pose and intrinsic parameters are adjusted to minimize this error resulting in the camera parameters and the respective error for them. Across the photogrammetry community, this error is widely used to evaluate to global consistency and quality of the reconstruction [4]- a lower error indicates improved reconstruction and camera localization. In the case of calibration, we can similarly use it to evaluate the calibration per-

Target		Avg points/frame	Monocular						Stereo		
			Left			Right					
			RMSE	μ	σ	RMSE	μ	σ	RMSE	μ	σ
Checkerboard		98	0.288	0.255	0.135	0.314	0.279	0.144	0.302	0.267	0.140
Fractal	All points	1017.4	0.242	0.191	0.149	0.255	0.200	0.159	0.249	0.196	0.154
	Subset	98	**0.059**	**0.049**	**0.032**	**0.059**	**0.049**	0.033	**0.059**	**0.049**	**0.033**
Fractal Occluded	All points	500.0	0.232	0.180	0.146	0.251	0.194	0.159	0.242	0.187	0.153
	Subset	98	0.065	0.053	0.036	0.060	0.050	**0.032**	0.062	0.051	0.034

Table 1: Results table for the mean re-projection error (pixels) of checkerboard and fractal target calibrations in both monocular and stereo configurations

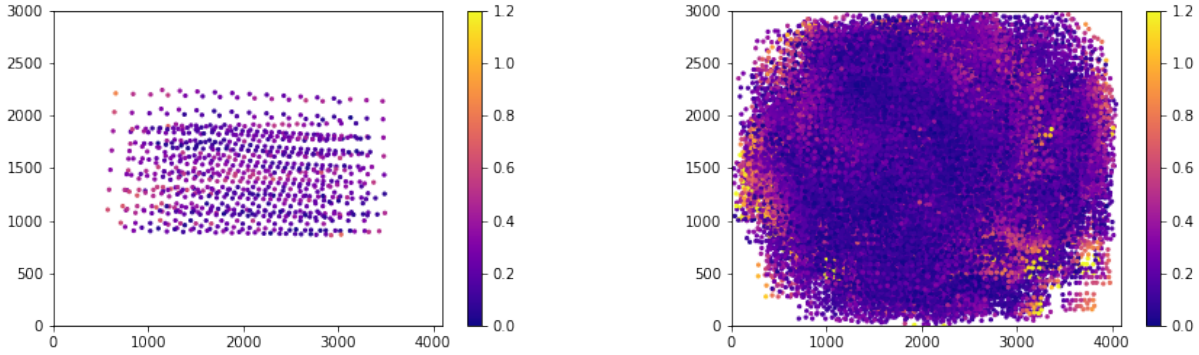

Figure 11: Image coverage of points detected from Checkerboard (left) and Fractal target (right) over accumulated 50 frames. Color magnitude represents the reprojection error (pixels).

formance. Since this error is often on the order of sub-pixel values, it is necessary to verify both that the localized image features are accurately localized before evaluating the calibration. The first set of experiments that tested the effect of projective (Figure 7b) and radial distortions (Figure 7c) on the localization accuracy highlighted that traditional corner localization techniques only approach accuracies of around 0.2 pixels at best. The ellipse based detections from this work are able to reach localization accuracies nearly an order of magnitude better, localizing centers to around 0.01 pixels. Furthermore, the results suggest that increasing the magnitude of geometric distortions affects both localization performances of corner points and elliptical features, with the latter remaining more robust over greater distortions. It is likely that extreme distortions which would yield worse localizations are rejected through the feature quality metric of elliptical fitting, a particular advantage over corner features that are more difficult to validate. These results conclude that the localization performance of elliptical features is greatly superior over corner features.

The second component of calibration evaluation focuses on the calibration quality which were evaluated separately on monocular and stereo calibrations in a real-world setting. From the results, we note that:

- Our feature extractor allows us to set up better feature point correspondences across different frames, which gives us a significantly lower mean reprojection error as compared to the standard checkerboard based calibrator (0.06 pixels vs 0.3 pixels).
- The uniquely identifiable patches allow the target to be partially visible, providing easier sampling of feature points across corners and edges of the FoV. Figure 11 shows the feature distribution for the checkerboard and fractal target respectively.
- Despite the improved and validated feature localization, there remains larger reprojection errors across corner and edge segments of the image coordinates, indicative of an outstanding systematic error in the camera model fitting. Additional radial parameters and local surface deformations may be a future improvement to address the calibration performance.

The overall reduction in this error is indicative of improved overall calibration since there is more local agreement to the solution. While the solution is limited to numerical optimization methods such as the Levenberg-Marquardt, there may be cases where the converged minimum is a solution that is not indicative to ideal geometry.

Conclusion

This paper introduced a fractal calibration target that uniquely encodes circular features for robust multi-view camera systems. Elliptical feature detection is used to improve the localization accuracy over traditional corner features. Projective and radial distortion experiments confirmed the ability of the detector to localize to 0.01 pixels, improving from around 0.2 pixel localization accuracy in traditional corner detection methods. Monocular and stereo calibrations were tested using a proposed fractal target, improving the reprojection errors compared to a checkerboard calibration. Despite the improvements in calibration performance, systematic errors indicate the limitation of existing camera models for geometric calibration.

Acknowledgments

This publication is based on work supported by the US Army Corps of Engineers under research Cooperative Agreement W912HZ-17-2-0024, NIST Award #70NANB17H211, as well as NSF award #CNS-1338192, MRI: Development of Advanced Visualization Instrumentation for the Collaborative Exploration of Big Data, under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, under the 2020 Qualcomm Innovation Fellowship, and by the LLNL-LDRD Program under Project No. 20-SI-005. We thank all collaborators at the Dronelab, Qualcomm Institute, the Contextual Robotics Institute, Autonomous Vehicle Laboratory, Engineers for Exploration, UC San Diego, as well as all other contributors to ideas, suggestions and comments. Opinions, findings, and conclusions from this study are those of the authors and do not necessarily reflect the opinions of the research sponsors.

