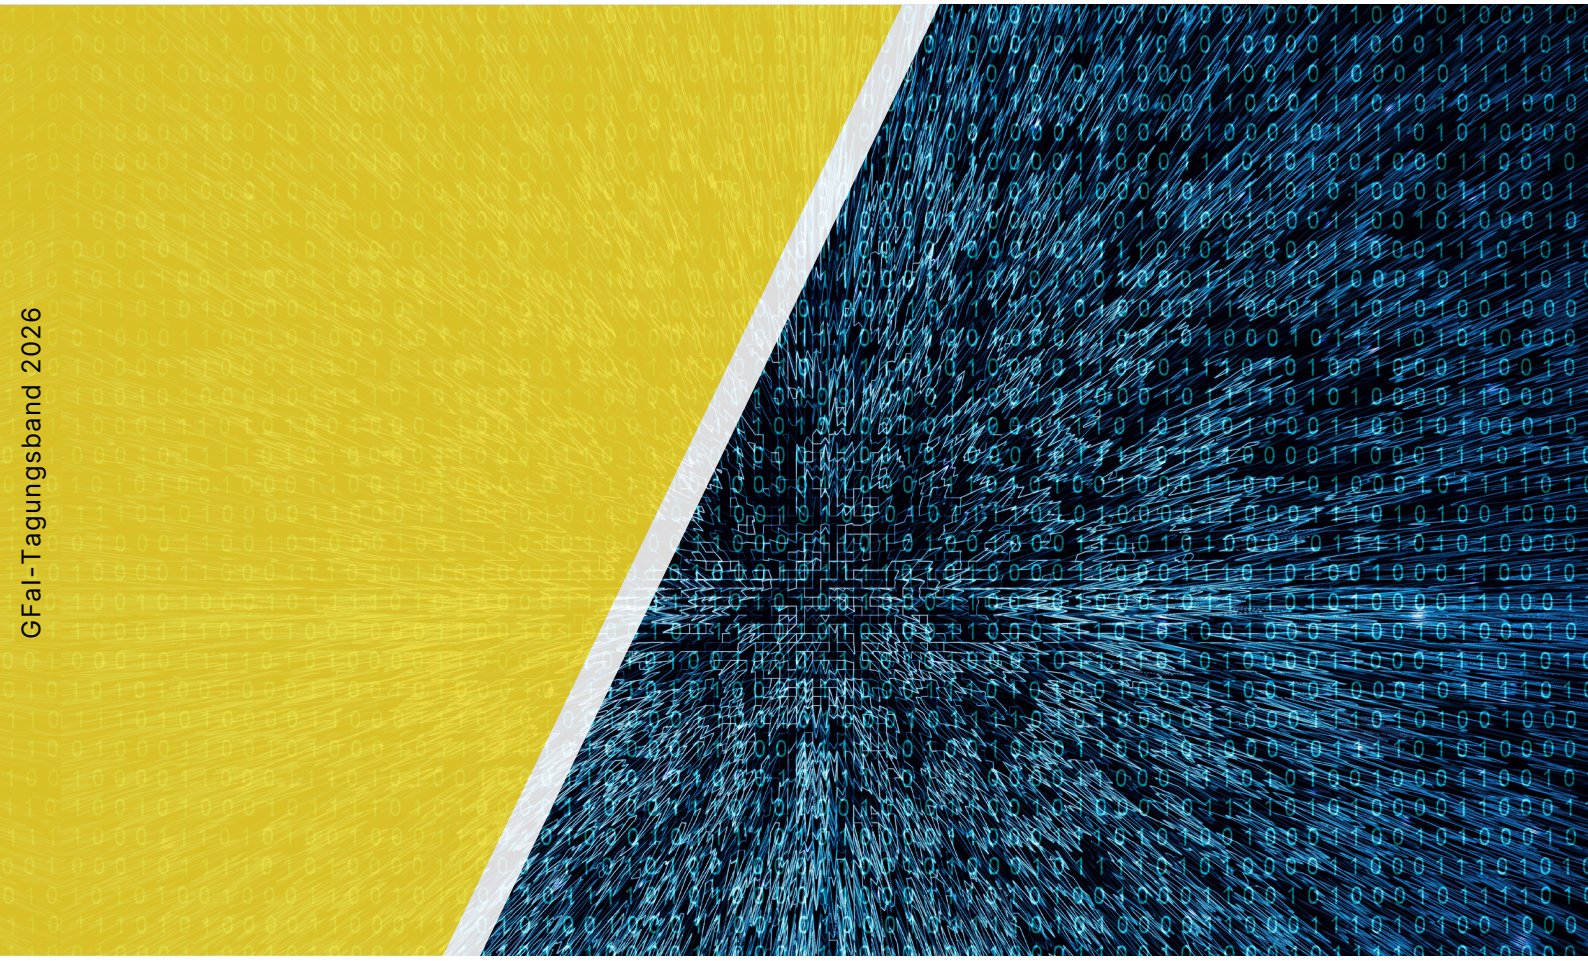
GFaI Future Tech Days

Tagungsband 2026



GFaI-Tagungsband 2026

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

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

Gesellschaft zur Förderung
angewandter Informatik e.V.

