GFaI Future Tech Days
Conference Proceedings 2025

AI4EA Workshop (Berlin Workshop on Artificial Intelligence for Engineering Applications)

3D-iSA (3D in Science & Applications)

Gesellschaft zur Förderung
angewandter Informatik e.V.

TABLE OF CONTENTS

GFaI Conference Proceedings 2025

Accurate Projector-Camera-Calibration from Structured Light Dot Patterns via Subpixel Plane Reconstruction

David Reese, Evgeny Degtyarev, Rico Nestler

Zentrum für Bild- und Signalverarbeitung e. V.
Werner-von-Siemens-Straße 12, 98693 Ilmenau
david.reese@zbs-ilmenau.de, evgeny.degtyarev@zbs-ilmenau.de, rico.nestler@zbs-ilmenau.de
www.zbs-ilmenau.de

Abstract: We propose a novel method for geometric calibration of active 3D camera-projector arrangements and projectors with fixed structured light dot patterns. While traditional calibration techniques often rely on projected pattern sequences such as Gray-code patterns, our approach enables accurate estimation of camera intrinsics and extrinsics using a single, static projection while maintaining high accuracy over the full field of view. This makes it highly suitable for stereo-based 3D imaging systems employing fixed laser-based dot projectors.

Keywords: Projector-Camera-Calibration; Single-Shot-3D; Structured Light; Dot Pattern Projection; 3D-based Augmented-Reality

1 Scientific Background and Motivation

Compact 3D imaging systems based on active depth sensing are increasingly used in applications such as robotics, industrial automation, and augmented reality. Two major depth-sensing technologies dominate the field: Time-of-Flight (ToF) sensors (such as the Microsoft Azure Kinect DK) and structured light systems that infer 3D geometry by projecting known patterns.

ToF cameras offer practical advantages, including ease of integration and frame-based depth acquisition. However, they are also subject to inherent limitations, most notably motion blur and multipath interference, where light reflected along indirect paths causes depth distortion, especially around concave shapes and sharp edges [1].

Structured light systems, by contrast, can provide more accurate geometry reconstruction in such challenging regions. This is particularly true for single-shot systems using laser and diffractive optical elements (DOE) to project dense, static dot patterns – an approach originally introduced in the Kinect v1. Our research is focused on developing and refining such fixed-pattern structured light sensors for use in complex and dynamic scenes, where reliable geometry is critical. Fig. 1 illustrates the advantages of our method by comparing a ToF sensor capture to a single-shot structured light capture on a combined radiometric-geometric calibration target [2].

A crucial requirement for these systems is accurate projector-camera calibration, i.e. estimating both the internal parameters of the camera and projector and their relative pose. Unlike modulated projection systems (e.g. those using Gray-code sequences), fixed-pattern projectors do not allow the direct estimation of dense pixel correspondences. Therefore, many calibration approaches rely on indirect reconstruction of the projector view by projecting onto a planar calibration target and using chessboard corner detections to estimate a local homography between the target plane and the projector image.

While chessboard corners can often be detected with good subpixel precision, the homographic mapping of projector dots based on sparse corner references introduces systematic errors – especially in regions far from the corners or when the texture of the board introduces local interference. In this work, we present a new calibration method specifically designed for fixed-pattern structured light systems. Instead of relying on sparse corner features, we directly use the projected dot centers as reference points, leveraging their known positions in the projector image and their subpixel-accurate localization in the camera image. This results in a dense and accurate calibration dataset without requiring projection sequences or external texture features. The method improves accuracy of the calibration process for fixed structured light systems intended for single-shot 3D-acquisition.

