photonics

MDPI

Article

Paired-Yaw Identifiability and Finite Yaw-Set Design for Coupled Abbe and Beam-Pointing Calibration in Four-Channel Laser Interferometry: A Numerical Study

Chunzheng Wang and Hao Du *

School of Optoelectronic Engineering, Guilin University of Electronic Technology, Guilin 541004, China; 15863553898@163.com

* Correspondence: 1244452759@mails.guet.edu.cn

Abstract

Geometric calibration of a four-channel plane-mirror interferometer requires separating translation–yaw coupling from beam-pointing and mirror-orientation effects. This numerical study examines paired ranges acquired at matched positions with nonzero and zero yaw. A scaled analytic Jacobian identifies weak parameter combinations, and an equal-budget search compares 70 four-level yaw sets. The eight-parameter model has full local numerical rank but a condition number of 9.42×10^5. Fixing two mirror angles at their synthetic truth reduces the condition number to 2519. The selected yaw levels are $\{-2.4, -1.8, 1.8, 2.4\}$ mrad. Four solver configurations are compared over 100 synthetic recalibrations. Their similar held-out residuals conceal differences in parameter recovery: an additive decomposition assigns 3.67 nm RMS to yaw-reference mismatch within a 3.70 nm total residual. Controlled noise/yaw switches further separate calibration-induced parameter error from validation-side mismatch. These results identify survey and angular-reference requirements for the declared geometry and provide a local numerical design screen; physical accuracy remains to be established experimentally.

Keywords: laser interferometry; Abbe error; cosine error; parameter identifiability; Fisher information; trajectory design; nonlinear least squares; synthetic calibration

1. Introduction

Precision motion stages require simultaneous translation and angular sensing because rotation about an offset measurement line changes the inferred displacement [1–4]. Plane-mirror, grating, fibre-coupled, and multi-beam interferometers address this need with different optical observables [5–11]. The Abbe contribution scales approximately as the transverse offset multiplied by the rotation angle [12]. For example, a 9.2 mm offset and 2.4 mrad yaw give a first-order scale of 22.1 µm. This scale motivates geometric calibration, but does not establish that geometry dominates the residual error after correction. Cyclic nonlinearity, air-path scale, and optical aberrations can then become decisive [13–17].

Geometric error models already account for Abbe offsets, beam alignment, coordinate transformations, and mirror installation [18–21]. Redundant interferometry and reversal methods estimate quantities that are difficult to survey directly [22–24]. In particular, Zhang et al. use an overdetermined interferometer arrangement for geometric calibration [22]; Ye et al. address coupled errors in a wafer-stage grating-interferometer model [25]; and Wen et al. analyze nonlinear installation errors and parameter uniqueness [26]. Recent skew-ray modeling further emphasizes the dependence of the inverse problem on the optical

Check for updates

Received: 7 September 2026
Revised: 2 November 2026
Accepted: 19 November 2026
Published: 26 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

layout [27]. These studies establish the relevant modeling and calibration approaches. The question here is narrower: Which geometric combinations remain weakly separated in a four-channel paired-yaw observable, and how much can finite angular excitation improve that separation?

Local rank, scaled singular values, and parameter correlation are established identifiability tools [28]. Information-based experiment design is also well established [29–31]. Likewise, Levenberg–Marquardt (L–M) and bounded trust-region methods are standard inverse solvers [32–35]. We use these tools to distinguish three questions: whether the declared parameters are locally distinguishable, which available yaw levels improve the weakest directions, and whether the validation residual can reveal differences between fitted geometries. A small residual alone answers none of these questions reliably when pose-reference error is present.

The contributions are model-specific findings rather than new generic optical or optimization algorithms. First, an exact paired-yaw model exposes near-collinearity between Abbe arms and mirror angles despite full local numerical rank; fixing two mirror angles makes the assumed survey information explicit. Second, an equal-budget search over 70 yaw sets quantifies the information available within a prescribed motion grid. Third, additive residual decomposition and controlled noise/yaw switches show why nearly identical validation residuals can conceal different parameter errors. A five-layout screen tests the persistence of the rank–conditioning distinction within this architecture. All results are numerical. The study provides a conditional calibration-design analysis, not evidence that yaw coupling dominates every physical interferometer.

A companion FTCS 2026 conference manuscript, "Error-Aware Four-Channel Differential Laser Interferometric Sensing for Three-DOF Precision Motion Stages," examines pose reconstruction when geometry is known. This article instead estimates unknown geometry and studies its identifiability, yaw-set selection, and held-out recovery. None of the companion manuscript's principal numerical data, figures, tables, or model-order results are reused. The companion is disclosed as related work without asserting a proceedings publication.

2. Materials and Methods

2.1. Evidence Boundary and Declared Topology

Figure 1 shows the idealized topology and analysis workflow. The reference frame is right handed, with in-plane stage pose

$$\boldsymbol{q} = \begin{bmatrix} X & Y & \theta \end{bmatrix}^{\mathsf{T}}, \qquad R(\theta) = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}, \tag{1}$$

where positive θ is counter-clockwise. The four channels are ordered (X_1, X_2, Y_1, Y_2). The two X-family rays nominally propagate along $+X$, and the two Y-family rays along $+Y$. Fixed channel origins and small differential direction trims are declared constants. Each stage-mounted mirror family is represented as an infinite line in the two-dimensional inverse model. This idealization excludes return-beam walk-off, polarization leakage, diffraction, beam-shape effects, and out-of-plane stage motion.