INHALTSVERZEICHNIS

GFaI-Tagungsband 2026

VORWORT

3 Vorwort der Workshopleitung

ASSISTENZ, MONITORING UND SICHERHEIT

4 EnaHuf – an AI powered embedded system for remote monitoring of horses

11 Was ich mir als Autofahrer bei der Bootsnavigation wünsche – Konzeption eines optischen Assistenzsystems für die Sportschifffahrt

17 Aktueller Arbeitsstand zur Patientenpositionierung in der Strahlentherapie - geometrische Systemmodellierung des Demonstrators

KI IN SIMULATION UND FERTIGUNG

23 From Data-Driven to Physics-Informed: Benchmarking Neural Surrogate Architectures for Transient Thermal Prediction in Industrial Casting

30 Additively Manufactured Sheet Metal Tools for High-Strength Steel Processing

3D-ERFASSUNG, REKONSTRUKTION UND KALIBRIERUNG

36 Per-Dot Calibration Approach for Single-Shot 3D Sensors Using Structured Light

43 Markerbasiertes Trackingsystem zur roboterunabhängigen Trajektoriengenerierung für automatisierte Beschichtungsprozesse

49 Entwicklung eines Rückenmodell-Generators zur robusteren Klassifizierung von Skoliose

56 Photogrammetric 3D Model-Based Animal Behavior Classification

3D-INSPEKTION, DEFEKT UND REPARATUR

64 Automated robust 3D defect reconstruction for additive repair processes

70 Evaluation of an AI-Based Concept for 2D-to-3D Defect Detection in Industrial Selective Repair Processes

76 Methoden zur Qualitätssicherung urbaner digitaler Zwillinge

MACHINE LEARNING UND DATENGETRIEBENE METHODEN

82 Detection of Anomalous Energy Consumption Patterns in Industrial Machines through Time Series Analysis to Improve Energy Efficiency

89 From LSTM Baselines to Edge-Deployed Mixture-of-Experts: Five Years of Neural Automatic Passenger Counting

95 RAISSE-Methodology for Interpretable Transfer Learning in Restricted Environments with an Illustrative Example on Gait Phase Classification

101 A data-driven optimization approach towards better automation in industrial fused granular fabrication

ZUKUNFTSTECHNOLOGIEN UND KI-SYSTEME

107 Dark Company: Autonomous Multi-Agent Systems and the Decoupling of Corporate Output from Headcount

113 OpenUSD as a Semantic Framework for GMP-Compliant Automated Optical Inspection

Per-Dot Calibration Approach for Single-Shot 3D Sensors Using Structured Light

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 present advances in the calibration pipeline for single-shot 3D sensors based on structured light dot pattern projection, building upon our previous work on projector-camera calibration via subpixel plane reconstruction. The primary contribution is a per-dot disparity-to-depth mapping that associates each projected dot with an individual hyperbolic depth model, enabling spatially resolved compensation of systematic depth errors across the full sensor working volume. To support the dense target sampling required for this calibration, we additionally present an improved chessboard target detection method related to OpenCV's Radon-transform-based corner detector, overcoming its key limitations regarding partially visible targets at image borders while also improving detection throughput for interactive use. Results are evaluated by comparing the systematic depth error of two sensor variants before and after the new calibration, providing a direct measure of improvement in 3D data quality.

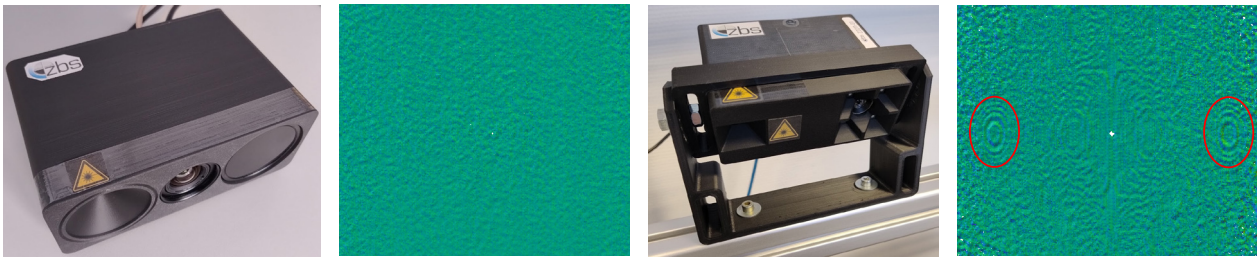
Keywords: Single-Shot 3D Imaging; Structured Light Calibration; Dot Pattern Projection; Per-Dot Disparity-to-Depth Mapping; Chessboard Target Detection; Radon Target Detection