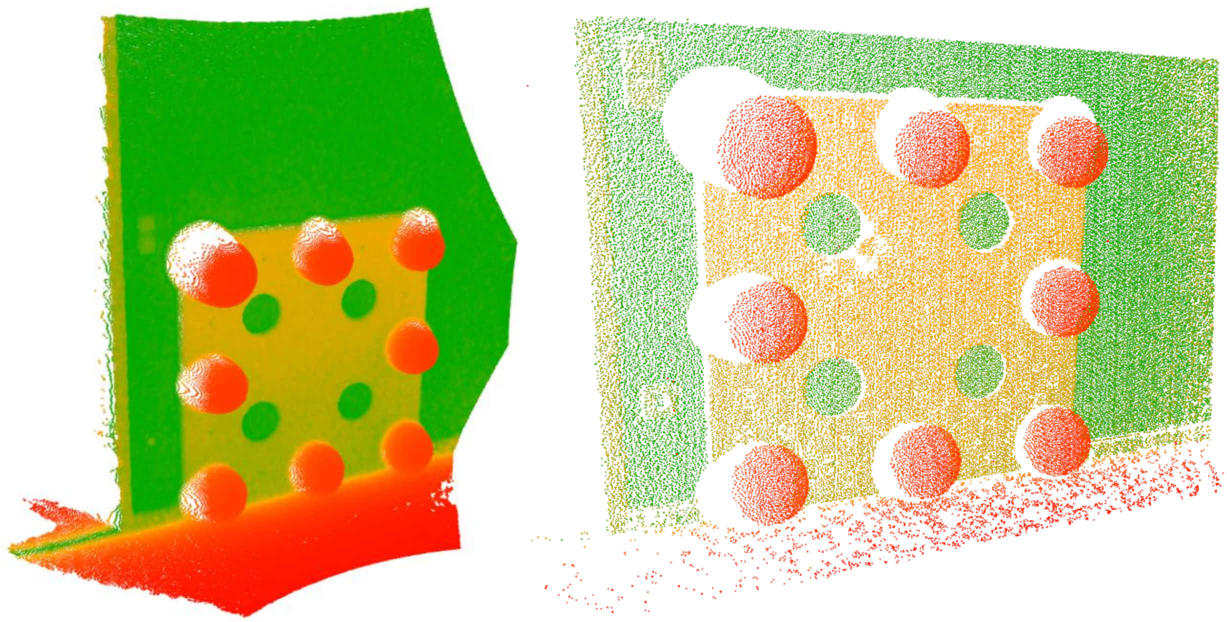

Fig. 1. 3D-pointclouds of a calibration target with geometric features captured by ToF-based Kinect Azure DK (left) and our in-house single-shot structured light sensor (right). The color-coding shows signed distances (green to red) of points to a plane reconstructed from the wall.

2 Methods Used

Standard Projector-Camera Calibration using temporal pattern Decoding (by Moreno and Taubin [3])

Projector-camera calibration is a critical step in active stereo 3D reconstruction systems. It involves determining the intrinsic parameters of both camera and projector, as well as their relative pose (extrinsics). The projector is modeled as an inverse camera and is calibrated through correspondences between camera image points and projector coordinates. A well-established method for this process is presented by Moreno and Taubin [3] and is implemented in our in-house calibration toolkit 3D-EasyCalib™ (see Fig. 2 and 3).