References

[1] Gary Bradski and Adrian Kaehler. Opencv. *Dr. Dobb's journal of software tools*, 3, 2000.

[2] Andrew W Fitzgibbon, Maurizio Pilu, and Robert B Fisher. Direct least squares fitting of ellipses. In *Proceedings of 13th Inter-*

national Conference on Pattern Recognition, volume 1, pages 253–257. IEEE, 1996.

[3] Wolfgang Förstner and Eberhard Gülch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proceedings of ISPRS intercommission conference on fast processing of photogrammetric data*, pages 281–305. Interlaken, 1987.

[4] Wolfgang Förstner and Bernhard P Wrobel. *Photogrammetric computer vision*. Springer, 2016.

[5] William T Freeman and Edward H Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence*, 13(9):891–906, 1991.

[6] Janne Heikkila. Geometric camera calibration using circular control points. *IEEE Transactions on pattern analysis and machine intelligence*, 22(10):1066–1077, 2000.

[7] Jianbo Shi and Carlo Tomasi. Good features to track. In *Proceedings of IEEE conference on computer vision and pattern recognition*, pages 593–600. IEEE, 1994.

[8] Hendrik Siedelmann, Maximilian Diebold, Marcel Gutsche, Hamza Aziz-Ahmad, and Bernd Jähne. A fractal calibration pattern for improved camera calibration. In *Forum Bildverarbeitung 2016*, volume 84, page 1. KIT Scientific Publishing, 2016.

[9] Xiaowei Wan, Gangyi Wang, Xinguo Wei, Jian Li, and Guangjun Zhang. Star centroiding based on fast gaussian fitting for star sensors. *Sensors*, 18(9):2836, 2018.

[10] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334, 2000.

[11] Richard Hartley Andrew Zisserman. Multiple view geometry in computer vision. 2004.

Author Biography

Dominique E. Meyer is a PhD candidate at the University of California San Diego, where he previously received his B.S. in Physics in 2017 and his M.S. in Computer Science Engineering in 2019. He is working with Professor Kuester and Professor Christensen on the development of real-time 3D imaging systems for robotics and 3D mapping.

Pranav Verma is currently pursuing his Masters of Science in Computer Science and Engineering at UC San Diego, focusing on Computer Vision. He's currently working as a Graduate Student Researcher in the UC San Diego DroneLab, with Prof. Falko Kuester, where he's working on visual SLAM systems, light fields, computer vision simulation pipelines, and semi-supervised object localization. Prior to this, he completed his Bachelors in Electrical Engineering from IIT Delhi in 2019, where he worked on using deep learning to solve domain-specific problems in computer vision, NLP and power systems.

Falko Kuester is a Professor in the Jacobs School of Engineering at UC San Diego and the director of the UC San Diego DroneLab, advancing research in robotics, remote imaging, large-scale visual analytics and virtual reality. He received his Ph.D. in Computer Science from UC Davis in 2001 and a M.S.E. in Mechanical Engineering as well as Computer Science and Engineering from the University of Michigan, Ann Arbor, in 1994 and 1995 respectively.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

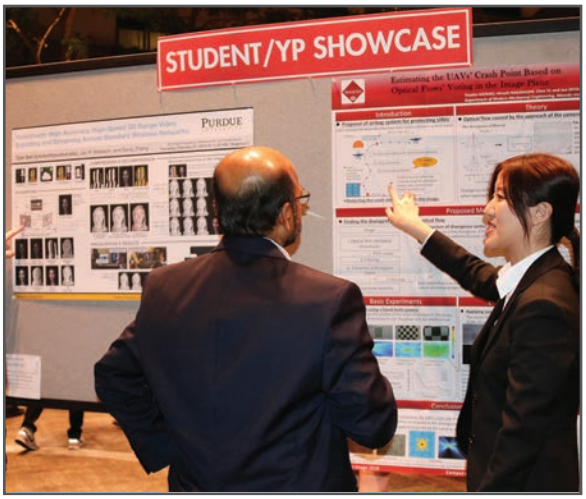

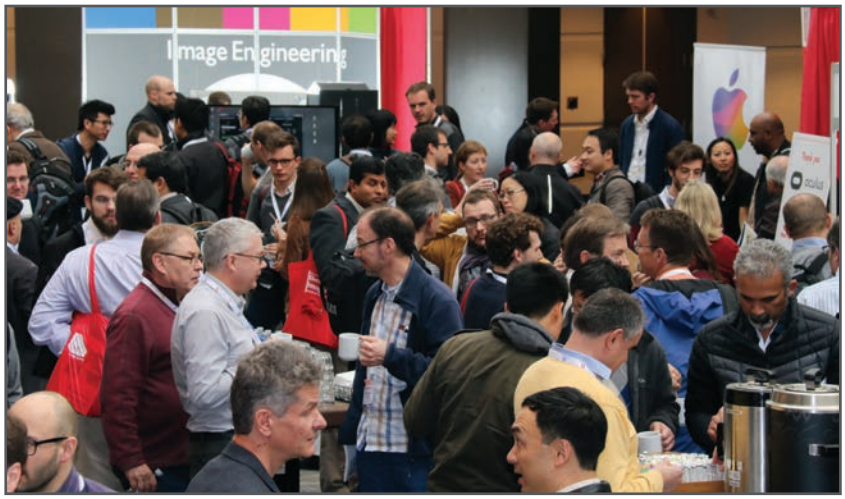

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org