1 Scientific Background and Motivation

Single-shot 3D sensors based on structured light dot pattern projection offer practical advantages for compact, real-time depth sensing. Since no sequential projection is required, these systems are inherently robust to scene motion. To exploit their full potential, for instance in 3D scanning and quality inspection applications, accurate geometric calibration is essential, both of the projector-camera system [1] and of the sensor's depth response across its working volume. This article evaluates two sensor variants built around laser projectors with diffractive optical elements (DOEs) projecting a fixed pseudorandom dot pattern. Both systems were intrinsically and extrinsically calibrated following the procedure described in [1] and using our calibration toolkit 3D-EasyCalib™ [2].

System1 (Fig. 1a) uses a single 830 nm dot pattern projector together with a NIR camera and additionally incorporates an RGB camera in the center. Instead of an RGB camera, *System2* (Fig. 1c) integrates a flood illumination unit alongside a 940 nm dot pattern projector, acquiring dot-pattern and flood-illuminated frames in alternation during operation. System2 is built with tighter size and cost constraints in mind, using a smaller back-illuminated sensor and a compact M12 lens. The two systems are furthermore configured for different working-distance ranges, System1's being shifted somewhat farther than System2's.

Figs. 1b and 1d show, for a flat reference surface, the color-coded directions of the reconstructed surface normal vectors when imaged with pattern I (System1) and pattern II (System2). Deviations from a uniform color indicate depth error, which is comparatively small for pattern I, whereas pattern II shows markedly larger deviations with a distinct wave-like structure. The generally higher error level of System2 is consistent with its lower-cost camera components. The wave-like structure, however, could clearly be attributed to the projector, being stationary in its view and independent of distance. System2's projector also covers a wider field of illumination, with a diagonal extent of approximately $86°$ versus $78°$. We suspect this contributes to the effect: a periodic DOE can only direct light into diffraction orders that are equally spaced in direction-cosine space [3], so that in IFTA-based design each dot of the target pattern is placed at the nearest available order [4]. The resulting residual position errors vary systematically across the field and form a moiré-like structure, which a wider field of illumination is likely to amplify.



(a) (b) (c) (d)

Fig. 1: Sensor heads of System1 (a), with an additional RGB camera in the center, and System2 (c), with an additional flood illumination unit. (b) and (d) show surface normal images (with color-coded normal vector directions), reconstructed from the projector reference frames of the respective systems. In (d), regions where the mentioned wave-like structures are best visible are marked in red. Gaps in the otherwise visible structure of the projected pattern are inpainted for better visualization.

Across both systems, 3D point cloud generation begins by locating, for each projected dot, its corresponding position in the NIR image. The displacement between a dot's position in the projected pattern and its observed position in the image yields that dot's disparity. Its component along the sensor rows varies with the distance of the illuminated surface and is the quantity used for the disparity-to-depth mapping. The identified point is then reprojected into space either via triangulation, using the shared intrinsic and extrinsic calibration obtained as in [1], or via the per-dot disparity-to-depth mapping presented in this article.

In our previous work [1], we demonstrated that replacing chessboard corner-based projector correspondences with directly localized dot centers yields more accurate projector-camera calibration with improved spatial consistency of the reconstructed 3D geometry. However, the standard parametric camera model, which describes sensor behaviour globally within the measurement volume using a limited set of intrinsic and extrinsic parameters, cannot capture position-dependent deviations in the disparity-to-depth conversion for individual projected dots, arising for example from imperfect lens geometries, manufacturing tolerances in the projector optics, or the DOE-related effects discussed above. Per-dot calibration addresses this remaining source of systematic depth error by modeling the disparity-to-depth relationship individually for each dot of the projected pattern.

To estimate the per-dot model parameters, target observations at many distances are required, covering the full sensor working volume, including boundary regions where the target is inevitably clipped by the image frame. Standard chessboard detectors, including OpenCV's `findChessboardCornersSB` as well as ChArUco-based methods, systematically discard valid border corners: the former requires detected corners to form a complete rectangular grid, and the latter only reports corners when both adjacent ArUco markers are fully visible. This motivated the improved target detection method described in Section 2.