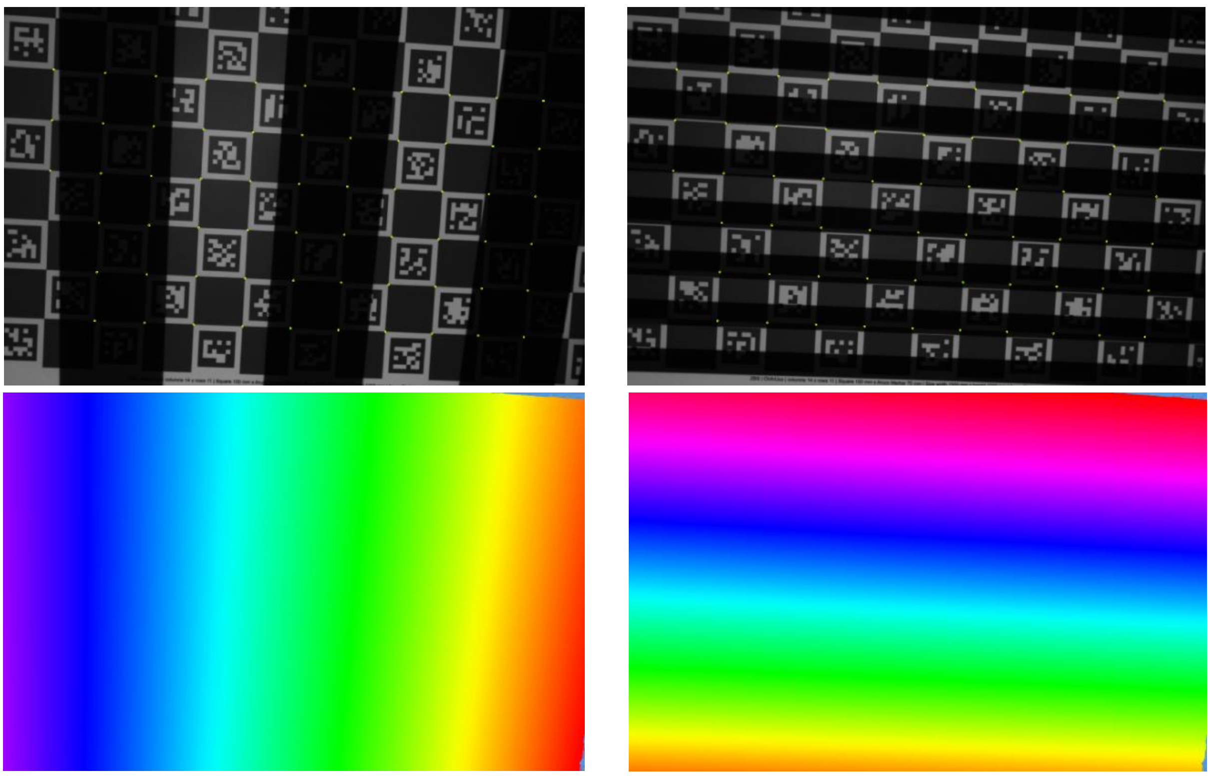

Fig. 2. Top: Sample-captures of a projected Gray-code sequence for a single target pose with corner-detections (homogeneous illumination used) of the ChArUco-Target. Bottom: Result of decoding the Gray-code sequence for this target pose. The Color-coded coordinates (left: x, right: y) map camera coordinates to projector coordinates.

In this approach, a sequence of Gray-code patterns is projected onto a planar calibration target (here a ChArUco-Target), and the sequence is captured by the camera. The known spatial coding of the sequence allows per-pixel decoding of projector coordinates corresponding to the observed pattern. These decoded projector coordinates are then mapped to the corresponding known 3D object points (of the target) via the camera image (see Fig. 3 left). To associate projector coordinates with 3D object points, the method estimates a local homography around each chessboard corner. This allows the projector to be calibrated by effectively projecting the known 3D structure into the projector's image space. Optimization over all correspondences across multiple target views e.g. by using the standard calibration approach by Zhang [4] yields the final intrinsic and extrinsic calibration parameters, modeled using a pinhole projection model with radial and optionally tangential lens distortion.

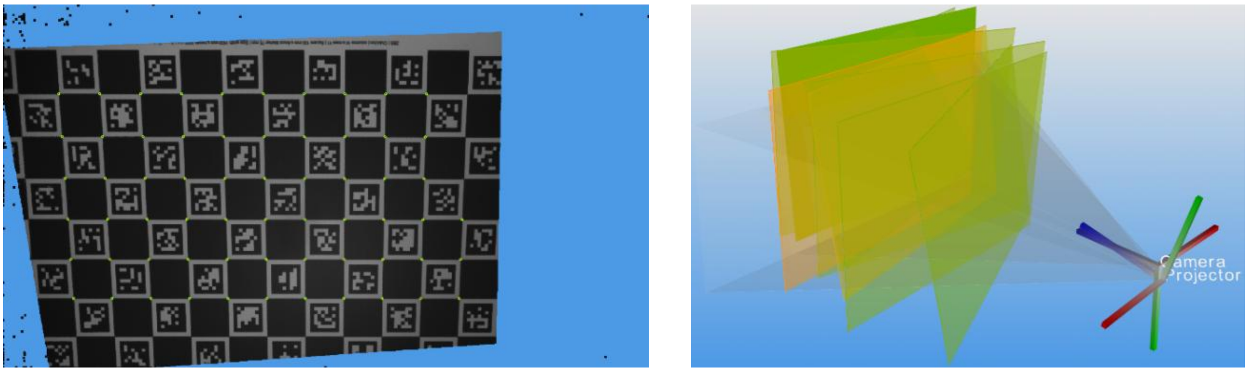

Fig. 3. Left: Reconstructed view of the target in the projector image space. The target pose is the same as in Fig. 2. Right: Results of the projector-camera calibration and view of the target poses used for calibration. The colors of the target poses visualize the reprojection error (low-to-high ≙ green-to-red, with pure yellow = 1.0 px). The overall root mean square (rms) reprojection of our calibration with eight target poses error was 0.625 px.

