

Full length article

Depth-adaptive calibration of line-laser 3D measurement systems using rational polynomial mapping and depth-independent misalignment compensation

Ling Cao^{a,b}, Zidong Lin^{b,c}, Yong Yang^c, Quan Tang^b, Zhipeng Gao^b, Daquan Feng^a, Wei Pan^{b,d,*}

^a College of Electronics and Information Engineering, Shenzhen University, Nanshan District, Shenzhen, Guangdong Province, China

^b OPT Machine Vision, Department of R&D, 66 Xingfa South Road, Dongguan, Guangdong Province, China

^c Guangdong Polytechnic Normal University, School of Mechatronic Engineering, No. 293 Zhongshan Avenue, Guangzhou, Guangdong Province, China

^d OPT Machine Vision Tech Co., LTD. Japan, Department of R&D, Aomi 2-7-4-735, Koto, Tokyo, Japan

ARTICLE INFO

Keywords:

Laser triangulation
Non-contact optical metrology
Depth-adaptive calibration
Rational polynomial mapping
Sensor-to-motion misalignment compensation

ABSTRACT

Line-laser triangulation is a principal modality in non-contact 3D optical metrology, yet two coupled error sources erode measurement accuracy across the working depth of field: depth-dependent drift in the pixel-to-physical mapping, and systematic geometric distortion arising from sensor-to-motion-frame rotational misalignment. Existing calibration pipelines typically address either the internal pixel mapping or the external installation geometry, but their joint effect across the depth of field is seldom modelled explicitly, forcing depth-dependent retuning in practice. We propose an automated two-stage calibration framework that explicitly decouples these error sources. The internal stage drives a precision Z-axis stage to acquire a planar target across the full measurement volume and fits a 22-coefficient rational polynomial (RPC) model that fuses a perspective backbone with a bivariate cubic residual, capturing depth-dependent nonlinearities that single-plane homographies cannot represent. The external stage decomposes the installation misalignment into roll-pitch-yaw rotations and estimates the shear-producing roll and yaw angles α and γ from height-block, profile, and checkerboard measurements, and applies a depth-independent inverse rotation in Z - Y - X order to remove installation-induced distortion globally. On a commercial line-laser profiler, the proposed framework reduces the X - and Z -direction mean absolute errors by 17.5% and 11.8% relative to the HALCON baseline, suppresses the depth-mean angle error of a 25° standard block from 0.7591° to 0.0387° (94.90%), and reduces checkerboard orthogonality error from 10.007° to 0.309° (96.92%). The depth-wise suppression rate stays within 94.69%–96.97% across six discrete depth positions Z_1 – Z_6 , indicating that the proposed calibration strategy improves both local mapping accuracy and global geometric consistency within the tested measurement volume.

1. Introduction

Line-laser triangulation is one of the primary modalities in non-contact 3D optical metrology: it combines compact hardware, fast line-profile acquisition, and strong robustness to ambient conditions, and is consequently applied to precision dimensional measurement, profile inspection, and surface geometry characterisation [1–3]. Measurement accuracy is, however, entirely determined by calibration quality. In a laser triangulation sensor, the calibrated mapping from stripe-pixel coordinates to physical coordinates determines the measurement uncertainty within the depth of field; if this mapping cannot be maintained

consistently across the working depth of field, both local profile fidelity and the global geometric consistency of the reconstructed point cloud degrade, limiting the utility of line-laser optical metrology in demanding measurement applications.

Most existing camera and structured-light calibration methods are built on planar-target paradigms that descend from Zhang's flexible technique and its plane-based extensions [4–6], and a substantial body of follow-up work refines plane-target strategies for line-structured-light and laser-stripe sensors [3,7,8]. These pipelines, however, mostly assume a fixed imaging configuration or a locally linear mapping; in many practical settings they still require multi-pose target acquisitions

* Corresponding author at: OPT Machine Vision, Department of R&D, 66 Xingfa South Road, Dongguan, Guangdong Province, China.
Email address: vp@foxmail.com (W. Pan).

<https://doi.org/10.1016/j.optlastec.2026.116090>

Received 25 May 2026; Received in revised form 2 July 2026; Accepted 30 July 2026

Available online 7 August 2026

0030-3992/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

with substantial manual intervention, and their accuracy degrades when working distance changes or when non-ideal mechanical factors are introduced. Recent work that introduces automated procedures [9,10] mainly targets operational efficiency and does not address how the internal mapping error and the global installation distortion couple as the depth changes.

A more critical challenge arises because line-laser sensors are integrated with a motion mechanism and reconstruct 3D point clouds by stitching multiple frames. Internal calibration only fixes the pixel-to-local mapping within a single frame; it cannot guarantee geometric consistency between the sensor frame and the motion axis. When the two are misaligned, the stitched cloud carries a systematic *global* affine shear—not a small local residual—and under varying depth this shear cannot be treated as a stack of independent planar errors without multiplying the parameter count and eroding practicality. Robotic vision offers rich work on extrinsic geometry, notably hand-eye calibration [11–14], but those formulations do not resolve the depth-varying shear that originates from sensor–motion misalignment in line-laser measurement. The distinction is structural rather than incidental: hand-eye calibration solves for a single rigid 6-DoF transform $\mathbf{X}$ from an $\mathbf{AX} = \mathbf{XB}$ system relating end-effector and camera motions, treating the sensor as a black-box pose, whereas the error we target is the depth-varying shear that the same fixed misalignment imprints on a *stitched* multi-frame point cloud. Our external stage estimates the shear-producing installation angles directly from geometric standards and inverts them in closed form, so it neither requires a robot’s controlled motion sequence nor resolves an extrinsic pose per workpiece; the internal RPC stage, which has no counterpart in hand-eye formulations, separately repairs the depth-dependent pixel-to-physical mapping that hand-eye calibration leaves untouched. Recent unified calibration frameworks for line-structured-light scanning [15] still do not jointly address depth-adaptive internal mapping and depth-independent installation compensation within a single measurement-volume framework.

To address this gap, we introduce a depth-adaptive automated framework that decouples the internal pixel-to-physical mapping from the external installation compensation. “Depth-adaptive” here is operational: the internal RPC mapping is fitted from correspondences spanning the full working depth of field and evaluated on held-out depth planes, while the external Z – Y – X compensator is depth-independent and reused across the whole volume, so the framework generalises across depths without recomputing per-plane parameters. Because the two error families differ in physical origin and in the level at which they act, folding them into one empirical correction blurs their boundaries; we instead handle them in two cooperating stages. The first uses a precision Z -axis stage to automatically sample a planar target across the full depth of field and fits a rational polynomial mapping from laser-stripe pixels to local physical coordinates—a motorised-stage acquisition pattern recurring in line-structured-light and galvanometric calibration [7,16–18] that preserves the planar-target foundation [4,5] while adding the residual capacity needed for non-ideal industrial behaviour. The second models external installation as three rotational misalignments, estimates them from height-block, profile, and checkerboard observations, and applies a single Z – Y – X inverse rotation to the reconstructed cloud, so the same compensator restores geometry at every depth without per-plane retuning.

Experiments on the tested commercial line-laser profiler (OPT-LPF2-006-01) show that this two-stage decoupling delivers micrometre-level local mapping accuracy and sub-degree global geometric consistency across the working depth of field, validating the theoretical structure and calibration accuracy of the framework within the tested measurement volume. The main contributions are:

1. A rational polynomial camera (RPC) mapping model for line-laser systems that fuses a perspective skeleton with a bivariate cubic

residual, raising internal mapping accuracy and absorbing nonlinear deviations that single-plane homographies cannot represent.

2. A depth-decoupled external compensation mechanism that estimates the shear-producing installation angles (roll α and yaw γ) independently and corrects the point cloud with a single Z – Y – X inverse rotation, restoring shear-free geometry uniformly across the depth range without per-depth fitting.
3. An automated calibration workflow that integrates the two stages on a precision Z -stage and a push-broom rig, replacing the multi-pose target operations of conventional pipelines with programmed motion-triggered acquisition while achieving micrometre-level local mapping accuracy on the tested system.

The remainder of the paper is organised as follows. Section 2 analyses the challenges in line-laser calibration and depth adaptivity. Section 3 presents the proposed two-stage method. Section 4 reports experiments and results. Section 5 concludes.

2. Challenges in line-laser 3D calibration and depth adaptivity

The accuracy of a line-laser 3D vision system in real engineering use is dictated by the quality of its calibration, but conventional pipelines reveal clear limits in three areas: depth-of-field adaptability, installation-error compensation, and global error control. This section traces these limits to their theoretical origins and motivates the two-stage framework introduced in Section 3.

2.1. Perspective geometry and the single-plane mapping assumption

The theoretical foundation of line-laser 3D measurement is optical triangulation under a perspective imaging model, which provides a unified starting point for both calibration (intrinsic and distortion) and light-plane calibration (laser stripe to camera coordinates) [3,19].

If distortion is ignored and the laser plane is fixed, the measurement points are confined to a 2D laser plane, and the relation between the pixel coordinates (u, v) and the local physical coordinates (X, Z) collapses to a strict homogeneous-coordinate linear mapping [4], expressible by a single homography matrix $\mathbf{H}$:

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{K} [\mathbf{R} \mid \mathbf{t}] \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}, \quad (1)$$

where s is a scale factor, $\mathbf{K}$ is the intrinsic matrix, and $(\mathbf{R}, \mathbf{t})$ are the extrinsic rotation and translation. For line-structured-light systems, calibration is conventionally split into camera calibration (intrinsic and distortion) and light-plane calibration (laser stripe to camera coordinates) [3,19].

If distortion is ignored and the laser plane is fixed, the measurement points are confined to a 2D laser plane, and the relation between the pixel coordinates (u, v) and the local physical coordinates (X, Z) collapses to a strict homogeneous-coordinate linear mapping [4], expressible by a single homography matrix $\mathbf{H}$:

$$s \begin{bmatrix} X \\ Z \\ 1 \end{bmatrix} = \mathbf{H} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix}. \quad (2)$$

A small set of non-collinear correspondences from a planar target is sufficient to solve $\mathbf{H}$ in closed form; this single-plane geometric prior is shared by mainstream line-structured-light calibration approaches that use planar, circular, or cylindrical targets and that further introduce unified extrinsic estimation or scanning-direction calibration on top of it [4,8,19–22]. Fig. 1 illustrates this mapping under ideal conditions.

This perspective-plus-single-plane formulation makes the geometric starting point of internal calibration clear, but it rests on three idealisations: a flat target, an ideal alignment, and a locally linear approximation. Real line-laser pipelines break these idealisations through finite acquisition efficiency, manual involvement, and non-ideal error structures, which is why the next subsection examines the engineering limits of the traditional internal-calibration paradigm.




Fig. 1. Ideal mapping in line-laser triangulation. A spatial point $P(X, Z)$ on the laser plane projects to a pixel $p(u, v)$ on the camera sensor under a strictly linear perspective relation, in the absence of distortion and with a fixed light plane.

2.2. Limits of the traditional internal-calibration paradigm

Although single-plane homographies establish a usable pixel-to-physical mapping under ideal conditions, their engineering applicability is constrained when realistic non-idealities are introduced. Recent work refines the internal calibration of line-structured-light sensors along several directions, including spatial-geometry-based light-plane calibration [19], high-precision planar-target calibration [8], large-range integrated calibration-measurement [18], circular-target calibration [20], freely placed cylindrical-target calibration [21], unified extrinsic estimation for shallow-depth high-precision measurement [22], and virtual-point-constraint calibration [23]. These advances improve target design, parameter solving, and operational flexibility, yet they all centre on the explicit identification of camera and light-plane parameters.

For line-laser 3D measurement, however, internal calibration is not only a question of model solvability; it must also balance acquisition efficiency, manual involvement, parameter stability, and the ability to express realistic errors. While methods that change target geometry, optimise feature construction, or adopt joint estimation [8,18–23] push performance upward, most still depend on specific targets, specific acquisition scripts, and an idealised geometric story. Once a system enters a real production line, small mechanical tilts in the lifting stage, local non-flatness of the target, radial and tangential lens distortion, and stripe-centre extraction noise all inject residual errors that drive observations away from the ideal perspective. A purely geometric homography then captures the dominant trend but cannot represent the full non-linear deviation.