2 Methods Used

Per-dot disparity-to-depth calibration

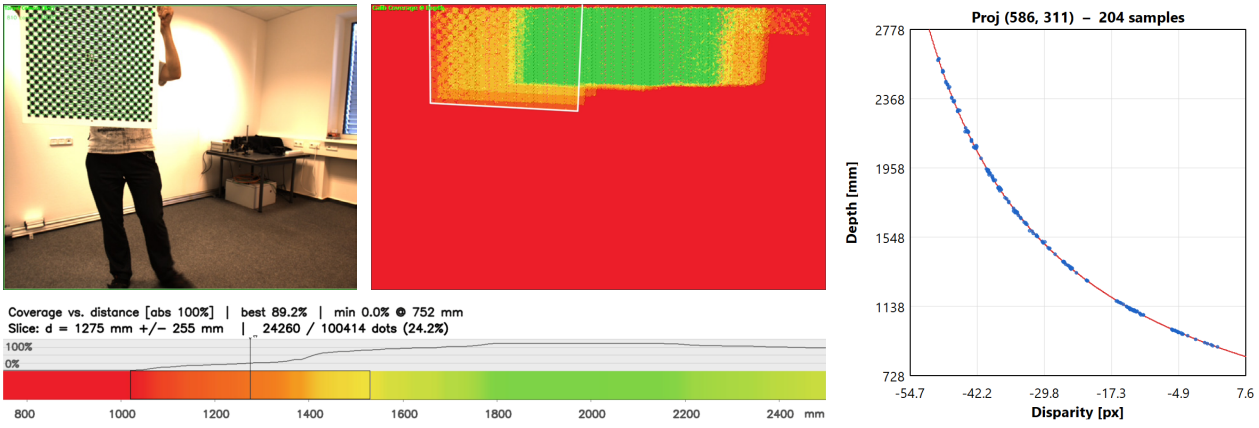
For a simplified triangulation-based system with parallel optical axes and baseline B, treating the projector-camera pair analogously to a stereo camera system and assuming a pinhole camera model in the sense of Tsai [5], depth Z and disparity d are related by

$$Z(d) = \frac{fB}{d},$$

with f denoting the camera constant, also referred to as the principal distance. This idealized relationship is a rectangular hyperbola in d (i.e. $Zd = fB = \text{const}$), motivating our choice of a hyperbolic function as the underlying model for the disparity-to-depth mapping. Real systems, however, deviate from this idealized geometry, e.g. through non-parallel optical axes, lens distortion, or other projector-specific systematic effects. A single hyperbola parameterized only by f and B can therefore not capture the true, spatially varying relationship. We therefore model the disparity-to-depth relationship for each projected dot by an individually fitted hyperbolic function, retaining the functional

form motivated above while allowing its parameters to vary across the pattern and thereby absorb these local systematic deviations. Parameters are estimated by capturing a planar target at multiple known distances throughout the working volume and fitting the model per dot using the reconstructed depth values. Data collection is supported by a semi-automatic sampling procedure: while the target is moved manually, frame acquisition is automated based on criteria such as target orientation and motion velocity, with visualization of spatial coverage.

The depth values to be sampled are defined in the NIR camera's coordinate system. For each frame, the target's pose is estimated from its detected corners and the intrinsic camera parameters using OpenCV's `solvePnP`. If detection succeeds and the pose satisfies the acceptance criteria discussed below, it defines a reference plane in the NIR camera frame, and intersecting each projector ray with it gives the corresponding surface point, whose distance along the NIR camera's optical axis serves as the reference depth for that dot. Obtaining the target pose reliably, and expressed in this frame, differs between the two sensor variants and constitutes the main practical difficulty of the sampling procedure. System1's auxiliary RGB camera with NIR-cutoff filter allows the target to be detected independently of the dot pattern, with its pose transformed into the NIR camera frame via the extrinsic RGB-to-NIR calibration, so that pose and depth data originate from the same instant. System2 has no such auxiliary view. Target detection instead runs on flood-illuminated frames acquired in alternation with the dot-pattern frames used for depth reconstruction, introducing a temporal offset between the two. Since any target motion during this offset directly biases the reconstructed reference depth, the sampling procedure discourages fast motion in particular along the sensor's depth axis, as an example of such a mitigating constraint. A lateral displacement of comparable magnitude affects the reference depth far less. Frames with excessive target tilt relative to the optical axis are excluded as well, since oblique views increase the sensitivity of both the plane fit and the resulting ray-plane intersection to small corner-localization errors. Fig. 2a shows the live visualization presented to the operator during sampling: detected target corners in the current camera frame (top left), together with two complementary views of accumulated sampling coverage, namely a projector-view slice restricted to a tolerance band around the current target distance (top right) and a profile of coverage against distance across the full working volume (bottom). This lets under-sampled regions of the pattern or working range be identified and revisited within the same session. Fig. 2b shows, alongside this, the accumulated disparity-depth samples and fitted hyperbolic model for a single, interactively selected dot, supporting a live check of individual per-dot fit quality.