<u>DOE dot projector calibration using spatial pattern decoding (by Vehar et al [5])</u>

A variant of this method was adapted in our previous work [5] using a fixed laser + DOE dot projector (see Fig. 4). Since the static dot pattern cannot be temporally modulated like Gray-code sequences, a spatial decoding or matching approach is employed instead: a calibration board is imaged both with and without the projected dot pattern, and the projector coordinates of individual dots are determined by spatially matching the observed pattern to a reference layout. These coordinates are then assigned to known target points using local geometric approximations through homographies computed in the vicinity of chessboard corners.

Compared to standard Gray-code-based projector-camera calibration methods, this spatial decoding approach introduces additional constraints: the camera intrinsics must be calibrated in advance, and all camera images must be undistorted prior to decoding. In the remainder of this paper, we assume that all camera images are preprocessed and undistorted accordingly. Despite these constraints, the core principle remains the same: estimating projector coordinates corresponding to known 3D object points to enable accurate geometric calibration of the projector-camera system.

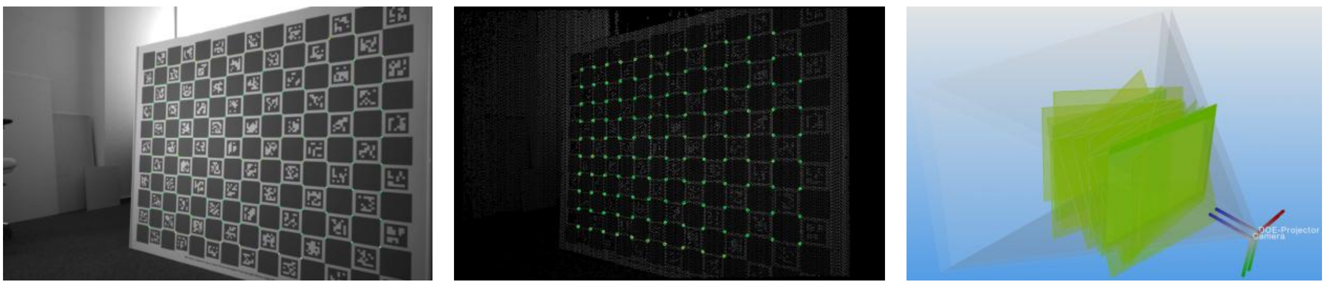

Fig. 4. Left: Camera view of target with detected corners. Center: Reconstructed projector view of the same target pose with transferred corners. Right: Results of the projector-camera calibration and view of the target poses used for calibration. The colors of the target poses visualize the reprojection error (green to red, with pure yellow = 1.0 px). The same colorization of target poses was used as in Fig. 3 and Fig. 5. The overall rms reprojection error of our calibration with 16 target poses was 0.321 px.

Proposed method for improving accuracy of DOE dot projector calibration

Our proposed method builds upon the approach by Vehar et al. [5] and retains much of the existing pipeline structure. Again, the goal is to recover accurate correspondences between camera and projector views to perform joint intrinsic and extrinsic calibration using the standard approach by Zhang [4]. We capture multiple target poses, acquiring image pairs with and without projection, and use spatial matching to identify corresponding projector coordinates for each detected dot.

The key innovation in our method lies in replacing chessboard corner-based references with directly projected dot correspondences. Unlike previous approaches that rely on local homographies around corners, we reconstruct projector correspondences directly using the dots projected to the target surface. To enable highly accurate subpixel localization of dot centers in the camera image and ensure robustness of the calibration, we introduce several filtering stages of the dot correspondences (see Fig. 5 left):

- Matching confidence filtering: Only high-confidence dot matches from the spatial decoder are retained.
- Homogeneity filtering: Candidate dot positions are evaluated using edge detection to exclude those near texture boundaries.
- High-density rejection: To avoid spatial bias, dots in locally overpopulated regions are selectively removed.

Importantly, valid correspondences may lie outside the texture area of the calibration board, provided they fall within the physical planar target. For the remaining matches, we retain the precise, known integer positions in the projector coordinate system and subpixel-accurate locations in the camera view.