Recent studies further show that residual error in internal calibration is not produced by light-plane fitting alone; feature localisation noise, shallow-depth blur, and noise re-injected by per-frame extrinsic estimation also contribute substantially [22,23]. The implication is that the goal of internal calibration is not to drop the geometric model, but to retain the perspective constraint while adding a representation that absorbs realistic non-ideal errors. Line-laser internal calibration therefore needs a hybrid model that is simultaneously geometrically grounded and expressive enough to describe non-ideal residuals, in order to deliver a stable and accurate pixel-to-physical mapping in industrial conditions.

Even with internal mapping improved, however, a system deployed on the factory floor still suffers from external geometric distortion when

the sensor is not perfectly aligned with the motion mechanism, so improving local mapping alone does not guarantee global measurement consistency.

2.3. Cross-depth shear from installation misalignment demands a unified external compensator

External installation misalignment is a separate and equally important error source. In a typical industrial scan, full 3D point clouds are obtained by combining mechanical motion with multi-frame stripe stitching. Internal calibration only fixes the local pixel-to-physical mapping inside one frame and leaves the spatial pose between the camera frame and the motion frame unconstrained. As illustrated in Fig. 2(a), installation tilt causes the laser plane to deviate from the stage motion direction, resulting in stripe defocus, systematic shear deformation, and rotational distortion in the stitched point cloud. Throughout this paper, we use “shear” to describe the affine appearance of this distortion in the stitched coordinate system, while the physical source is rotational misalignment between the sensor frame and the motion frame; the corrective transform is therefore a rigid inverse rotation in 3D rather than a generic 2D shear of the stitched plane. Recent work on scanning-direction calibration, rotational-axis estimation, and hand-eye calibration in line-structured-light systems has shown that inaccurate external geometry directly breaks the coordinate consistency and stitching accuracy of 3D reconstruction [24–27]; in a closely related setting, automatic unwarping of rotating line-scan imagery has likewise been shown to require explicit recovery of the sensor–motion geometry, with circular-pattern elliptical distortion serving as the calibration cue [28], which further confirms that line-scan systems cannot rely on internal calibration alone.

This installation deviation can be decomposed into rotations α , β , γ around the three orthogonal axes X , Y , Z , and naïve sequential stitching turns the deviation into an affine deformation. As an example, when the mechanism scans an orthogonal rectangular workpiece along the Y -axis with $\gamma \neq 0$, laser points sharing the same true Y coordinate are pushed in X as a linear function of the stitching distance ($\Delta x = \Delta y \tan \gamma$), so the stitched rectangle becomes a parallelogram (Fig. 2(b)). Several existing approaches model this from different angles—hinge-connected targets [29], planar-target hand-eye calibration for line-structured-light robots [30], position-constrained robot




Fig. 2. Geometric origin of stitched-cloud shear. (a) The camera world coordinate frame is misaligned with the actual mechanism rotation direction by an angle θ . (b) Ideal orthogonal scanning preserves the rectangular checkerboard, while installation tilt turns the reconstructed grid into a parallelogram.

compensation [31], and free-placed cylindrical scanning reconstruction [27]—and they all share the conclusion that sensor-to-motion geometry must be modelled explicitly. The shear distortion they describe is a systematic physical error caused by a non-orthogonal motion frame; it cannot be removed by tuning intrinsics or distortion coefficients, and only an external coordinate-frame model with an inverse spatial transform can correct it.

Under varying depth of field, this geometric distortion should not be corrected by per-depth compensation parameters. Doing so multiplies the parameter count, breaks model uniformity, and adds operational fragility in deployment. Studies on scanning-direction calibration and unified extrinsic estimation suggest a stronger alternative: enforce a consistency constraint directly on the external geometry so that the same compensator restores global coordinates across all sampling conditions [24,26,27]. The right design target for line-laser systems is therefore a depth-decoupled external shear compensator, not a depth-indexed family of local fixes. Beyond detecting shear, the harder question is how to fold its compensation into a unified external framework that simultaneously preserves global geometric consistency and keeps the compensation procedure simple under varying depth—which in turn means internal mapping and external compensation must be organised cooperatively, not in isolation.

2.4. Lack of a unified two-stage framework coupling internal mapping and external shear

The error sources analysed above belong to two physically distinct families: pixel-to-local mapping deviations inside each frame, and global geometric distortion produced when frames are stitched under installation misalignment. They differ in physical origin, in the level at which they act on the measurement chain, and in their error profile, and without a clean stage boundary between them, system error is hard to describe or correct in a targeted way. Existing surveys of extrinsic and multi-sensor calibration likewise note that engineering systems rarely balance accuracy, automation, and robustness simultaneously, and this tension is most acute when multiple error sources couple [32].

Most existing line-laser calibration pipelines, however, treat internal mapping and external compensation separately. One family focuses on camera or light-plane parameters and on local geometric mapping while leaving the stitching-stage external pose unmodelled; another family applies post-hoc corrections after observing global distortion, but those corrections are not consistently tied to the internal mapping model. Work on simultaneous robot-world/hand-eye calibration, simultaneous hand-eye/robot-world/camera-IMU calibration, and unified parameter estimation for line-structured-light systems repeatedly shows that handling the extrinsic and error-budget pieces in isolation blurs error boundaries, aggravates parameter coupling, and amplifies error propagation along the measurement chain [9,13–15].

The problem is sharper still under varying depth of field. Treating installation-induced distortion as a stack of independent planar errors and adding a compensation parameter per depth multiplies the parameter count, breaks model uniformity, and erodes engineering deployability. Recent work begins to explore closed-loop alternatives via integrated calibration, self-calibration, and unified extrinsic estimation. Some studies fold sensor pose into the inspection task itself to compensate for pose-induced error [33]; others use line-structured-light sensors to enable robot self-calibration and reduce the cost and uncertainty of sequential calibration in the field [10]; in the broader structured-light family, fringe-projection systems have likewise moved toward self-calibrating pipelines that recover internal and external parameters jointly from the projected pattern [34], again pointing to the same closed-loop design principle. The shared message is that the system needs a single framework that ties internal coordinate recovery and external geometric compensation together, rather than a sequence of local patches.

High-precision line-laser calibration is therefore not a single-model problem. It should be cast as a two-stage unified framework: an internal mapping stage that recovers local coordinates accurately, and an external shear-compensation stage that enforces global consistency across the working depth of field. A loose “calibrate first, patch later” chain is not enough; the full measurement chain needs explicit parameter relations, clear error boundaries, and a controlled engineering procedure—which is precisely the design starting point of the method introduced next.

To position the proposed framework against recent line-structured-light calibration work, Table 1 contrasts representative methods along the axes that matter for depth-of-field-wide industrial measurement: target type, degree of acquisition automation, whether the internal pixel-to-physical mapping is made depth-adaptive, whether sensor-to-motion installation misalignment is explicitly compensated, and whether accuracy is validated across multiple depth planes rather than at a single working distance. Recent internal-calibration methods push target design and parameter solving [8,19–23] but treat the mapping at essentially one configuration and do not model external installation shear; recent external/scanning-direction and hand-eye methods recover the motion geometry [15,26,27] but leave the internal depth-dependent mapping to a conventional single-plane model. The proposed framework is the one entry that couples a depth-adaptive internal mapping with a depth-independent external compensator and validates both across six depth planes, which addresses the gap identified above.

3. Proposed line-laser 3D calibration method

We instantiate the two-stage framework as a depth-adaptive line-laser calibration method. Stage one builds a depth-adaptive RPC mapping that absorbs nonlinear residuals while preserving the perspective backbone. Stage two estimates and inverts external installation angles

Table 1

Qualitative comparison with representative recent line-structured-light calibration methods. *Depth-adaptive map.*: internal pixel-to-physical mapping fitted/validated across the depth of field rather than at a single plane. *Ext. misalign. comp.*: explicit compensation of sensor-to-motion installation misalignment. *Cross-depth valid.*: accuracy reported at multiple depth planes. ✓ = addressed, × = not addressed, ~ = partial.

Method	Target	Automated	Depth-adaptive map.	Ext. misalign. comp.	Cross-depth valid.
Zhang et al. [8]	planar	~	×	×	×
Gao et al. [19]	planar	×	×	×	×
Fu et al. [20]	circular	~	×	×	×
Guo et al. [21]	cylinder	✓	×	×	×
Li et al. [22]	planar	~	~	×	×
Mu et al. [23]	planar	~	×	×	×
Li et al. [26]	planar	~	×	✓	×
Deng et al. [27]	cylinder	✓	×	✓	×
Wang et al. [15]	planar	~	×	✓	×
Proposed	planar + block + checkerboard	✓	✓	✓	✓

Table 2

Notation used in the method (Section 3) and experiments (Section 4).

Symbol	Meaning
(u, v)	laser-stripe centre pixel coordinates
(X, Z)	local physical coordinates in a single laser-plane frame
(x, y, z)	corrected global 3D coordinates of the stitched cloud
(x', y', z')	distorted stitched coordinates observed before correction
α, β, γ	roll (X), pitch (Y), yaw (Z) installation angles
$\mathbf{R}_x, \mathbf{R}_y, \mathbf{R}_z$	elementary rotation matrices used in the inverse compensation
L	RANSAC reference line fitted on the empty bottom-plane stripe
$\mathbf{e}_x$	platform x -axis unit vector (laser-plane in-plane direction at mounting time)
H, H'	nominal and measured upper-face height of the standard height block
W	block width along the platform x -axis
$Z_1 - Z_6$	six discrete depth positions sampled by the precision Z -axis stage

in a closed-form Z - Y - X rotation. The two stages are wired together by an automated motion-triggered workflow that drives a precision Z -axis stage for internal sampling and a push-broom rig for external evaluation, so that the framework runs end-to-end without manual multi-pose target placement.

Symbols used throughout the method and experiments are summarised in Table 2. The convention is: (u, v) for pixel coordinates extracted from the laser-stripe image; (X, Z) for the local physical coordinates returned by the internal RPC mapping inside a single laser-plane frame; (x, y, z) for the corrected global 3D coordinates of the stitched cloud after the external Z - Y - X inverse rotation; and (x', y', z') for the distorted stitched coordinates observed before correction.

3.1. Depth-adaptive RPC mapping

The internal stage establishes the mapping from the laser-centre pixel coordinates (u, v) to the local physical coordinates (X, Z) . For a fixed measurement plane and an approximately ideal alignment, the homography in Eq. (2) captures the perspective relation well and provides a stable geometric backbone for further compensation. Reusing the homogeneous-coordinate form, the local mapping satisfies

$$\begin{bmatrix} X \\ Z \\ 1 \end{bmatrix} \propto \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix}. \quad (3)$$

In real systems, slight tilt of the lifting stage, local non-flatness of the target, and radial-tangential lens distortion push the observations away from this ideal perspective. We make the deviation explicit by defining a height-direction residual function $r(u, v)$ that captures the gap between the perspective-predicted height and the measured height Z :

$$Z = \frac{h_{21}u + h_{22}v + h_{23}}{h_{31}u + h_{32}v + 1} + r(u, v). \quad (4)$$

We then approximate $r(u, v)$ by a bivariate cubic polynomial whose terms span the dominant non-linear deformation directions while staying numerically stable:

$$r(u, v) = \sum_{i=0}^3 \sum_{j=0}^{3-i} a_{ij} u^i v^j. \quad (5)$$

Here the coefficients a_{ij} are the unknown weights of the bivariate-cubic height residual: each a_{ij} scales the monomial $u^i v^j$ over the index set $\{(i, j) : 0 \leq i + j \leq 3\}$ (ten terms in total), and they are not prescribed in advance but are solved jointly with the perspective parameters in the least-squares fit below. The transition from the residual model to the rational form is then a direct algebraic merge. Substituting Eq. (5) into Eq. (4) writes Z as the perspective ratio plus the cubic residual; placing the perspective numerator of Eq. (3) and the cubic residual $r(u, v)$ over the same linear denominator $1 + s_{10}u + s_{01}v$ and collecting like monomials in the numerator collapses the two contributions into a single ratio of a bivariate-cubic numerator to a linear denominator. Carrying out the same merge for both output coordinates yields the rational polynomial camera (RPC) model that maps (u, v) to the world coordinates X and Z :

$$X = \frac{\sum_{i=0}^3 \sum_{j=0}^{3-i} p_{ij} u^i v^j}{1 + s_{10}u + s_{01}v}, \quad (6)$$

$$Z = \frac{\sum_{i=0}^3 \sum_{j=0}^{3-i} q_{ij} u^i v^j}{1 + s_{10}u + s_{01}v}. \quad (7)$$