(a) (b)

Fig. 2: Live sampling visualization. (a) Operator view during sampling, combining detected target corners in the current camera frame (top left), sampling coverage in projector view restricted to a tolerance band around the target's current distance (top right, color-coded from not yet sampled (red) to fully sampled (green)), and the corresponding coverage profile against distance across the working volume (bottom). (b) Disparity-depth samples and fitted hyperbolic model for a single, manually selected dot after the completed sampling process.

Reliable operator feedback during sampling depends on target detection that is both fast enough for interactive frame rates and tolerant of the frequent border-clipping that occurs while sweeping the target through the working volume, particularly near the edges of the pattern. Standard chessboard detectors satisfy neither requirement well, as discussed next.

Improved chessboard target detection

Our detection method uses the same corner feature extraction as OpenCV's `findChessboardCornersSB` and is related to it in its overall structure: a fast, point-symmetric candidate corner detector evaluating a localized Radon transform along multiple orientations (`FastX` in OpenCV's implementation) provides candidates for sub-pixel corner locations, following [6], whose evaluation of corner localization accuracy therefore applies comparably here. The primary implementation changes are a generalization of the output data structure and the underlying detection pipeline: rather than requiring detected corners to form a complete rectangular grid, our detector returns all corners whose board coordinates can be unambiguously determined via a set of circular markers, including those at clipped board edges and under partial occlusion elsewhere on the target. The `FastX` computation is additionally ported to OpenCL for GPU acceleration, improving detection throughput for interactive calibration workflows.

OpenCV's `findChessboardCornersSB` (tested with OpenCV 4.11.0 in C++) does in principle support recovering corners from partially visible targets with circular markers when using the `CALIB_CB_MARKER` and `CALIB_CB_LARGER` flags together with setting a small target size (e.g. 3×3 corners) as parameters, with the actual target being larger. The `meta` output array reports, per corner, whether it coincides with a marker, which resolves the board origin. This combination is comparatively undocumented and, in our experience, not widely known. However, the corners returned this way always form a single, complete rectangular sub-grid of the physical board: if any corner within that sub-grid cannot be localized, the corresponding row and column, and everything beyond it, is discarded as well, even where those corners are themselves unaffected and readily detectable. Figs. 3a and 3c show this for a border-clipped and a glare-affected target capture, respectively.



(a) (b) (c) (d)

Fig. 3: Comparison between the OpenCV detector `findChessboardCornersSB` (a, c) and the proposed method (b, d) for two failure cases: a target partially clipped by the image border (a, b) and a target with specular glare, while the dot markers remain visible (c, d). Detected corners are color-coded by their deviation from the respective estimated grid model, from green (small deviation) to red (large deviation).

Our detector shares its initial stages with `findChessboardCornersSB`: `FastX` supplies sub-pixel corner candidates across multiple scales, and initial board candidates are generated from these in essentially the same way. The differences lie in how such a candidate is subsequently grown into a board. Growth proceeds outward from a small seed region: for each board position adjacent to the already-detected set, an expected image location is predicted from previously detected neighbours along each of the two board axes. Where predictions from both axes are available, they must agree before the corner is accepted. Predicted locations are matched against the `FastX` candidates. Because growth follows connectivity rather than filling a bounding rectangle, the detected region can take on arbitrary, non-rectangular shapes instead of collapsing to the largest fully intact sub-grid.

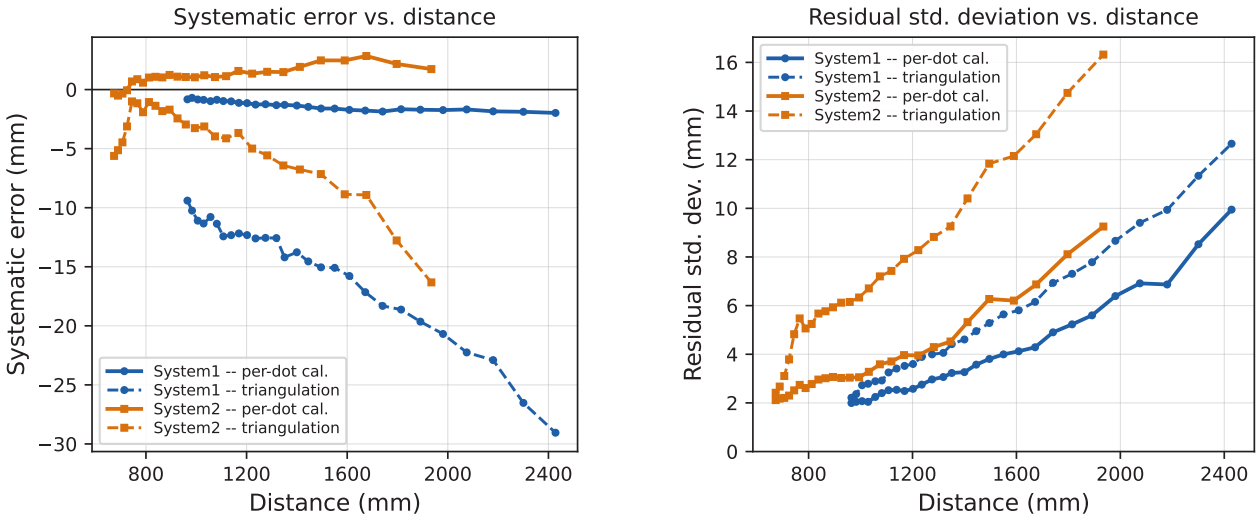
Where growth stalls at gaps, e.g. caused by occlusion, glare or the marker cutout itself, a second pass interpolates the missing positions from the surrounding geometry, again requiring agreement between independent estimates so that the fill remains confined to regions with supporting context. Interpolated corners are flagged as such in the output, so that downstream steps needing precise, image-measured locations – such as the per-dot calibration presented in this paper – can rely solely on directly detected corners, while interpolated ones still support board connectivity and origin resolution. Together, these measures recover the corners discarded by the original detector in both failure cases, as shown in Figs. 3b and 3d.