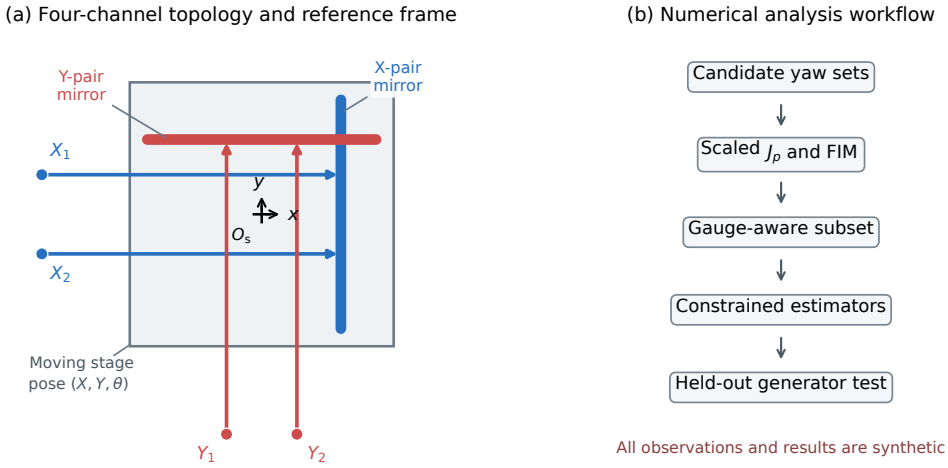

Figure 1. Declared theoretical system and workflow. (**a**) Blue and red denote the X- and Y-channel pairs, respectively; filled circles mark the fixed ray origins, arrowheads show nominal propagation, and O_s is the stage pose datum. The labels X_1, X_2, Y_1, Y_2 identify channels rather than measured experimental signals. (**b**) Candidate yaw sets are evaluated with a scaled Jacobian and Fisher information, followed by gauge-aware subset selection, nonlinear estimation, and a held-out synthetic generator test. Every observation and result in the workflow is synthetic.

2.2. Exact Paired-Yaw Observation Model

For axis family $j \in \{X, Y\}$, define the nominal axis unit vectors $\boldsymbol{e}_X = (1,0)^\mathsf{T}$ and $\boldsymbol{e}_Y = (0,1)^\mathsf{T}$. The geometric parameter block is

$$\boldsymbol{p}_j = \begin{bmatrix} a_j & c_j & \alpha_j & \beta_j \end{bmatrix}^\mathsf{T}, \tag{2}$$

where a_j is a signed transverse Abbe arm, c_j is a signed longitudinal mirror offset, α_j is a mirror-normal installation angle, and β_j is a common pointing angle for that channel family. Channel-specific trims δ_h are fixed. With fixed base origin $\bar{\boldsymbol{o}}_h$, the ray origin is

$$\boldsymbol{o}_h = \bar{\boldsymbol{o}}_h + a_j \boldsymbol{t}_j, \quad \boldsymbol{t}_X = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad \boldsymbol{t}_Y = \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{3}$$

The outgoing direction and moving mirror normal are

$$\boldsymbol{u}_h = R(\beta_j + \delta_h)\boldsymbol{e}_j, \qquad \boldsymbol{n}_j(\theta) = R(\theta + \alpha_j)\boldsymbol{e}_j. \tag{4}$$

The point defining the mirror line at pose $\boldsymbol{q}$ is

$$\boldsymbol{m}_j(\boldsymbol{q}) = \begin{bmatrix} X \\ Y \end{bmatrix} + R(\theta) c_j \boldsymbol{e}_j. \tag{5}$$

For ray $\boldsymbol{r}_h(s) = \boldsymbol{o}_h + s\boldsymbol{u}_h$, direct ray–line intersection gives the one-way range

$$\tau_h(\boldsymbol{q}, \boldsymbol{p}_j) = \frac{\boldsymbol{n}_j^\mathsf{T}(\boldsymbol{m}_j - \boldsymbol{o}_h)}{\boldsymbol{n}_j^\mathsf{T}\boldsymbol{u}_h}, \qquad |\boldsymbol{n}_j^\mathsf{T}\boldsymbol{u}_h| > 0.25. \tag{6}$$

The executable observation is not a raw phase record. It is a synthetic paired-yaw difference between two poses having the same (X, Y) and nonzero and zero yaw, respectively:

$$\Delta\tau_h(X, Y, \theta; \boldsymbol{p}_j) = \tau_h(X, Y, \theta; \boldsymbol{p}_j) - \tau_h(X, Y, 0; \boldsymbol{p}_j). \tag{7}$$

This construction represents two matched simulated runs and cancels a constant range datum. It also makes every zero-yaw row identically zero; the resulting zero-yaw rank statement is therefore specific to Equation (7) and is not a general assertion about raw interferometer signals. If λ_{air} is prescribed, the ideal double-pass phase representation is

$$\Delta\phi_h = \frac{4\pi}{\lambda_{\mathrm{air}}}\Delta\tau_h. \tag{8}$$

The saved datasets contain both phase and one-way-equivalent range columns. Phase is a redundant double-pass encoding retained for a unit audit; estimation reads the stored one-way-equivalent range columns directly and does not perform a phase-decoding step. Raw phase zeros, time alignment between physical runs, wavelength-scale estimation, and drift estimation are outside this inverse model.