Here (u, v) are the laser-line centre pixel coordinates; X and Z are the corresponding world coordinates, with Z the physical height. The numerator coefficients p_{ij} and q_{ij} contribute 10 terms each and absorb high-order distortion; the shared denominator carries the perspective scale through the two coefficients s_{10} and s_{01} , giving a total of 22 unknowns. Concretely, each q_{ij} is the sum of the corresponding perspective-numerator term and the residual weight a_{ij} of Eq. (5), so the a_{ij} are absorbed into the q_{ij} and never appear as free parameters of the final model; p_{ij} is formed analogously for the X coordinate. The 20 numerator coefficients plus the 2 denominator coefficients are therefore exactly the 22 unknowns counted above. We linearise Eqs. (6)–(7) by cross-multiplication, $X(1 + s_{10}u + s_{01}v) = \sum p_{ij} u^i v^j$ and similarly for Z , which makes the unknowns $\{p_{ij}, q_{ij}, s_{10}, s_{01}\}$ linear in the residual; assembling the multi-depth correspondences then yields an overdetermined linear system that is solved by ordinary least squares. To keep the cubic Vandermonde matrix well-conditioned, we normalise the pixel coordinates to $[-1, 1]$ via $\tilde{u} = (u - u_0)/u_s$ and $\tilde{v} = (v - v_0)/v_s$, where (u_0, v_0) is the image centre and (u_s, v_s) is the half-width of the image, and the world coordinates (X, Z) are likewise centred on the working volume; the cubic monomials are then evaluated on $(\tilde{u}, \tilde{v})$, the linear system is solved by SVD-based least squares for numerical stability, and the recovered coefficients are de-normalised back to the pixel domain. Across the full

working depth of field this assembly contributes correspondences from all sampled Z -positions, well above the 22 unknowns, so the system is heavily overdetermined and converges without explicit Tikhonov regularisation; we further hold out one depth plane during fitting and report all metrics in Section 4.2 on this held-out test plane to guard against overfitting. By coupling a geometric constraint with a data-driven residual, the model attains sub-pixel accuracy without explicit recovery of physical optical parameters.

3.2. Decoupled estimation and Z - Y - X inverse compensation of installation angles

Internal calibration only resolves the pixel-to-local mapping, but industrial scanning relies on the combination of mechanical motion and multi-frame stitching, so the spatial pose between the camera frame and the motion frame must be modelled explicitly. When the stripe-stitching reference and the motion direction are not parallel, the originally orthogonal physical structure undergoes affine deformation in the stitched cloud—the systematic shear distortion analysed in Section 2.3. We use the term “shear” descriptively to denote the visual character of this distortion in the stitched cloud, but the underlying cause is rotational misalignment of the sensor frame with respect to the motion frame; the corrective transform is therefore a rigid inverse rotation rather than a generic affine shear, and this distortion cannot be eliminated by tuning intrinsics. Accordingly, we decompose the misalignment into rotations around the three orthogonal axes and invert them in 3D space.

We decompose the installation deviation into three orthogonal rotation angles and assign each one a clear physical meaning:

- γ (Z -axis): the deviation of the camera coordinate system in the XY plane relative to the motion direction;
- β (Y -axis): the deviation in the XZ plane (camera pitch) relative to the motion direction;
- α (X -axis): the deviation in the YZ plane (camera roll).

The shear example from Section 2.3 makes this concrete: scanning by Δy along the y -axis with a γ angle between the stitching direction and the y -axis pushes points by $\Delta x = \Delta y \tan \gamma$, an offset that no internal model can absorb. We therefore estimate the shear-producing angles α and γ separately, analyse β as part of the physical decomposition, and invert the effective misalignment on the reconstructed cloud through the Z - Y - X formulation, with β set to zero in the tested push-broom configuration.

3.2.1. Decoupled angle estimation

We first define a reference line L shared by the α and β estimators: scanning the empty platform bottom plane yields a single-line profile $\{(x_i, 0, z_i)\}$ that collapses to $\{(x_i, z_i)\}$ on the laser plane, and a robust RANSAC fit on these points returns the bottom-plane reference line L .

We adopt RANSAC for every line- and plane-fitting step in this stage because the points returned by line-laser scanning are corrupted by structured, non-Gaussian outliers—edge rounding at block boundaries, specular bright spots, occlusion drop-outs, and locally non-planar surface patches—rather than by zero-mean Gaussian noise alone. Under such gross outliers an ordinary least-squares fit is pulled toward the contaminated points, and even bounded-influence M-estimators (Huber, Tukey) only down-weight rather than reject them, so a cluster of edge-rounding points still biases the geometry; the least-median-of-squares (LMedS) estimator tolerates outliers but breaks down once the contaminated fraction approaches 50% and offers no explicit inlier set for downstream refinement. RANSAC instead partitions inliers from outliers by maximising the consensus set and returns a model supported only by the geometrically consistent points, which matches the failure mode of laser-stripe data and yields a clean inlier set for sub-pixel re-fitting. In our implementation a candidate model is hypothesised from a minimal sample, scored with a point-to-model distance threshold of 0.05 mm, iterated for up to 1000 hypotheses at a target confidence of 0.99, and accepted only when the inlier ratio exceeds 0.7; the surviving

inliers are then refined by total least squares. The same threshold/confidence/inlier-ratio settings are reused for the upper-face plane fit below and for the per-cell checkerboard fits in the γ estimator.

α (X -axis) estimation. A standard height block of nominal height H is then scanned with one face resting on the bottom plane, with its long edge oriented along the platform x -axis (the laser-plane in-plane direction defined by the camera mounting), so that W denotes the block width along x . The average vertical distance from the stripe centre points on the upper face to L , denoted H' , is computed by RANSAC plane fitting on the upper-face stripe points to suppress local non-flatness, contamination, and edge-rounding outliers. Target flatness nonetheless sets a floor on the achievable accuracy and propagates into the calibration in two distinct ways. A local out-of-plane deviation δ on the upper face shifts the mean height H' and hence biases the roll estimate by $\delta\alpha \approx \delta/W$ through Eq. (8), while a non-flat calibration plate enters the internal RPC fit directly as a spurious Z -residual that the cubic term cannot distinguish from genuine optical distortion. For the gauge-grade grid distortion target used here, whose NIST-traceable certificate specifies a substrate flatness of a few micrometres over the field of view, the first term contributes well below 0.01° to α at the block width W of this setup, and the RANSAC plane fit together with multi-depth averaging further attenuates isolated non-flat patches; a coarser machined target would raise both error floors in proportion to its flatness, which is the reason a certified flat target is used for all quantitative results. The roll angle is

$$\alpha = \arctan\left(\frac{H' - H}{W}\right). \quad (8)$$


Eq. (8) reads α from the apparent height-versus-width slope and is exact when $\beta = \gamma = 0$. For finite but small β and γ , the upper-face mean height H' is biased only at second order in β and γ , so within the nominal industrial mounting tolerance $|\alpha|, |\beta|, |\gamma| \lesssim 5^\circ$ the three estimators decouple to first order and can be solved independently in a single pass. This 5° figure marks where the first-order decoupling is provably clean rather than a hard bound on validity: the single-pass estimator continues to recover the angles and remove the dominant shear well beyond it—validated to a roll of $\alpha = 20^\circ$ in Section 4.3.4—while at extreme tilt an optional joint iteration of the three estimators is available to drive the second-order residual back down.

β (Y -axis) estimation. The same line L also encodes the pitch deviation. We define the platform x -axis as the in-plane direction of the laser-plane xz frame anchored at sensor mounting time: it is the horizontal axis returned by the camera intrinsic calibration when the bottom plane is taken as the reference, and remains fixed across calibration sessions. Under an ideally aligned mounting, L would coincide with this x -axis; the signed angle between L and the platform x -axis therefore directly reads off the camera pitch around the Y -axis,


$$\beta = \angle(L, \mathbf{e}_x), \quad (9)$$

where $\mathbf{e}_x$ denotes the platform x -axis unit vector. Decoupling from α is again first-order: roll perturbs L only through the projection of out-of-plane points onto the laser plane, an effect that scales as $\sin \alpha \sin \beta$ and is below the stripe-extraction noise floor for $|\alpha|, |\beta| \lesssim 5^\circ$. In the push-broom configuration studied here, however, β rotates the laser plane about its own axial direction and therefore imprints no in-plane shear on the stitched cloud. Eq. (9) is retained for completeness of the physical decomposition, while the compensation in Section 3.2.2 sets $\beta = 0$ and corrects only the shear-producing roll and yaw (Section 4.3).

γ (Z -axis) estimation. The yaw misalignment γ between the laser plane normal and the stage motion direction acts as an affine shear on the stitched coordinates rather than as a rigid 2D rotation; we make this physical model explicit and derive the estimator from it, then specialise the implementation form for continuity with the 3D rigid compensation in Section 3.2.2 (Fig. 3).



(a) Before correction: the four-point quadrilateral $ABCD$ on the scanned checkerboard is sheared.



(b) After correction: $A'B'C'D'$ are restored to right angles by the recovered γ .

Fig. 3. Closed-form γ estimation from a checkerboard target. Enforcing $\overrightarrow{A'B'} \cdot \overrightarrow{A'D'} = 0$ on the corrected coordinates yields the quadratic in $\sin \gamma$ used in Eq. (14).

Forward stitching map. Let $\mathbf{e}_x$ be the laser-plane in-plane direction aligned with world X at $\gamma = 0$, and let the stage move along world Y . The derivation below uses the laser-plane coordinate convention, in which x' is the projection of the world point onto the rotated laser-plane in-plane direction and y' is the stage encoder coordinate biased by the lateral intersection offset; this is equivalent, up to an exchange of which axis carries the linear coupling, to the illustrative $\Delta x = \Delta y \tan \gamma$ shear used in Section 2.3. Under this convention, for a world point $(X, Y, 0)$ on the bottom plane, a sensor rotated by γ around the world Z -axis observes

$$x' = X \cos \gamma - Y \sin \gamma, \quad y' = X \tan \gamma + Y, \quad (10)$$

where x' is the laser-plane in-plane coordinate and y' is the stage encoder position at which the laser plane intersects the world point. Eq. (10) is an affine map with determinant $\sec \gamma$; the cross-term $X \tan \gamma$ in y' is the source of the visible shear in the stitched cloud and prevents the distortion from being inverted by any pure 2D rotation (which would preserve inner products). **Inverse recovery map.** Inverting Eq. (10) yields

$$X = (x' + y' \sin \gamma) \cos \gamma, \quad (11)$$

$$Y = -x' \sin \gamma + y' \cos^2 \gamma. \quad (12)$$

This inverse contains a γ -dependent shear that actively redistributes the inner product between adjacent edge vectors. Applying Eqs. (11)–(12) to the four checkerboard corner points $A(x_1, y_1)$, $B(x_2, y_2)$, $C(x_3, y_3)$, $D(x_4, y_4)$ observed in stitched coordinates, denoting $\Delta x_{AB} = x_2 - x_1$ etc., and enforcing orthogonality of adjacent recovered edges, $\overrightarrow{A_w B_w} \cdot \overrightarrow{A_w D_w} = 0$, the cross-coupling terms in $\sin \gamma \cos \gamma$ cancel via $\sin^2 \gamma + \cos^2 \gamma = 1$, leaving

$$\Delta x_{AB} \Delta x_{AD} + \cos^2 \gamma \Delta y_{AB} \Delta y_{AD} = 0, \quad (13)$$

which is equivalent to the quadratic in $\sin \gamma$

$$A_\gamma \sin^2 \gamma + B_\gamma \sin \gamma + C_\gamma = 0, \quad (14)$$

with coefficients $A_\gamma = -\Delta y_{AB} \Delta y_{AD}$, $B_\gamma = 0$ in the bottom-plane case (a non-zero B_γ first-order term arises when the checkerboard departs from the bottom plane), and $C_\gamma = \Delta x_{AB} \Delta x_{AD} + \Delta y_{AB} \Delta y_{AD}$, all determined from the four observed coordinates. Eq. (14) has at most one root in $[-1, 1]$ within the installation range $|\gamma| \lesssim 10^\circ$ relevant here; if both algebraic roots fall in $[-1, 1]$ (a degenerate case caused by near-square checkerboards), we retain the one with smaller $|\sin \gamma|$, since the other root corresponds to the spurious 90° -rotated solution. Applying $\arcsin$ then recovers γ . To suppress per-corner detection noise and local non-flatness of the target, we apply Eq. (14) to every interior 2×2 corner cell of the 7×11 checkerboard rather than a single quadrilateral, reject cells whose corner-detection RANSAC inlier ratio falls below 0.7,

and report the median of the surviving per-cell estimates as the final $\hat{\gamma}$ together with the median absolute deviation across cells. The 5° figure quoted throughout is the vendor's *nominal* mounting tolerance for the OPT-LPF2-006-01 kit rather than a hard limit on the estimator: the root selection above stays single-valued and the recovered $\hat{\gamma}$ remains accurate up to the 10° yaw tested in Section 4.3.4, which is the largest the fixture can mechanically realise.