A related, more general approach to recovering occluded checkerboard corners is presented by Hillen et al. [7],

who fit a Gaussian process regressing image pixel coordinates from board coordinates as a detector-agnostic post-processing step on top of an initial corner detection, using the fitted model to fill in occluded or missing corners, extend the board beyond the image border, and re-assign initially misassigned corners. In contrast, our propagation is not a separate post-processing stage decoupled from the underlying detector, but performs interpolation directly during growth, using local, per-direction geometric estimates rather than a globally fitted regression model. Combined with the spatial keypoint index, this keeps the computational cost low enough for our targeted use case of interactive, high-throughput disparity-to-depth sampling.

3 Evaluation and Results

Depth accuracy of the 3D sensor systems is assessed using the same target-as-reference-plane principle as the sampling procedure (see Methods), but for measurement rather than calibration: for each accepted frame, the depth predicted from every matched dot correspondence is compared against the reference depth given by the intersection of the corresponding projector ray with the detected target plane, yielding a per-point residual. Two depth estimation methods are evaluated side by side on the same correspondences: the per-dot calibration presented in this paper, and the underlying triangulation-based approach it replaces, i.e. depth reconstruction from the shared intrinsic and extrinsic calibration [1] without per-dot disparity-to-depth mapping. Residuals are aggregated over many frames into the mean residual (systematic error) and the residual standard deviation at each distance, evaluated separately for each of the two systems described above. The target was kept within the central field of view so that each distance is evaluated over a comparable set of dots, since the aim here is to demonstrate the principle on functional prototype systems rather than to characterize them exhaustively. Owing to the different working-distance ranges noted above, the evaluated distances differ between the two systems. Fig. 4 shows both quantities across the full evaluated range, for both systems and both methods.



(a) (b)

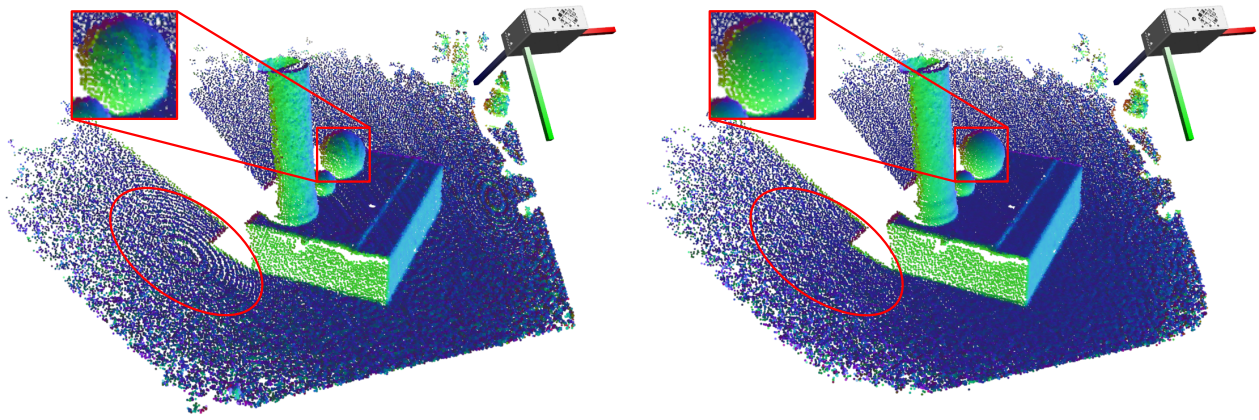
Fig. 4: Systematic error (a) and residual standard deviation (b) versus distance, for both systems and both depth estimation methods, over the full evaluated range. Solid lines: per-dot calibration. Dashed lines: triangulation. As described above, System1 and System2 differ in enough other respects that their curves should not be read as a direct ranking between the two designs. The relevant comparison is calibrated versus uncalibrated within each system.