Since the 3D object coordinates of the projected dots are not directly known, we recover them by first estimating the pose of the planar target using solvePnP [6], applied to the chessboard corners detected in the reference image (projector off) and known camera intrinsics. From this pose, we derive the plane equation of the board in the camera coordinate system and compute the 3D coordinates of each dot via ray-plane intersection, using rays cast through the subpixel-accurate dot positions in the image.

This approach avoids the inaccuracies introduced by local homographies around sparse corners and instead reconstructs projector correspondences directly from the dense dot pattern projected onto the surface. By combining precise subpixel detection in homogeneous regions with a selective sampling of reliable correspondences, the method achieves improved geometric accuracy. A visual comparison between the previous and proposed approaches is shown in Fig. 4 and Fig. 5.

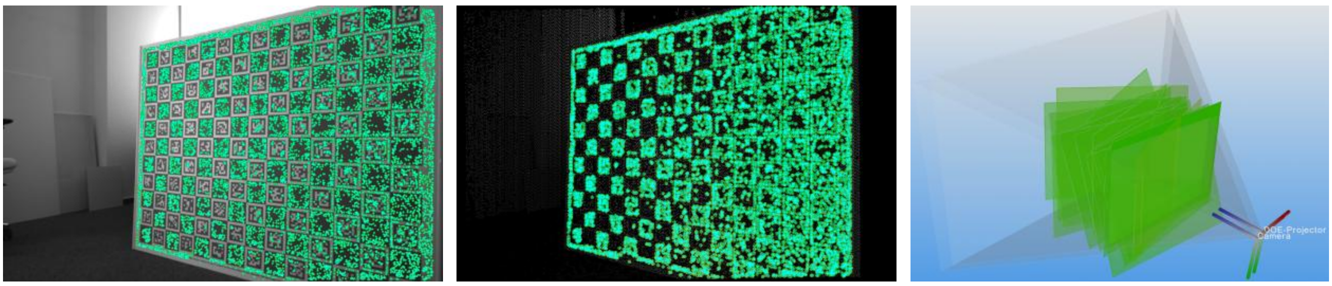

Fig. 5. Left: Camera view of target with sampled centers of projected dots. Center: Reconstructed projector view of the same target pose with corresponding dots. Right: Results of the projector-camera calibration and view of the target poses used for calibration. The same colorization of target poses was used as in Fig. 3 and Fig. 4. The overall rms reprojection error of our calibration with 16 target poses was 0.146 px.

3 Evaluation and Results

To assess the effectiveness of the proposed calibration method, we conducted a comparative evaluation against our previous calibration procedure using chessboard corner correspondences. Both methods use the same captured data and rely on our established spatial matching scheme to associate the fixed dot pattern with projector coordinates. An initial indication of improved accuracy is given by the root mean square (rms) reprojection error after stereo calibration, which is reduced from 0.321 px with chessboard corner correspondences to 0.146 px when using dot center correspondences. However, the reprojection error reflects different sets of input correspondences.

To enable a more meaningful comparison, we evaluated both calibration results on a new set of images of the planar calibration target at various unseen poses. For each pose, a 3D-pointcloud was reconstructed from matched camera-projector dot correspondences using triangulation. The plane equation of the target was estimated via solvePnP from the chessboard corners (projector off), and each 3D point's signed distance to the plane was computed. This allows us to quantify the geometric consistency of the calibration result by examining how well the reconstructed 3D dots lie on the known planar surface.

Fig. 6 shows representative results for five target poses, comparing the point-to-plane errors color-coded for the old calibration method (top row) and the proposed method (bottom row). The target poses cover a range of distances. Notably, the old method exhibits pronounced systematic deviations in the corners of the image (see the two rightmost columns), whereas the new method shows more spatially uniform error distributions.