Small-angle implementation form. For consistency with the 3D rigid Z - Y - X compensation in Section 3.2.2, which operates on full 3D points and is naturally expressed as a rotation, the in-plane γ estimator is implemented as a small-angle inverse rotation that maps the observed stitched corner (x'_i, y'_i) to its corrected counterpart $(\tilde{x}_i, \tilde{y}_i)$:

$$\tilde{x}_i = x'_i \cos \gamma - y'_i \sin \gamma, \quad (15)$$

$$\tilde{y}_i = x'_i \sin \gamma + y'_i \cos \gamma. \quad (16)$$

The apparent difference between the affine forward map (determinant $\sec \gamma \neq 1$, Eq. (10)) and the rigid inverse rotation (Eq. 20) reflects two distinct levels of description and does not represent a theoretical inconsistency. The affine model characterises the *observed distortion signature* in the stitched coordinate system: because the laser plane intersects the motion path at an oblique angle, the stitched y' coordinate carries a cross-coupling term proportional to $X \tan \gamma$, producing an apparent non-unit-determinant mapping. The rigid 3D rotation, by contrast, is the physically correct corrective transform for the underlying cause—a sensor-frame misalignment that is a rigid spatial rotation. For $|\gamma| \lesssim 5^\circ$ the rotation matrix in Eqs. (15)–(16) matches the affine inverse in Eqs. (11)–(12) to first order in γ (with $\sin \gamma$, $\tan \gamma$, and $1 - \cos \gamma$ agreeing to within 0.4%), and the orthogonality condition under either form yields a numerically indistinguishable root of Eq. (14) within and somewhat beyond the nominal tolerance; the second-order residual between the two inverse forms is bounded by $O(\gamma^2) < 0.4\%$ and falls below the stripe-extraction noise floor. We therefore report the affine derivation as the forward distortion model, derive the γ estimator from the orthogonality constraint (Eq. 14), and implement the correction as the rigid 3D rotation for consistency with the full Z - Y - X inverse compensation in Eq. (20).

3.2.2. Inverse Z - Y - X rotation

To restore an undeformed point $P(x, y, z)$ from the distorted observation $P'(x', y', z')$, we invert the three rotations in the order $Z \rightarrow Y \rightarrow X$, which matches the physical chain in which the misalignments enter the sensor mounting (Fig. 4): the yaw γ around Z is set first when the sensor head is rotated to the working orientation, the pitch β around Y is then introduced by the front-to-back tilt of the fixture, and the roll α around X is the residual side-to-side tilt of the bracket. Because the three angles lie within the nominal small-angle regime $|\alpha|, |\beta|, |\gamma| \lesssim 5^\circ$, the rotation matrices commute to first order ($\sin \theta \approx \theta$, $\cos \theta \approx 1 -$




Fig. 4. Decoupled Z - Y - X inverse rotation. The mismatched stitching axis (Y') is first re-aligned to the world Y -axis through $\mathbf{R}_z(\gamma)$, then pitch and roll are removed by $\mathbf{R}_y(\beta)$ and $\mathbf{R}_x(\alpha)$; in the tested push-broom setup β is set to zero because it produces no in-plane shear, so the effective correction is $\mathbf{R}_x(\alpha)\mathbf{R}_z(\gamma)$. The composition in Eq. (20) is independent of working depth.

$\theta^2/2$), so any chosen order recovers the same angles up to second-order cross-terms below the noise floor. The six possible orderings of three elementary rotations— Z - Y - X , Z - X - Y , Y - Z - X , Y - X - Z , X - Y - Z , and X - Z - Y —differ only by these second-order cross-terms $O(\alpha\beta)$, $O(\beta\gamma)$, $O(\alpha\gamma)$. Moreover, any residual difference between two orderings is itself a rigid rotation, which leaves the intrinsic angle metrics E_θ and E_{orth} invariant, so the orderings are interchangeable for the reported geometry rather than merely within a narrow tolerance band. We deliberately select Z - Y - X among them because it is the exact inverse of the physical order in which the misalignments are imprinted at mounting time (yaw set first, then fixture pitch, then bracket roll, as detailed below), so each inverse step removes the most recently applied physical rotation first and the decomposition retains a clean physical reading. This interchangeability is confirmed numerically across all six orders in Section 4.4.2. The three step matrices are

$$\mathbf{R}_z(\gamma) = \begin{bmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (17)$$

$$\mathbf{R}_y(\beta) = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix}, \quad (18)$$

$$\mathbf{R}_x(\alpha) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{bmatrix}. \quad (19)$$

The corrected coordinate is then

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \mathbf{R}_x(\alpha)\mathbf{R}_y(\beta)\mathbf{R}_z(\gamma) \begin{bmatrix} x' \\ y' \\ z' \end{bmatrix}. \quad (20)$$

Because Eq. (20) is depth-independent, the same compensator restores geometry at every working distance once the misalignment angles are estimated, which removes the need for per-depth re-fitting. In the tested push-broom configuration the pitch β imprints no in-plane shear (Section 3.2.1) and is set to zero in Eq. (20), so the effective correction reduces to $\mathbf{R}_x(\alpha)\mathbf{R}_z(\gamma)$.

The 5° figure is the regime in which the three estimators are guaranteed to decouple to first order, not a hard ceiling on the compensator; it is nonetheless worth stating explicitly how the second-order terms grow as the installation tilt exceeds it. Three effects compound with growing $|\alpha|$, $|\beta|$, $|\gamma|$. First, the cross-axis bias that couples the α and β estimators scales as $\sin \alpha \sin \beta$: it stays below the stripe-extraction noise floor at 5° but rises into the 0.1° range by 10° . Second, the gap between the affine forward-shear model and the rigid inverse rotation, bounded by $O(\gamma^2)$

and under 0.4% at 5° , grows quadratically to a few percent near 10° . Third, the assumption that the three angles can be estimated independently in a single pass weakens, so a small residual shear can leak into the corrected cloud. These terms grow only gradually, however, and the single-pass closed-form compensator continues to remove the dominant shear far beyond the nominal tolerance: the direct tilt sweep in Section 4.3.4 confirms that a roll of $\alpha = 20^\circ$ (four times the tolerance) and a yaw of $\gamma = 10^\circ$ are still corrected to $E_\theta \leq 0.10^\circ$ and $E_{\text{orth}} \leq 0.6^\circ$, with the recovered $\hat{\alpha}$ tracking the physical roll to within 0.11° . Only if sub- 0.1° residuals are required at such extreme tilt does it become worthwhile to iterate the estimators jointly—re-estimating each angle from the partially corrected cloud until convergence—which trades the closed-form single-pass property for a further reduction of the second-order residual; within and well beyond the vendor tolerance the single-pass form used throughout is sufficient.

3.3. Automated calibration workflow

We integrate the two stages into a four-step automated workflow that separates the one-off factory geometric calibration from the on-site installation compensation. Fig. 5 shows the data flow from raw images to a calibrated 3D point cloud.

- 1. Automated multi-depth acquisition.** The camera and line laser are rigidly mounted on a precision Z -axis stage; precision linear stages have long been used to translate planar targets through controlled depths for line-structured-light and galvanometric calibration [7,16–18], and we adopt the same hardware pattern under closed-loop motion-triggered control. The stage moves through a programmed start position and step size, and at each height the system triggers the laser-on/laser-off sequence to acquire image pairs that yield clean stripe centrelines.
- 2. Internal RPC fitting.** From the multi-height image stack, calibration-target feature points and stripe centre pixel coordinates are extracted; cross-multiplying Eqs. (6)–(7) assembles an overdetermined linear system in the 22 RPC unknowns, which is solved by ordinary least squares.
- 3. Installation-angle estimation.** A standard height block, the platform bottom plane, and a checkerboard target are scanned via the push-broom rig, and the shear-producing angles α and γ are estimated independently using Eqs. (8)–(14), with β set to zero as it produces no in-plane shear (Section 3.2.1).
- 4. Global compensation.** Every reconstructed point is mapped through Eq. (20) to remove shear globally, producing a depth-consistent calibrated 3D point cloud.




Fig. 5. End-to-end two-stage workflow. *Stage 1 (factory calibration)*: a precision Z-stage drives multi-depth target acquisition; the laser-on/laser-off pair yields clean stripe centres for feature extraction; the $(u, v) \rightarrow (X, Z)$ correspondences are stacked, Eqs. (6)–(7) are cross-multiplied into a linear system, and the 22 RPC coefficients are solved by ordinary least squares. *Stage 2 (on-site compensation)*: a brightness-image checkerboard is mapped to the depth image, the deflection angles are computed by Eqs. (8)–(14), and the inverse Z – Y – X rotation in Eq. (20) produces the calibrated 3D point cloud.

3.4. Design choices and rationale

Three design choices in Sections 3.1 and 3.2 are made on physical and engineering grounds: the bivariate cubic residual order in the RPC model, the Z – Y – X inverse-rotation order for shear compensation, and the strict separation of the internal and external stages. Section 4 reports ablation evidence: the residual-order sensitivity (Section 4.4.3), the rotation-order interchangeability (Section 4.4.2), and the contribution of the internal vs. external stages in isolation versus combined (Section 4.4.1).

The bivariate cubic residual order deserves explicit justification, since it is neither the simplest nor the most expressive option. We choose cubic as the lowest order that captures the dominant non-ideal deformation while remaining numerically stable. A first-order (homography-only) residual cannot represent the curved height-direction deviation that lens distortion and stage tilt introduce; a quadratic residual still underfits the antisymmetric edge-distortion lobes characteristic of line-laser optics, because even-order monomials alone cannot model the sign change across the field centre. A quartic or higher residual, conversely, adds five or more monomials per axis that inflate the condition number of the cubic Vandermonde system and begin to absorb stripe-extraction noise, so training error keeps falling while held-out-depth error rises—the classic over-fitting signature. The cubic order therefore sits at the bias–variance knee: rich enough to fit the third-order optical-edge distortion, restrained enough to stay well-conditioned under ordinary least squares. An order sweep over quadratic, cubic, and quartic residuals that substantiates this choice is reported in Section 4.4.3.

4. Experiments and results

We evaluate the two stages of the proposed framework against the design goals from Section 2. Section 4.1 describes the hardware, target set, and acquisition protocol. Section 4.2 evaluates the internal RPC mapping against the HALCON baseline on X - and Z -direction repeatability and mean absolute error. Section 4.3 evaluates the external compensator across six depth positions on two complementary geometric tasks (angle measurement and orthogonality), under three installation states (normal / tilted-uncompensated / tilted-compensated).

Table 3

Depth positions Z_1 – Z_6 used for internal calibration and external evaluation. The positions are commanded by the closed-loop precision Z-stage at a uniform 30 mm encoder step, spanning a 150 mm working-depth range. Z_3 is held out as the internal test plane.

Position	Z_1	Z_2	Z_3	Z_4	Z_5	Z_6	Step
Encoder reading (mm)	440	470	500	530	560	590	30

Train/test depth split. The six depth positions Z_1 – Z_6 are produced by the precision Z-stage at a uniform encoder step of 30 mm, spanning a total working-depth range of 150 mm from the lower- to the upper-extreme acquisition plane (Table 3); they uniformly sample the working volume so that the calibration is exercised across the whole depth of field rather than near a single best-focus plane. The internal RPC mapping is fitted on five of the six discrete depth positions Z_1 – Z_6 , and the remaining depth, Z_3 , is held out as the test plane for the metrics reported in Table 5 and Fig. 8. The held-out plane is chosen near the centre of the working volume so that it falls inside the convex hull of the training depths; the held-out plane therefore probes interpolation accuracy within the working volume, not boundary-depth extrapolation, and the latter is not claimed in this experiment. The same train/test split is used for the HALCON baseline, so neither pipeline sees the test plane during fitting. The external angle estimators in Section 3.2 are calibrated once from the bottom-plane and standard-block scan at the central depth Z_3 , and the resulting estimate ($\hat{\alpha} = 5.150^\circ$, $\hat{\gamma} = 9.830^\circ$) is then frozen, after which the depth-wise external evaluation in Section 4.3 and the three-way ablation in Section 4.4.1 reuse the same frozen angles across all six depths Z_1 – Z_6 ; the six depths in the external evaluation therefore act as test depths for the external compensator.