For both systems, the per-dot calibration reduces the systematic error by roughly an order of magnitude at the far end of the working range. The residual standard deviation is likewise clearly reduced, by roughly a factor of two across most of System2's working range and by roughly a factor of 1.3 to 1.4 for System1. Residual error which is partially caused by sensor noise grows with distance for either method, as is expected for a triangulation-based system. Systematic error is largely reduced by the per-dot calibration, especially for larger distances.

4 Discussion

At each distance, residuals are pooled over all evaluated dots before computing their mean and standard deviation. The mean therefore reflects error components common to all dots, such as a global depth offset. Deviations that vary across the field of view, such as the wave-like structure, instead show up in the standard deviation. This structure is systematic: it is fixed in the projector's view and does not change with distance, so an individually fitted model per dot can absorb it. The marked reduction in standard deviation is therefore mainly due to the correction of this structure, and is correspondingly larger for System2, whose distortion was more pronounced. What remains is a stochastic component, caused for example by sensor noise. It grows with distance, since the disparity-to-depth mapping becomes increasingly sensitive to disparity errors. Being random, it cannot be compensated by any calibration.

Although the quantitative evaluation is restricted to the central region of the camera's field of view, the benefit of the per-dot calibration is not confined to that region. Fig. 5 shows pointclouds with color-coded surface normal directions for System2 on a scene containing several non-planar objects, reconstructed once by triangulation and once using the per-dot calibration. The wave-like structure characteristic of the uncalibrated system is visibly reduced across the whole field of view, including areas well away from the image centre, and object surfaces appear correspondingly more uniform.



(a) (b)

Fig. 5: Pointclouds with color-coded surface normal vector directions for System2 on a scene containing several non-planar objects, reconstructed via triangulation (a) and via the per-dot disparity-to-depth calibration (b). Notable differences are highlighted in red: The wave-like structure visible in (a) practically disappears with the per-dot disparity-to-depth calibration. Also note the improved surface quality e.g. on the larger sphere.

Strictly speaking, the per-dot calibration does not depend on the full calibration procedure of [1]: only the NIR camera's intrinsic parameters are required, to back-project each identified dot into the NIR camera's coordinate system using its per-dot depth and its direction from the intrinsic parameters. The extrinsic parameters remain necessary, however, whenever the point cloud is required in the projector's coordinate system, where the fixed pattern structure is directly available and supports efficient downstream processing such as outlier removal and surface normal computation.

The per-dot calibration involves a substantially larger overall parameter set than the handful of global intrinsic and extrinsic parameters used by triangulation. The per-point runtime cost can nevertheless be lower: in our case it replaces an iterative per-point undistortion and triangulation which takes considerably longer than a simple hyperbolic parameter lookup and evaluation. Its accuracy is nonetheless bounded by the quality of the underlying camera calibration, since the reference pose used during sampling is derived from it: errors in the intrinsic parameters of both systems, and additionally in System1's extrinsic RGB-to-NIR calibration, propagate into that pose. Real lenses moreover deviate somewhat from an ideal pinhole model, so distance-dependent intrinsic calibrations could be a worthwhile refinement where very high accuracy is required.

It is worth emphasising that the improvements reported here were obtained with a purely manual sampling procedure. Despite the error contributions this entails, a substantial reduction of systematic depth error is achieved,

which suggests that automating the acquisition offers room for further gains, for instance by moving the camera on a pan-tilt head and linear rail rather than the target.

5 Conclusions

We presented a per-dot disparity-to-depth calibration for single-shot structured-light 3D sensors, together with an improved chessboard target detection method that supports the dense, boundary-inclusive target sampling this calibration requires. Evaluated on two sensor variants, the approach substantially reduces systematic depth error compared to the underlying triangulation-based reconstruction, particularly at longer working distances, while a semi-automatic sampling procedure keeps the additional calibration effort manageable. These gains directly benefit applications such as 3D scanning and quality inspection, where spatially accurate depth across the full working volume is essential.

The individual contributions are not tied to the two sensor variants evaluated here. The per-dot calibration requires only a fixed, reproducible projected pattern and a way to establish correspondences, and is therefore applicable to such sensors in general; it is particularly attractive wherever cost constraints on the optical components lead to systematic deviations that global parametric models cannot capture. The semi-automatic sampling procedure and its coverage visualization transfer to any calibration task requiring dense, spatially resolved observations across a working volume. The improved target detection, finally, is of general use for camera calibration and extends the range of methods available in our calibration toolkit 3D-EasyCalib™ [2].