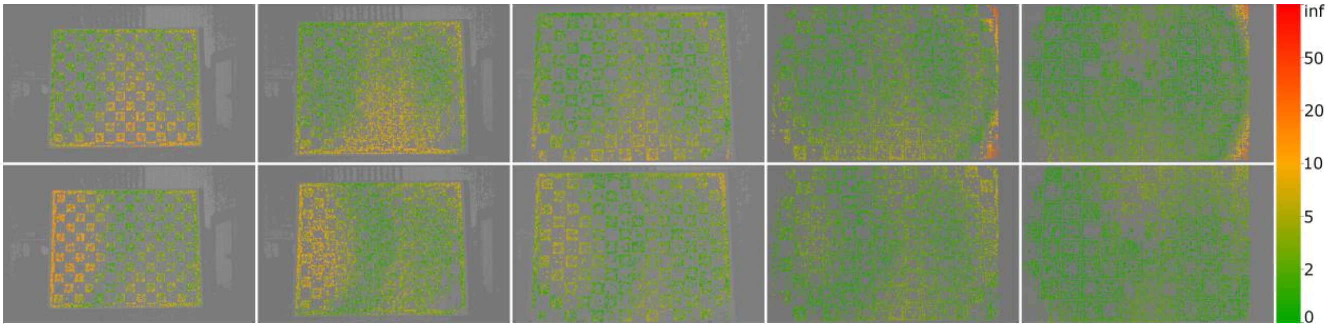

Fig. 6. Color-coded signed distances of points to the target plane, comparing calibration results of the prior approach (top row) to the proposed approach (bottom row). The color table shows plane distance error in mm.

Quantitatively, averaging over the same five target poses used in Fig. 6, the old method yields a mean point-to-plane error of 3.28 mm with a standard deviation of 2.63 mm, while the proposed method results in a similar mean error of 3.24 mm and a reduced standard deviation of 2.33 mm. Although the difference in mean error is not significant – especially considering that slight is warping was present in the physical calibration target used for both calibration and evaluation – the lower standard deviation for the proposed method indicates a more spatially consistent reconstruction quality. This improvement can be attributed to the avoidance of systematic calibration errors, particularly those occurring near the image boundaries in the previous approach.

4 Discussion

The proposed calibration method adapts established projector-camera calibration principles to the constraints of fixed-pattern structured light systems. By relying on subpixel-accurate detection of projected dot centers, it removes the need for relying solely on corner-based features, which can be affected by local texture interference. This is especially relevant for systems where dynamic structured light projection (e.g. Gray-code) is not possible due to fixed hardware. Our evaluation shows that the method achieves geometrically consistent calibration, with notably reduced systematic errors near image borders. This results in lower depth variance and more reliable 3D reconstructions, which is critical for applications requiring accurate geometry.

Fig. 7 demonstrates the benefits of precise calibration in a practical setup: a planar target with embedded spheres is captured by the active stereo sensor, and fused with RGB imagery. The alignment between depth and color data emphasizes the importance of accurate calibration not only for geometry but also for multi-sensor integration.

A key limitation, however, is that the fixed NIR dot pattern is invisible to RGB cameras, making corner-based methods still necessary for full system calibration. In our case, we used a ChArUco target, though standard detection omits some border corners due to marker pairing constraints. For future work, we plan to explore the use of Radon targets, which integrate circular markers into a chessboard layout and allow for the detection of corners even when the target is partially out of frame. This may further improve calibration robustness and coverage, especially in wide-baseline or off-axis configurations.