4.1. Experimental platform and targets

The experimental platform integrates a commercial line-laser 3D camera with a programmable Z-stage and a push-broom rig (Fig. 6). The camera is an OPT-LPF2-006-01 (Opt Machine Vision Tech Co., Ltd., Guangdong, China). The high-precision Z-stage drives multi-height

Table 4

Traceable metrology of the reference standards used for ground truth. Each artefact is an independently certified physical standard, so the evaluation compares the reconstruction against an external invariant rather than against the calibration output.

Reference standard	Certified accuracy	Traceability
Grid distortion target (internal dot array)	pitch $\pm 1 \mu\text{m}$, dot size $\pm 3 \mu\text{m}$, flatness a few μm	NIST-traceable (MIL-STD-45662A)
Metal checkerboard (honeycomb aluminium)	25 mm cells, position $\pm 10 \mu\text{m}$, flatness $50 \mu\text{m}$	vendor certificate
25° angle block (E_θ)	nominal 25°, $\pm 20''$ tolerance	calibration certificate
Height block (α estimate)	gauge-grade height standard	manufacturer specification
Precision Z-stage (depth)	encoder res. / repeatability $0.5 \mu\text{m}$	closed-loop encoder

Table 5

Internal calibration accuracy: proposed RPC mapping vs. HALCON baseline. Lower is better for all metrics. For MAE_X and MAE_Z the entries report the mean over $N = 100$ samples with the sample standard deviation in parentheses; R_X and R_Z are point repeatability statistics for which no per-sample distribution is available.

Method	R_X (μm)	R_Z (μm)	MAE_X (μm)	MAE_Z (μm)
HALCON baseline	7.184	0.821	10.600 (2.092)	4.332 (0.263)
Proposed RPC	4.754	0.246	8.745 (1.367)	3.820 (0.067)
Improvement (%)	33.8	70.0	17.5	11.8

sampling for internal calibration, and the linear push-broom rig delivers multi-frame profile stitching for external evaluation. The target set (Fig. 7) contains a diffuse-reflectance grid distortion target (Edmund Optics) carrying a 49×49 chrome-dot array at 1 mm pitch and a honeycomb-aluminium metal checkerboard with an 11×7 grid of 25 mm cells (overall 375×275 mm); standard angle blocks supply ground-truth angles for the external geometric evaluation.

The reliability of the reference values used for error evaluation rests on traceable metrology rather than on self-consistency with the device under test; Table 4 summarises the certified accuracy and traceability of each standard. The internal dot array is a diffuse-reflectance grid distortion target carrying a serialised NIST-traceable certificate of accuracy (per MIL-STD-45662A) with a centre-to-centre dot-pitch accuracy of $\pm 1 \mu\text{m}$, a dot-size accuracy of $\pm 3 \mu\text{m}$, and a glass substrate flat to within a few micrometres over the array; the metal checkerboard is a honeycomb-aluminium panel certified to a ± 0.01 mm ($10 \mu\text{m}$) feature-position accuracy and a $50 \mu\text{m}$ flatness over the board; the 25° angle block (for the E_θ evaluation) is certified to a $\pm 20''$ (arc-second) nominal-angle tolerance, and the standard height block (for the α estimate) is a gauge-grade height standard with a micrometre-level tolerance; and the depth ground truth is provided by the closed-loop precision Z-stage, whose $0.5 \mu\text{m}$ encoder resolution and bidirectional repeatability bound the depth-axis truth. Crucially, each evaluation task compares the reconstructed geometry against an *independently* certified physical invariant—a known dihedral angle, a known 90° corner, a known dot pitch—rather than against the calibration output itself, so the metrics are free of the circularity that would arise if the system under test were used to define its own ground truth.

4.2. Internal RPC mapping vs. HALCON baseline

This experiment quantifies whether the proposed RPC mapping improves the internal pixel-to-physical correspondence over a main-stream baseline. Under identical acquisition conditions, we run the proposed method and the HALCON reference calibration on the same data: 20 groups of laser-stripe profile images plus a single-frame dot-target calibration image are collected, and stripe-centre and feature pixel

coordinates are extracted using a Steger-style sub-pixel ridge detector shared by both pipelines. Specifically, the Steger-style sub-pixel ridge detector was configured with a line filter width of 5 pixels and a transition pixel width of 7 pixels to adequately smooth the laser profile and suppress high-frequency noise. The smoothing contrast threshold was set to 40 grayscale levels to robustly separate the valid laser stripe from the background. Furthermore, to eliminate isolated noise points and repair disconnected stripe segments, we applied a point filter with a kernel size of 91 pixels, a morphological operation kernel of 3 pixels, and a point offset value of 9 pixels. On the hardware acquisition side, the 3D camera exposure time was fixed at 200 μs to ensure a stable signal-to-noise ratio during continuous scanning. We report two metrics in both X and Z : repeatability (range across repeats) and mean absolute error (MAE).

HALCON baseline configuration. The HALCON baseline is configured to be a fair, vendor-recommended reference for line-laser pixel-to-physical mapping rather than an intentionally weakened baseline. We use HALCON's planar dot-target operator chain (find_calib_object $\rightarrow$ calibrate_cameras) for camera intrinsic and lens-distortion identification with the area-scan polynomial distortion model (area_scan_polynomial) enabled, followed by the Sheet-of-Light calibration family (calibrate_sheet_of_light) for the light-plane-to-camera transform, and the resulting calibrated mapping is queried via image_points_to_world_plane (with fit_3d_plane_xyz as the underlying plane fit) to obtain $(u, v) \rightarrow (X, Z)$ correspondences. Both pipelines consume the identical multi-depth image stack used for the proposed RPC fit—the same six Z -positions Z_1 – Z_6 , the same grid distortion target, the same stripe-centre extractor, and the same train/test depth split (five depths used for fitting, one depth held out for testing in both pipelines)—and parameters were set according to the vendor-recommended configuration for line-laser scanning where applicable, rather than left at their generic defaults. The visible accuracy gap between HALCON and the proposed RPC is therefore attributable to the modelling choice—specifically, that HALCON enforces a single 3D plane through fit_3d_plane_xyz whereas the proposed RPC absorbs depth-varying nonlinearities via its bivariate cubic residual—rather than to baseline tuning. We avoid claiming that HALCON is incapable of multi-plane modelling in general; the contrast here is between the standard vendor-recommended single-plane configuration and the proposed cubic-residual RPC.

The repeatability is defined as the range of repeated measurements at a given test point,

$$R_X = \max_i(x_i) - \min_i(x_i), \quad R_Z = \max_i(z_i) - \min_i(z_i), \quad (21)$$

and the MAE measures the mean deviation between measured values and ground truth,

$$\text{MAE}_X = \frac{1}{N} \sum_{i=1}^N |x_i^{\text{meas}} - x_i^{\text{gt}}|, \quad \text{MAE}_Z = \frac{1}{N} \sum_{i=1}^N |z_i^{\text{meas}} - z_i^{\text{gt}}|. \quad (22)$$

Table 5 and Fig. 8 summarise the results. The HALCON baseline yields $R_X = 7.184 \mu\text{m}$, $R_Z = 0.821 \mu\text{m}$, $\text{MAE}_X = 10.600 \mu\text{m}$ ($\sigma = 2.092$), and $\text{MAE}_Z = 4.332 \mu\text{m}$ ($\sigma = 0.263$), while the proposed RPC mapping yields $R_X = 4.754 \mu\text{m}$, $R_Z = 0.246 \mu\text{m}$, $\text{MAE}_X = 8.745 \mu\text{m}$ ($\sigma = 1.367$), and $\text{MAE}_Z = 3.820 \mu\text{m}$ ($\sigma = 0.067$) over $N = 100$ samples—improvements of 33.8% and 70.0% in X/Z repeatability and 17.5% and 11.8% in X/Z MAE. A Welch two-sample t -test confirms these gains exceed the measurement noise: $t \approx 7.4$, Cohen's $d \approx 1.05$ for MAE_X and $t \approx 18.9$, $d \approx 2.67$ for MAE_Z (both $p < 0.001$), so even the smaller $0.512 \mu\text{m}$ MAE_Z gap is unambiguous. The pattern of gains is diagnostic: the largest improvement (70.0%) falls on R_Z , the metric most sensitive to height-direction non-linearity—exactly the deviation the cubic residual $r(u, v)$ is designed to absorb—and it is mirrored by a 74.5% tightening of σ_Z ($0.263 \rightarrow 0.067 \mu\text{m}$, versus 34.7% for σ_X). The much smaller MAE_Z gain reflects that the perspective backbone already removes most of the absolute height bias, leaving the residual to repair






Fig. 6. Experimental platform. The OPT-LPF2-006-01 line-laser profiler is mounted on a precision Z-axis motorised stage (used for multi-depth internal-calibration sampling). A precision Y-axis motorised stage drives the push-broom scanning of targets placed on the worktable, and a programmable motion controller synchronises triggering and image capture.



(a) Diffuse-reflectance grid distortion target (chrome dots on a diffusely-reflecting glass substrate), 49×49 array at 1 mm pitch, used to build pixel-physical correspondences for the internal RPC mapping.



(b) Metal checkerboard (honeycomb-aluminium substrate), 11×7 grid of 25 mm cells, used to estimate γ and to evaluate orthogonality.

Fig. 7. Calibration targets used in the two stages of the proposed framework.

variance rather than bias. This bias/variance decomposition matches the geometric-plus-data-driven structure of the RPC model and provides a reliable internal foundation for the external compensation in Section 4.3.

4.3. Depth-wise external shear compensation

This experiment tests whether a single external compensator suppresses installation-induced shear consistently across the working depth of field. Because the shear accumulates with stitching and varies with position, a single-depth test cannot establish cross-depth stability, so we deploy the six depth positions Z_1 – Z_6 spanning the lower-extreme, standard, and upper-extreme planes. At each depth we run three states:

Normal (machining/assembly tolerances only), *Tilted (uncompensated)*, and *Tilted (compensated)* (the same tilt corrected by Eq. (20)). The intentional tilt is held fixed throughout this experiment at a physical installation of $\alpha \approx 5^\circ$ (roll about X , which produces the vertical shear that distorts the angle block) and $\gamma \approx 10^\circ$ (yaw about Z , which produces the in-plane horizontal shear that distorts the checkerboard). The pitch β (about Y) is deliberately excluded: it rotates the laser plane about its own axial direction and therefore imprints no in-plane shear on the stitched cloud (Section 3.2), so it neither needs nor enters the compensation. Table 6 reports per-depth diagnostic re-estimates of the angles, acquired to evaluate the repeatability of the estimator; the compensation applied in Tables 7 and 8 uses the single frozen Z_3 estimate




Fig. 8. Internal calibration accuracy. (a) Absolute R_X , R_Z , MAE_X , MAE_Z for HALCON and the proposed RPC mapping (lower is better; values in μm). (b) Relative improvement of the proposed method over the baseline (higher is better). The largest improvement appears in R_Z (70.0%), the metric most sensitive to height-direction nonlinearity that the cubic residual term is designed to absorb.

Table 6

External installation angles estimated by the proposed method at each depth, for the intentional tilt state (physical $\alpha \approx 5^\circ$, $\gamma \approx 10^\circ$). The roll α (about X) causes the vertical shear measured by the 25° block; the yaw γ (about Z) causes the in-plane shear measured by the checkerboard. The pitch β (about Y) acts only along the laser-line axial direction, produces no in-plane shear, and is therefore not estimated or compensated.

Estimate	Z_1	Z_2	Z_3	Z_4	Z_5	Z_6	Mean
$\hat{\alpha}$ ($^\circ$)	4.981	5.043	5.150	4.938	4.921	5.113	5.024
$\hat{\gamma}$ ($^\circ$)	10.026	10.029	9.830	10.033	10.034	10.037	9.998

Table 7

Depth-wise angle measurement error of the 25° standard block (E_θ , in degrees). Per-cell entries report the mean over $N = 10$ trials with the trial-wise standard deviation in parentheses. The *Reduction* column is computed from the depth-mean (mean of trials). The bottom row reports the depth-mean of each column with the standard deviation across the six depth-means in parentheses, summarising depth-to-depth scatter.