The results presented in this article are based on our work within the joint project "Implicit mobile human-robot communication for spatial action coordination with action-specific semantic environment modelling" (EDISON), grant no. 13N17577. The project is funded by the German Federal Ministry of Research, Technology and Space (BMFTR). EDISON takes place within the framework of the funding program "Photonik für die digitalisierte und automatisierte Produktion".

References

[1] D. Reese, E. Degtyarev, and R. Nestler, "Accurate projector-camera-calibration from structured light dot patterns via subpixel plane reconstruction," in *Anwendungsbezogener Workshop „3D in Science & Applications" (3D-iSA), Tagungsband, GFaI e. V.*, (Berlin), 2025.

[2] D. Vehar, R. Nestler, and K.-H. Franke, "3D-EasyCalib – Toolkit zur geometrischen Kalibrierung von Kameras und Robotern," in *22. Anwendungsbezogener Workshop zur Erfassung, Modellierung, Verarbeitung und Auswertung von 3D-Daten „3D-NordOst"*, (Berlin), p. 15, GFaI Gesellschaft zur Förderung angewandter Informatik e. V., Dec. 2019.

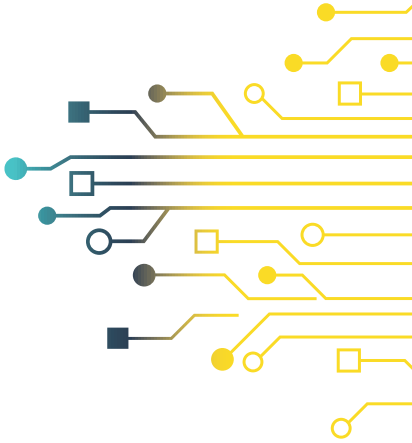
[3] J. E. Harvey and R. N. Pfisterer, "Understanding diffraction grating behavior: including conical diffraction and rayleigh anomalies from transmission gratings," *Optical Engineering*, vol. 58, no. 8, p. 087105, 2019.

[4] X. Liu, Y. Hu, S. Tu, C. Kuang, X. Liu, and X. Hao, "Fast generation of arbitrary optical focus array," *Optics and Lasers in Engineering*, vol. 162, p. 107405, 2023.

[5] R. Y. 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*, vol. 3, no. 4, pp. 323–344, 1987.

[6] A. Duda and U. Frese, "Accurate detection and localization of checkerboard corners for calibration," in *British Machine Vision Conference (BMVC)*, (Newcastle, UK), p. 126, 2018.

[7] M. Hillen, I. D. Boi, T. D. Kerf, S. Sels, E. C. D. L. Hoz, J. Gladines, G. Steenackers, R. Penne, and S. Vanlanduit, "Enhanced checkerboard detection using gaussian processes," *Mathematics*, vol. 11, no. 22, p. 4568, 2023.



Über uns:

Die Gesellschaft zur Förderung angewandter Informatik e. V. (GFaI) ist eine industrienahe Forschungseinrichtung mit über 30 Jahren Erfahrung in der angewandten Informatik. Wir entwickeln praxisorientierte IT-Lösungen für Unternehmen und Institutionen, insbesondere in den Bereichen optische Mess- und Prüftechnik, Schall und Schwingungen, 3D-Datenverarbeitung, Robotik, Energie- und Umwelttechnik, Netzwerkinfrastruktur sowie wissensbasierte Dialogsysteme.

Neben der Forschung engagiert sich die GFaI in der Ausbildung und Qualifizierung von Fachkräften. In Zusammenarbeit mit Hochschulen fördern wir den Wissenstransfer und begleiten Studierende sowie Nachwuchstalente auf ihrem Weg in die Praxis.

Zusammen mit ihren Tochterunternehmen gfai tech GmbH und Adalogic GmbH unterstützt die GFaI den Transfer von Innovationen in den Mittelstand Deutschlands und setzt Impulse für digitale und nachhaltige Entwicklungen.

Weitere Informationen finden Sie unter www.gfai.de.

Kontaktdaten

GFaI Future Tech Days
E-Mail: futuretechdays@gfai.de
Internet: www.gfai.de/aktuelles/veranstaltungen/gfai-future-tech-days/workshop

IMPRESSUM

Herausgeber / Organisation:
F. Böhm, B. Hohnhäuser, M. Bauer

Gesellschaft zur Förderung angewandter Informatik e. V. (GFaI)
Volmerstraße 3, 12489 Berlin

Telefon: +49 30 814563-300
Telefax: +49 30 814563-302
eMail: sekretariat@gfai.de
Internet: www.gfai.de

ISBN-Nummer:
978-3-942709-38-5