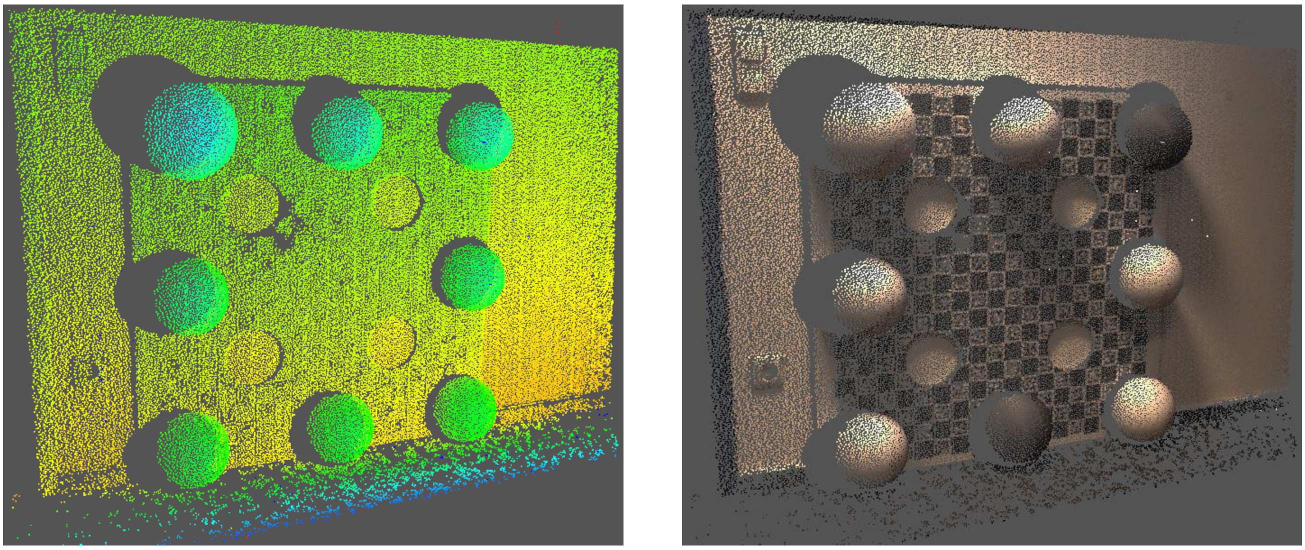

Fig. 7. 3D-pointcloud of our combined radiometric-geometric target captured by our in-house single-shot structured light sensor (left with color-coded depth in relation to the sensor position) and fusion with RGB camera data (right).

5 Conclusions

We presented a calibration method tailored for fixed-structured-light stereo systems that achieves high accuracy without requiring projection sequences or textured features. The approach is robust, practical, and well suited for scenarios where conventional calibration is infeasible. It also lays the groundwork for future self-contained or low-cost 3D stereo sensors with minimal setup and strong geometric fidelity.

The results presented in this article are based on our work within the joint project "Innovative Photonik für Automatische Kollaborative Systeme in dynamischen Waren-Transportprozessen" (AutoKoWaT), grant no. 13N16335. The project is funded by the German Federal Ministry of Research, Technology and Space (BMFTR). AutoKoWaT takes place within the framework of the funding program "Photonik für die digital vernetzte Welt – schnelle optische Kontrolle dynamischer Vorgänge".

References

[1] M. Hansard, S. Lee, O. Choi und R. P. Horaud, Time-of-flight cameras: principles, methods and applications, Springer Science & Business Media, 2012.

[2] E. Degtyarev, D. Reese und R. Nestler, „Modular Target and Approach for Geometric Calibration of Multimodal Imaging Systems," in 3D in Science & Applications" (3D-iSA), Berlin, 2024.

[3] D. Moreno und G. Taubin, „Simple, accurate, and robust projector-camera calibration," in Second International Conference on 3D Imaging, Modeling, Processing, Visualization & Transmission, 2012.

[4] Z. Zhang, „A flexible new technique for camera calibration," IEEE Transactions on pattern analysis and machine intelligence, pp. 1330-1334, 2002.

[5] D. Vehar, A. Hermerschmidt, R. Nestler und K.-H. Franke, Single-shot structured light with diffractive optic elements for real-time 3D imaging in collaborative logistic scenarios.

[6] L. Quan und Z. Lan, „Linear n-point camera pose determination," IEEE Transactions on pattern analysis and machine intelligence, Bd. 21, Nr. 8, pp. 774-780, 1999.

About us:

The Society for the Promotion of Applied Computer Science (GFaI) is an industry-oriented research institution with over 30 years of experience in applied computer science. We develop practical IT solutions for companies and institutions, particularly in the fields of optical measurement and testing technology, sound and vibrations, 3D data proces-sing,
robotics, energy and environmental technology, network infrastructure, and knowledge-based dialogue systems.

In addition to research, GFaI is involved in the training and qualification of skilled workers. In cooperation with universities, we promote knowledge transfer and support students and young talents on their way into practice.

Together with its subsidiaries gfai tech GmbH and Adalogic GmbH, GFaI supports the transfer of innovations to small and medium-sized enterprises in Germany and provides impetus for digital and sustainable developments.

Further information can be found at www.gfai.de.

Contact Information

3D-iSA & AI4EA:
eMail: 3disa-ai4ea@gfai.de
Web: www.gfai.de/aktuelles/veranstaltungen/gfaifuturetechdays

Conference Book GFaI Future Tech Days 2025

2nd edition

IMPRINT:

Publisher / Organization:
F. Böhm, B. Hohnhäuser, M. Bauer

Society for the Advancement of Applied Computer Science (GFaI)
Volmerstraße 3, 12489 Berlin

Phone : +49 30 814563-300
Fax: +49 30 814563-302
eMail: sekretariat@gfai.de
Web: www.gfai.de

ISBN-Number:
978-3-942709-37-8