Depth	Normal ($^\circ$)	Tilted, un-compensated ($^\circ$)	Tilted, compensated ($^\circ$)	Reduction (%)
Z_1	0.0409 (0.0027)	0.7684 (0.1122)	0.0388 (0.0039)	94.95
Z_2	0.0457 (0.0023)	0.8010 (0.0556)	0.0425 (0.0018)	94.69
Z_3	0.0347 (0.0019)	0.7068 (0.0839)	0.0368 (0.0014)	94.80
Z_4	0.0438 (0.0031)	0.8046 (0.0950)	0.0406 (0.0027)	94.95
Z_5	0.0396 (0.0019)	0.7521 (0.1115)	0.0379 (0.0021)	94.96
Z_6	0.0369 (0.0026)	0.7216 (0.1080)	0.0358 (0.0029)	95.04
Mean	0.0403 (0.0041)	0.7591 (0.0403)	0.0387 (0.0025)	94.90

($\hat{\alpha} = 5.150^\circ$, $\hat{\gamma} = 9.830^\circ$; Section 4). $\hat{\alpha}$ clusters tightly around 5° and $\hat{\gamma}$ around 10° across all six planes, with the per-depth dispersion bounding the angle-estimation repeatability. Each (depth, state) combination repeats two complementary geometric tasks $N = 10$ times; Tables 7 and 8 report the per-cell mean and trial-wise standard deviation.

4.3.1. Task 1: Angle measurement error on a 25° standard block

We use a 25° standard angle block (Fig. 9) to probe whether installation tilt distorts angular geometry in the reconstructed cloud. For each scan, we fit a plane to the inclined face and a plane to the bottom face and take the angle between them as the measured angle θ_{meas} . The error is

$$E_\theta = |\theta_{\text{meas}} - 25^\circ|. \quad (23)$$

Table 8

Depth-wise orthogonality error of the checkerboard point cloud (E_{orth} , in degrees). Per-cell entries report the mean over $N = 10$ trials with the trial-wise standard deviation in parentheses. The *Reduction* column is computed from the depth-mean (mean of trials). The bottom row reports the depth-mean of each column with the standard deviation across the six depth-means in parentheses, summarising depth-to-depth scatter.

Depth	Normal ($^\circ$)	Tilted, un-compensated ($^\circ$)	Tilted, compensated ($^\circ$)	Reduction (%)
Z_1	0.318 (0.027)	9.964 (1.066)	0.302 (0.033)	96.97
Z_2	0.341 (0.036)	10.432 (1.007)	0.321 (0.025)	96.92
Z_3	0.294 (0.025)	9.808 (1.060)	0.307 (0.051)	96.87
Z_4	0.327 (0.029)	10.106 (0.940)	0.316 (0.032)	96.87
Z_5	0.308 (0.026)	9.945 (1.027)	0.304 (0.039)	96.94
Z_6	0.298 (0.032)	9.786 (1.049)	0.301 (0.027)	96.93
Mean	0.314 (0.018)	10.007 (0.239)	0.309 (0.008)	96.92




Fig. 9. Task 1 setup: a 25° standard angle block placed on the metal checkerboard, scanned by the line-laser profiler in push-broom mode. The error E_θ is the absolute deviation of the dihedral angle between the fitted inclined plane and the fitted bottom plane from the nominal 25° .

Table 7 and Fig. 10 report E_θ at every depth and state. In the *Tilted (uncompensated)* state, the depth-mean angle error rises from 0.0403° in the normal state to 0.7591° , an $18.84\times$ amplification, which confirms that installation deviation directly destroys the angular geometry of

the reconstructed cloud. After applying the proposed external compensator, the depth-mean angle error drops to 0.0387° , a 94.90% reduction relative to the uncompensated tilted state and a level slightly better than the normal-installation reference. Across Z_1 – Z_6 , the depth-wise reduction stays in 94.69%–95.04% (a band narrower than 0.4 percentage points), indicating that the same compensator handles every working plane without per-depth retuning. The absolute compensated values track the normal-installation values almost point-by-point in Fig. 10(a), which suggests that the residual after compensation is dominated by sensor and stripe-extraction noise rather than by un-modelled external geometry.

4.3.2. Task 2: Orthogonality error on a standard checkerboard

We probe spatial orthogonality by scanning a metal checkerboard, fitting two pairs of mutually perpendicular edge lines from the reconstructed cloud (Fig. 11), and measuring the angle θ_{orth} between them. The orthogonality error is

$$E_{\text{orth}} = |\theta_{\text{orth}} - 90^\circ|. \quad (24)$$

Table 8 and Fig. 12 report E_{orth} across depths and states. In the *Tilted (uncompensated)* state, the depth-mean orthogonality error reaches 10.007° , $31.87\times$ above the normal-installation reference 0.314° , which makes installation-induced shear directly visible as severely non-orthogonal stitched geometry. After compensation, the depth-mean error falls to 0.309° , a 96.92% reduction relative to the uncompensated tilted state and again on par with normal installation. Across Z_1 – Z_6 , the depth-wise reduction stays in 96.87%–96.97% (a band narrower than 0.1 percentage points), showing that the proposed compensator restores global orthogonal consistency uniformly across the working depth of field. The orthogonality experiment amplifies the geometric distortion ($31.87\times$) more than the angle-block experiment ($18.84\times$) because E_{orth} accumulates the in-plane yaw-induced shear together with residual cross-axis coupling between the roll and yaw corrections in the stitched 3D geometry, while the dihedral E_θ only sees the projected angle along a single direction; the fact that the proposed compensator suppresses the larger amplification with a higher reduction percentage (96.92% vs. 94.90%) indicates that the physical roll–yaw decomposition and the Z – Y – X inverse rotation capture the dominant cross-axis coupling in the tested setup.

4.3.3. Discussion

Combining Tables 7 and 8, the proposed external compensator suppresses two complementary kinds of installation-induced geometric




Fig. 11. Task 2 visualisation: two mutually perpendicular edge lines (red overlay) extracted from the stitched checkerboard cloud. The angle between them is the measured θ_{orth} , and E_{orth} is its absolute deviation from 90° . This metric is far more sensitive than E_θ to multi-axis shear because the two lines lie in different parts of the scan, so any cross-axis distortion directly shifts the angle.

distortion (planar angle error and 3D orthogonality error) by 94.90% and 96.92% on average, with depth-wise variation under one percentage point (94.69%–95.04% for angle, 96.87%–96.97% for orthogonality). To verify that this reduction is statistically meaningful rather than a depth-mean artefact, we treat the six depth positions Z_1 – Z_6 as paired observations and run a paired t -test on the per-depth means between the *tilted, uncompensated* and *tilted, compensated* columns of Tables 7 and 8. For the angle task the mean reduction is 0.720° with paired $t(5) = 46.5$ and 95% confidence interval $[0.681^\circ, 0.760^\circ]$ ($p < 0.001$); for the orthogonality task the mean reduction is 9.698° with $t(5) = 102.7$ and CI $[9.46^\circ, 9.94^\circ]$ ($p < 0.001$). A non-parametric Wilcoxon signed-rank test is consistent (all six per-depth differences are positive, $W^+ = 21$, two-sided $p \approx 0.031$ at the $n = 6$ resolution floor). Depth uniformity is the central claim of Section 3.2: the same frozen α – γ estimate, applied through the depth-independent Z – Y – X formulation with $\beta = 0$ in Eq. (20), restores ideal-geometry-level accuracy at every working plane, which the experiments substantiate.




Fig. 10. Depth-wise angle measurement error of the 25° block. (a) Trajectories of E_θ across Z_1 – Z_6 for the three installation states (log scale). The compensated curve overlaps the normal-installation reference at every depth. (b) Depth-wise error reduction (%) achieved by the compensator; the band is tight (94.69%–95.04%, mean 94.90%), confirming depth uniformity.




Fig. 12. Depth-wise orthogonality error of the checkerboard cloud. (a) Trajectories of E_{orth} across Z_1 – Z_6 for the three installation states (log scale). The uncompensated curve sits over 30 times above the normal-installation reference, while the compensated curve overlaps it. (b) Depth-wise error reduction (%); the band (96.87%–96.97%, mean 96.92%) is narrower than 0.1 percentage points, again confirming depth uniformity.




Fig. 13. Geometric error amplification under installation tilt and recovery after compensation, expressed as multipliers of the normal-installation reference (log scale, depth-mean values). Installation tilt amplifies the angle and orthogonality errors by 18.84× and 31.87×, respectively; the proposed compensator pulls both metrics back to ~1×, i.e., within the noise floor of the normal-installation case.

Fig. 13 re-expresses Tables 7 and 8 as multipliers relative to the normal-installation reference. Tilt amplifies the angle and orthogonality errors to 18.84× and 31.87× of the normal level, respectively, while the proposed compensator pulls them back to 0.96× and 0.98×, i.e., back to within the noise floor of the normal-installation case. The fact that the two compensated multipliers fall slightly below unity is consistent with the compensator removing a small residual machining/assembly tolerance that is also present in the normal-installation reference, rather than overfitting; it should not be read as a claim of better-than-perfect calibration.

Together with the internal-stage gains in Table 5 and Fig. 8, this confirms that the two-stage decoupled framework simultaneously improves local mapping fidelity and restores global geometric consistency. The two stages also remain complementary rather than redundant: improving the internal RPC alone does not remove cross-axis shear, while applying the inverse rotation alone does not fix height-direction nonlinearity; the combined gains across Tables 5–8 show

that the proposed framework needs both stages to deliver micrometre-level local mapping accuracy and sub-degree depth-wise geometric consistency simultaneously on the tested device.

The quantitative gains are mirrored qualitatively in the reconstructed point clouds. Fig. 14 shows the stitched cloud of the checkerboard-plus-block scene at all six depths Z_1 – Z_6 , before and after external compensation. In the uncompensated rows the checkerboard grid is visibly sheared into a parallelogram and the cells are non-orthogonal, with the distortion persisting across every depth; after the depth-independent Z – Y – X inverse rotation, the grid is restored to an orthogonal lattice at every plane and the 25° angle block sits squarely on the recovered surface. The visual recovery is uniform across Z_1 – Z_6 , consistent with the single frozen compensator generalising across the working volume rather than requiring per-depth retuning.

Deploying the frozen compensator over long production runs raises two practical robustness questions that the static-misalignment model does not capture directly: thermal drift and platform vibration. Thermal expansion of the gantry and sensor bracket over a shift slowly changes the effective (α, β, γ) . Because the external estimation in Section 3.2 is a fast bottom-plane-plus-block scan that needs no manual posing, it can be re-run periodically as a lightweight in-process re-calibration: a drift threshold on the re-estimated angles—for instance a change exceeding the across-trial standard deviations reported in Tables 7 and 8—can trigger automatic re-compensation without stopping the line, and a fixed warm-up period before the first calibration removes the largest transient. High-frequency platform vibration, by contrast, injects approximately zero-mean jitter into individual stripe profiles rather than a systematic bias, so it is suppressed by the same RANSAC plane/line fits, by the N -trial averaging used throughout Section 4.3, and by encoder-synchronised triggering that ties each profile to a metered stage position; only vibration coherent with the scan motion would survive averaging and bias the geometry, which is mitigated mechanically rather than in calibration. A full thermal-cycling and vibration-spectrum characterisation on a production gantry is left to future deployment studies and is noted in the limitations of Section 5.

4.3.4. Behaviour beyond the 5° tolerance

The 5° value used above is the vendor's nominal mounting tolerance, not a hard validity bound. To probe the envelope directly, we hold the depth at Z_3 with the Tasks 1–2 protocol and sweep the physical roll through $\alpha \in \{5^\circ, 10^\circ, 20^\circ\}$ at a fixed yaw $\gamma \approx 10^\circ$ (the fixture's mechanical limit). Deepening α steepens the vertical shear, so the uncompensated




Fig. 14. Reconstructed point clouds of the checkerboard-plus-block scene at the six depths Z_1 – Z_6 , arranged in two blocks (Z_1 – Z_3 above, Z_4 – Z_6 below); within each block the top row is uncompensated and the bottom row is compensated by the external Z – Y – X shear compensation. Uncompensated, the checkerboard lattice is sheared into a non-orthogonal parallelogram at every depth; compensated, the lattice is restored to an orthogonal grid uniformly across the working volume. The red element is the 25° standard angle block.

E_θ climbs from 0.71° to 2.89° to 6.39° , while the uncompensated E_{orth} stays near 9.6° (driven by the fixed γ). Yet the single-pass closed-form compensator recovers the geometry at every level (Table 9, Fig. 15): E_θ is pulled back to 0.037° – 0.101° and E_{orth} to 0.16° – 0.59° —a 97–99% reduction that does not collapse even at $\alpha = 20^\circ$, four times the tolerance—and the recovered $\hat{\alpha}$ still tracks the physical roll to within 0.11° . In short, 5° marks where the estimators decouple to first order and the closed form is provably clean, but the method keeps removing the dominant shear well beyond it; the second-order terms analysed in Section 3.2.2 grow only gradually, with an optional joint iteration available for sub- 0.1° residuals at extreme tilt. The yaw was capped at 10° by the fixture, not by the method.

4.4. Ablation studies

This subsection isolates the contribution of each design choice introduced in Section 3.4.

Table 9

Compensator behaviour beyond the nominal 5° tolerance, at depth Z_3 . The physical roll α is swept to four times the nominal tolerance while the yaw is held near the fixture's 10° mechanical limit. $\hat{\alpha}$, $\hat{\gamma}$ are the estimated angles. E_θ (angle block) and E_{orth} (checkerboard) are reported before (uncompensated) and after (compensated) the single-pass closed-form correction.

Physical ($^\circ$)		Estimated ($^\circ$)		E_θ ($^\circ$)	E_{orth} ($^\circ$)
α	γ	$\hat{\alpha}$	$\hat{\gamma}$	before / after	before / after
5	10	5.150	9.830	0.7068 / 0.0368	9.808 / 0.307
10	10	10.060	9.835	2.8885 / 0.1006	9.610 / 0.158
20	10	19.895	9.266	6.3922 / 0.0492	9.568 / 0.593

4.4.1. Three-way component ablation

We isolate the contribution of each pipeline stage by running three configurations on the same depth-tiered tilted dataset that produced Tables 7 and 8: *internal-only* (proposed RPC mapping with the external




Fig. 15. Compensator behaviour as the physical roll α is swept to 20° (yaw fixed near the fixture limit $\gamma \approx 10^\circ$), at depth Z_3 (log scale). (a) Angle-block error E_θ : the uncompensated curve rises steeply with the deepening vertical shear, while the compensated curve stays near 0.05–0.10°. (b) Checkerboard orthogonality error E_{orth} : the uncompensated error stays near 9.6° (driven by the fixed γ), whereas the compensated error stays below 0.6°. The dashed line marks the nominal 5° tolerance; correction remains effective well beyond it.

Table 10

Three-way component ablation on the same depth-tiered tilted dataset used for Tables 7 and 8. Mean over Z_1 – Z_6 ; standard deviation across the six depths is shown in parentheses. *Internal-only*: proposed RPC mapping with the external compensator disabled (the resulting E_θ and E_{orth} therefore coincide with the uncompensated tilted state of Tables 7 and 8, because the RPC mapping does not alter cross-frame stitching geometry). *External-only*: HALCON's rigid single-plane baseline with the proposed external compensator applied on top. *Full*: both stages, i.e., the proposed pipeline. Bold marks the column minimum.

Configuration	MAE _x (μm)	MAE _z (μm)	E_θ (°)	E_{orth} (°)
Internal-only	8.847 (0.338)	3.827 (0.21)	0.7591 (0.0403)	10.007 (0.239)
External-only	10.600 (0.32)	4.333 (0.19)	0.0506 (0.0094)	0.519 (0.013)
Full (proposed)	8.83 (0.405)	3.820 (0.15)	0.0387 (0.0025)	0.309 (0.008)

compensator disabled, i.e., Eq. (20) skipped), *external-only* (HALCON's rigid single-plane baseline with the proposed external compensator applied on top), and *Full* (both stages, i.e., the proposed pipeline). For each configuration we measure mapping accuracy (MAE_x, MAE_z) and geometric consistency (E_θ on the 25° block, E_{orth} on the checkerboard) at the six depth positions Z_1 – Z_6 and report the depth-mean in Table 10.

The table separates two distinct failure modes. *Internal-only* ties *Full* on MAE_x and MAE_z to within 0.02 μm (well inside the ~0.4 μm depth-wise scatter), since both share the RPC mapping; but its E_θ (0.7591°) and E_{orth} (10.007°) sit one to two orders of magnitude above *Full*, exactly because the external compensator is disabled. Symmetrically, *external-only* drastically reduces E_θ and E_{orth} from their uncompensated extremes to levels comparable with *Full* (0.0506° and 0.519°). This improvement highlights the external compensator's effectiveness in removing global shear distortion. However, because this configuration relies on the HALCON calibration instead of the proposed RPC mapping, it cannot correct the underlying sub-pixel mapping residuals.

Consequently, its MAE_x (10.600 μm) and MAE_z (4.333 μm) regress back to the baseline level. This residual mapping bias propagates into the plane and edge fits, explaining why *external-only* fails to match the superior geometric consistency of the *Full* pipeline despite the use of the same inverse rotation. It is worth noting that the mapping accuracy of *Full* here (8.83 μm, 3.820 μm) represents the depth-mean across six working planes Z_1 – Z_6 . These values are highly consistent with the single-depth baseline results reported in Table 5 (8.745 μm, 3.820 μm), confirming that the proposed RPC mapping maintains its sub-pixel accuracy and robustness across the entire depth of field.

Across the four metrics, *Full* reaches or ties for the column minimum on every metric, while neither single-stage configuration is competitive

on more than two: *internal-only* matches *Full* only on the two mapping columns (within ± 0.02 μm noise), and *external-only* provides strong but suboptimal improvements only on the two geometric columns. The two stages are therefore complementary rather than redundant: the RPC mapping carries the mapping-accuracy column, the inverse rotation carries the geometric-consistency column, and removing either stage collapses its performance back to the unimproved baseline.

4.4.2. Rotation-order interchangeability

Section 3.2.2 shows analytically that the six orderings of the three elementary rotations are interchangeable for the reported geometry, and we confirm this numerically. With the pitch $\hat{\beta}$ excluded from the compensation (Section 4.3), the six inverse orderings applied to the angles estimated at Z_3 ($\hat{\alpha} = 5.150^\circ$, $\hat{\gamma} = 9.830^\circ$) are *not* identical as rotation matrices: they fall into two equivalence classes whose matrices differ by a single rigid rotation of $O(\hat{\alpha}\hat{\gamma}) = 0.88^\circ$ about Y . This offset does not affect the reported errors, because E_θ and E_{orth} are *intrinsic* angles—a dihedral angle and an edge-pair angle—invariant under any global rigid rotation of the cloud. Applying all six orderings to the same tilted cloud therefore yields the identical geometry, $E_\theta = 0.0368^\circ$ and $E_{\text{orth}} = 0.307^\circ$ for every ordering, agreeing to within 10^{-4}° —far below the trial-wise standard deviations in Tables 7 and 8. We retain Z – Y – X because it is the exact inverse of the physical mounting chain (Section 3.2.2), which keeps the decomposition physically interpretable.

4.4.3. Residual polynomial order

Section 3.1 argues on bias–variance grounds that the bivariate cubic residual is the lowest order that fits the third-order optical-edge distortion while remaining numerically stable. We verify this with a held-out order sweep: keeping the perspective backbone fixed, we replace the residual term with a quadratic, cubic, and quartic bivariate polynomial in turn, re-fit each on the same five training depths, and report MAE_x and MAE_z on both the training depths and the held-out depth Z_3 (Table 11). The quadratic residual underfits: its held-out errors (9.558 and 4.784 μm) are the worst of the three on both axes, because even-order monomials alone cannot represent the antisymmetric edge-distortion lobes. The quartic residual shows the textbook over-fitting signature: it attains the lowest *training* errors (8.124 and 3.419 μm), yet its held-out errors jump to 12.155 and 5.130 μm—a roughly 50% degradation from train to test on both axes—as the five extra monomials per axis inflate the Vandermonde condition number and begin to absorb stripe-extraction noise. The cubic residual sits at the bias–variance knee: its train and held-out errors are nearly identical (8.680 $\rightarrow$ 8.745 μm in X , 3.766 $\rightarrow$ 3.820 μm in Z), the smallest generalisation gap of the three, and its held-out values coincide with the single-depth figures of Table 5.

Table 11

Residual polynomial-order sweep. The perspective backbone is held fixed and only the bivariate residual order is varied; each model is fitted on the five training depths and evaluated on both the training depths and the held-out depth Z_3 . Lower is better. Bold marks the best held-out accuracy. The quartic model attains the lowest *training* error but over-fits (largest train→held-out gap), while the quadratic model underfits.

Residual order	MAE _X (μm)		MAE _Z (μm)	
	Train	Held-out	Train	Held-out
Quadratic	9.485	9.558	4.054	4.784
Cubic (this work)	8.680	8.745	3.766	3.820
Quartic	8.124	12.155	3.419	5.130

Table 12

Wall-clock time budget of one complete two-stage calibration. The two scanning durations are approximate values representative of the present setup that depend on the stage traverse speed and sampling density; the computation times are measured. The total is dominated by physical scanning; the numerical core of the method (RPC fit, angle estimation, and point-cloud compensation) adds only ≈ 2.3 s.

Stage	Step	Time (s)
Internal (Stage 1)	Multi-depth stripe acquisition, processing, and centre extraction	≈ 60
	RPC fit (22 coefficients, OLS)	0.005
External (Stage 2)	Checkerboard push-broom scan	≈ 600
	Installation-angle (α, γ) estimation	1.7
	Point-cloud Z - Y - X compensation	0.6
Total		≈ 662

The sweep therefore confirms the cubic choice empirically: it is the only order that is simultaneously accurate on the held-out plane and stable from train to test.

4.5. Calibration time budget

Table 12 breaks down the wall-clock time of one complete two-stage calibration on the experimental platform. The total is about 11 minutes, but it is dominated almost entirely by physical scanning rather than computation: the multi-depth stripe acquisition for the internal stage (≈ 60 s) and the checkerboard push-broom scan for the external stage (≈ 600 s) together account for over 99% of the budget, whereas the numerical core of the method—solving the 22 RPC coefficients (5 ms), estimating the installation angles (1.7 s), and applying the inverse Z - Y - X compensation to the point cloud (0.6 s)—adds only about 2.3 s in total. The internal RPC fit is a one-time factory step amortised over the sensor's service life, and the external angle estimation and compensation are fast enough to be re-run on-site as a lightweight in-process re-calibration (Section 4.3.4); the dominant scanning time can be reduced directly by a faster stage traverse or a coarser checkerboard sampling without altering the method.

5. Conclusion

We presented a depth-adaptive calibration method for line-laser 3D vision systems that decouples internal pixel-to-physical mapping from external installation compensation through a fully automated two-stage workflow. Stage one fits a 22-coefficient rational polynomial mapping that fuses a perspective skeleton with a bivariate cubic residual, which absorbs nonlinear deviations that single-plane homographies cannot represent. Stage two estimates the shear-producing installation angles (roll and yaw) independently and inverts them through a depth-independent Z - Y - X rotation, restoring shear-free geometry across the working depth of field with a single compensator.

On the tested commercial line-laser camera (OPT-LPF2-006-01), the method reduces the X - and Z -direction MAE by 17.5% and 11.8% relative to the HALCON baseline, suppresses the depth-mean angle error of a 25° block from 0.7591° to 0.0387° (94.90%), and the depth-mean checkerboard orthogonality error from 10.007° to 0.309° (96.92%). The depth-wise reduction stays within 94.69%–96.97% across six discrete depth positions, indicating that the proposed calibration strategy maintains consistent geometric correction across the tested depth range of this device.

The empirical claims above are scoped to this single tested system—one camera and one nominal working-distance range (with the installation tilt swept from the nominal 5° up to 20° , Section 4.3.4)—so the framework is best read as a depth-adaptive calibration pipeline validated on a representative industrial line-laser sensor rather than as a fully general result over arbitrary line-laser hardware. The following limitations indicate the natural extensions and the experiments needed to broaden this scope.

- *Single hardware platform.* The evaluation uses one OPT-LPF2-006-01 camera, one nominal working-distance range, the 25° standard block, and one checkerboard; the installation tilt is, however, swept across $\alpha \in \{5^\circ, 10^\circ, 20^\circ\}$ (Section 4.3.4). Cross-device validation on a second camera or a different optical configuration and additional standard angles (e.g., $15^\circ, 45^\circ$) are the natural next steps for separating the framework's contribution from device-specific characteristics.
- *Single-pass small-angle regime.* The affine-shear derivation in Section 3.2.1 and the rigid-rotation compensation in Section 3.2.2 decouple to first order within the nominal mounting tolerance $|\alpha|, |\beta|, |\gamma| \lesssim 5^\circ$. The tilt sweep in Section 4.3.4 shows the single-pass closed form still corrects effectively to four times this tolerance ($\alpha = 20^\circ$) and to the fixture's $\gamma = 10^\circ$ yaw limit; the residual refinement available beyond that regime is an optional joint iteration of the estimators, and pushing the yaw past 10° would require a fixture with a larger mechanical adjustment range.
- *Static-misalignment assumption.* The external compensator presumes that (α, β, γ) are stationary across a single calibration cycle and that the platform bottom plane is approximately level relative to the camera-mounting reference. Sensors mounted on long-stroke gantries with thermal drift may need lightweight in-process re-estimation, and systems with large platform tolerances should be paired with a coarse pre-levelling step so that residual platform tilt is not absorbed into the α and β estimates.
- *Limited cross-session statistical validation.* The internal-stage comparison reports per-sample Welch t -tests and Cohen's d (Section 4.2) and the external-stage results report paired t -tests and 95% confidence intervals over the six depth positions (Section 4.3); the additional check the paper does not yet report is the cross-session repeatability of the estimated (α, β, γ) from independent re-mountings of the sensor. Collecting and reporting this cross-session variability is a planned follow-up.

Beyond these extensions, adapting the framework to a different line-laser profiler is mechanical rather than structural, which speaks to its universality. The internal RPC mapping is dimensionless in form: porting it to another sensor requires only re-running the automated Z -stage acquisition over that sensor's own depth range and field of view and re-solving the same 22-coefficient linear system, with the pixel-normalisation bounds (u_0, v_0, u_s, v_s) and the working-volume centring updated to the new resolution; no model term changes. The external Z - Y - X compensator depends only on the mounting geometry rather than on the optics, so it transfers unchanged across the validated tilt range (Section 4.3.4). A profiler with a substantially larger depth range or stronger optical-edge distortion may call for more Z sampling planes or a higher residual order, but both are hyperparameters of the same

pipeline rather than a redesign, so the framework is expected to generalise across fields of view and depth ranges with only re-acquisition and re-fitting.

These directions are the natural next steps for a general industrial line-laser calibration framework built on the structure introduced here.

CRedit authorship contribution statement

Ling Cao: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Zidong Lin:** Writing – review & editing, Writing – original draft, Validation, Investigation, Data curation. **Yong Yang:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Quan Tang:** Writing – review & editing, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Zhipeng Gao:** Writing – review & editing, Visualization, Validation, Software, Methodology, Investigation, Data curation. **Daquan Feng:** Writing – review & editing, Supervision, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Wei Pan:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was funded by OPT Machine Vision under the Dongguan Key Research and Development Program (No. 20241200300122) and the Guangdong Provincial Key Research and Development Program (No. 2025B0101120001).

Data availability

Data will be made available upon request.

References

- [1] G. Sansoni, M. Trebeschi, F. Docchio, State-of-the-art and applications of 3D imaging sensors in industry, cultural heritage, medicine, and criminal investigation, *Sensors* 9 (1) (2009) 568–601, <https://doi.org/10.3390/s90100568>
- [2] J. Geng, Structured-light 3D surface imaging: a tutorial, *Adv. Opt. Photonics* 3 (2) (2011) 128–160, <https://doi.org/10.1364/AOP.3.000128>
- [3] F. Zhou, G. Zhang, Complete calibration of a structured light stripe vision sensor through planar target of unknown orientations, *Image Vis. Comput.* 23 (1) (2005) 59–67, <https://doi.org/10.1016/j.imavis.2004.07.006>
- [4] Z. Zhang, A flexible new technique for camera calibration, *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (11) (2000) 1330–1334, <https://doi.org/10.1109/34.888718>
- [5] J. Heikkilä, O. Silvén, A four-step camera calibration procedure with implicit image correction, in: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 1997, pp. 1106–1112, <https://doi.org/10.1109/CVPR.1997.609468>
- [6] P. Sturm, S. Maybank, On plane-based camera calibration: a general algorithm, singularities, applications, in: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 1999, pp. 432–437, <https://doi.org/10.1109/CVPR.1999.786974>
- [7] Q. Sun, Y. Hou, Q. Tan, G. Li, A flexible calibration method using the planar target with a square pattern for line-structured light vision system, *PLoS One* 9 (9) (2014) e106911, <https://doi.org/10.1371/journal.pone.0106911>
- [8] Z. Zhang, H. Zhang, Y. Wu, F. Zhang, Simple and precise calibration of the line-structured light vision system using a planar target, *J. Opt. Soc. Am. A* 40 (7) (2023) 1397–1408, <https://doi.org/10.1364/JOSAA.485907>
- [9] D. Li, Y. Wang, W. Ma, X. Liu, G. Ding, G. Zhang, J. Fang, Automatic hand-eye calibration method of welding robot based on linear structured light, *J. Robot. Mechatron.* 36 (2) (2024) 438–448, <https://doi.org/10.20965/jrm.2024.p0438>
- [10] H. Wang, T. Li, X. Xia, S. Xu, J. Liu, Robot self-calibration using line-structured light sensor, *Ind. Robot* 52 (6) (2025) 903–913, <https://doi.org/10.1108/IR-01-2025-0028>
- [11] R.Y. Tsai, R.K. Lenz, A new technique for fully autonomous and efficient 3D robotics hand/eye calibration, *IEEE Trans. Robot. Autom.* 5 (3) (1989) 345–358, <https://doi.org/10.1109/70.34770>
- [12] F.C. Park, B.J. Martin, Robot sensor calibration: solving $AX = XB$ on the euclidean group, *IEEE Trans. Robot. Autom.* 10 (5) (1994) 717–721, <https://doi.org/10.1109/70.326576>
- [13] X. Wang, J. Huang, H. Song, Simultaneous robot-world and hand-eye calibration based on a pair of dual equations, *Measurement* 181 (2021) 109623, <https://doi.org/10.1016/j.measurement.2021.109623>
- [14] J. Wu, M. Wang, Y. Jiang, B. Yi, R. Fan, M. Liu, Simultaneous hand-eye/robot-world/camera-IMU calibration, *IEEE/ASME Trans. Mechatron.* 27 (4) (2022) 2278–2289, <https://doi.org/10.1109/TMECH.2021.3103995>
- [15] X. Wang, W. Li, X. Gao, R. Jin, Z. Gao, S. Chen, L. Ning, W. Li, Y. Huang, Hand-eye parameter estimation and line-structured light scanning calibration within a unified framework, *Opt. Laser Technol.* 183 (2025) 112323, <https://doi.org/10.1016/j.optlastec.2024.112323>
- [16] S. Yang, L. Yang, G. Zhang, T. Wang, X. Yang, Modeling and calibration of the galvanometric laser scanning three-dimensional measurement system, *Nanomanufacturing Metrol.* 1 (3) (2018) 180–192, <https://doi.org/10.1007/s41871-018-0020-0>
- [17] X. Wu, B. Liu, P. Li, Y. Liu, A unified calibration method of 3D laser profile measurement with different laser-line lengths, *Robot. Intell. Autom.* 43 (1) (2023) 53–64, <https://doi.org/10.1108/RIA-06-2022-0159>
- [18] W.-W. Sheng, Precise measurement for line structure light vision sensor with large range, *Sci. Rep.* 13 (2023) 7234, <https://doi.org/10.1038/s41598-023-34428-w>
- [19] H. Gao, G. Xu, Z. Ma, A novel calibration method of line structured light plane using spatial geometry, *Sensors* 23 (13) (2023) 5929, <https://doi.org/10.3390/s23135929>
- [20] Y. Fu, X. Deng, G. Jiang, Y. Peng, K. Liu, Line structured light calibration method based on circular targets, *Opt. Eng.* 63 (5) (2024) 054107, <https://doi.org/10.1117/1.OE.63.5.054107>
- [21] D. Guo, J. Cui, Y. Chen, Y. Wu, Line structure light calibration based on a freely placed single cylindrical object, *Opt. Lasers Eng.* 178 (2024) 108236, <https://doi.org/10.1016/j.optlaseng.2024.108236>
- [22] T. Li, C. Liu, F. Duan, X. Fu, G. Niu, C. Liang, A. Chen, A light plane calibration method of line-structured light sensors based on unified extrinsic parameters estimation, *Opt. Lasers Eng.* 188 (2025) 108925, <https://doi.org/10.1016/j.optlaseng.2025.108925>
- [23] Y. Mu, Q. Mao, H. Xu, G. Wang, High-precision practical calibration method with virtual point constraint for line-structured light sensor, *Opt. Express* 33 (10) (2025) 20993–21015, <https://doi.org/10.1364/OE.562496>
- [24] C. Liu, F. Duan, J. Li, Y. Xu, S. Xing, A scanning direction calibration method of line-structured light three-dimensional sensors, *Chin. J. Lasers* 50 (5) (2023) 0504001, <https://doi.org/10.3788/CJL220791>
- [25] D. Cao, W. Liu, S. Liu, J. Chen, W. Liu, J. Ge, Z. Deng, Simultaneous calibration of hand-eye and kinematics for industrial robot using line-structured light sensor, *Measurement* 221 (2023) 113508, <https://doi.org/10.1016/j.measurement.2023.113508>
- [26] T. Li, F. Duan, C. Liu, X. Fu, A. Chen, C. Liang, J. Li, H. Zheng, G. Guo, Parameters estimation-based calibration method for rotational axis of line-structured light 3D reconstruction, *Opt. Express* 32 (1) (2024) 457–470, <https://doi.org/10.1364/OE.512299>
- [27] Z. Deng, Y. Ruan, F. Hao, T. Liu, Hand-eye calibration of line structured-light sensor by scanning and reconstruction of a free-placed standard cylindrical target, *Measurement* 229 (2024) 114487, <https://doi.org/10.1016/j.measurement.2024.114487>
- [28] L. Cao, W. Zhang, D. Feng, W. Pan, Automatic geometric calibration and unwarping of rotating line-scan imaging systems using elliptical distortion of circular patterns, *Opt. Lasers Eng.* 203 (2026) 109793, <https://doi.org/10.1016/j.optlaseng.2026.109793>
- [29] Z. Zhu, H. Liu, J. Zhang, Y. Zhou, Calibration method of line-structured light sensors based on a hinge-connected target with arbitrary pinch angles, *Appl. Opt.* 62 (7) (2023) 1695–1703, <https://doi.org/10.1364/AO.483595>
- [30] Q. Wu, J. Qiu, Z. Li, J. Liu, B. Wang, Hand-eye calibration method of line structured light vision sensor robot based on planar target, *Laser Optoelectron. Prog.* 60 (10) (2023) 1015002, <https://doi.org/10.3788/LOP220852>
- [31] J. Lin, Y. Feng, W. Ren, J. Feng, J. Zheng, Position-constrained calibration compensation for hand-eye calibration in industrial robots, *Sensors* 24 (23) (2024) 7554, <https://doi.org/10.3390/s24237554>
- [32] Z. Qiu, J. Martínez-Sánchez, P. Arias-Sánchez, R. Rashdi, External multi-modal imaging sensor calibration for sensor fusion: a review, *Inf. Fusion* 97 (2023) 101806, <https://doi.org/10.1016/j.inffus.2023.101806>
- [33] R. Wang, J. Bai, Y. Xue, The calibration method of line structured light sensor for integrated position and pose detection of highway guardrail inspection robots, *Wuhan Univ. J. Nat. Sci.* 30 (4) (2025) 367–378, <https://doi.org/10.1051/wujns/2025304367>
- [34] J. Sun, H. Zhang, L. Lu, Z. Su, W. Pan, Self-calibrating structured light three-dimensional imaging method with fringe projection, in: *International Conference on Optical Communication, Signal Processing, and Optical Engineering (OCSPOE 2026)*, vol. 14185, SPIE, 2026, pp. 268–275, <https://doi.org/10.1117/12.3114025>