2.3. Parameter Vector, Scales, Bounds, and Gauge Choices

The canonical geometry order is

$$\boldsymbol{p} = \begin{bmatrix} a_X & c_X & \alpha_X & \beta_X & a_Y & c_Y & \alpha_Y & \beta_Y \end{bmatrix}^{\mathsf{T}}. \tag{9}$$

The full local diagnostic is denoted P8. Longitudinal ray-origin coordinates are not estimated because translating an origin along its own ray changes an absolute range but cancels in Equation (7). Tangential motion of a point on an infinite mirror line is likewise a representation gauge, so only c_j is retained. The reference-frame orientation, channel ordering, and signs of the fixed origins are frozen before generation.

Preliminary singular-vector inspection showed that a_X and α_X, and separately a_Y and α_Y, form nearly collinear directions over the finite operating region. Therefore, P8 is retained as a confounding diagnostic but is not used for primary recovery. The fitted P6 subset is

$$\boldsymbol{p}_6 = \begin{bmatrix} a_X & c_X & \beta_X & a_Y & c_Y & \beta_Y \end{bmatrix}^{\mathsf{T}}, \tag{10}$$

with α_X and α_Y synthetically fixed at their scenario truth. Thus, the improved P6 information is partly supplied by a hypothetical exact survey; it is not attributed solely to the selected motion. A separate $+3$ µrad survey-offset scenario tests this dependence.

Table 1 lists the scenario values, parameter status, and bounds. These values are illustrative numerical inputs, not surveyed properties of an apparatus. Optimization uses diagonal scales $(10, 25, 50, 50, 10, 25, 50, 50)$ in units of mm or µrad matching the corresponding parameter type.

Table 1. Declared geometry parameters. Mirror angles are fixed exactly at their synthetic truth in the baseline P6 comparison. No listed value is a physical measurement.

Parameter	Role	Unit	Status	Scenario	Bounds
a_X	Estimated Abbe arm	mm	Estimated	9.2	$[-20, 20]$
c_X	Estimated longitudinal offset	mm	Estimated	22	[5, 45]
α_X	Survey-fixed mirror angle	µrad	Survey-fixed	34	$[-120, 120]$
β_X	Estimated common beam pointing	µrad	Estimated	-24	$[-120, 120]$
a_Y	Estimated Abbe arm	mm	Estimated	-8.4	$[-20, 20]$
c_Y	Estimated longitudinal offset	mm	Estimated	18.5	[5, 45]
α_Y	Survey-fixed mirror angle	µrad	Survey-fixed	-29	$[-120, 120]$
β_Y	Estimated common beam pointing	µrad	Estimated	19	$[-120, 120]$

2.4. Analytic Jacobian and Dimensionless Information

Let $A_h = \boldsymbol{n}_j^{\mathsf{T}}(\boldsymbol{m}_j - \boldsymbol{o}_h)$ and $D_h = \boldsymbol{n}_j^{\mathsf{T}}\boldsymbol{u}_h$. For any scalar parameter ξ, differentiation of Equation (6) gives

$$\frac{\partial \tau_h}{\partial \xi} = \frac{\left[(\partial_\xi \boldsymbol{n}_j)^{\mathsf{T}}(\boldsymbol{m}_j - \boldsymbol{o}_h) + \boldsymbol{n}_j^{\mathsf{T}}\partial_\xi(\boldsymbol{m}_j - \boldsymbol{o}_h)\right] D_h - A_h\left[(\partial_\xi \boldsymbol{n}_j)^{\mathsf{T}}\boldsymbol{u}_h + \boldsymbol{n}_j^{\mathsf{T}}\partial_\xi \boldsymbol{u}_h\right]}{D_h^2}. \tag{11}$$

The required elementary derivatives are

$$\begin{array}{c|ccc} \xi & \partial_\xi \boldsymbol{n}_j & \partial_\xi \boldsymbol{u}_h & \partial_\xi(\boldsymbol{m}_j - \boldsymbol{o}_h) \\ \hline a_j & \mathbf{0} & \mathbf{0} & -\boldsymbol{t}_j \\ c_j & \mathbf{0} & \mathbf{0} & R(\theta)\boldsymbol{e}_j \\ \alpha_j & G\boldsymbol{n}_j & \mathbf{0} & \mathbf{0} \\ \beta_j & \mathbf{0} & G\boldsymbol{u}_h & \mathbf{0} \end{array}, \quad G = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \tag{12}$$

The paired-yaw Jacobian is the difference of Equation (11) evaluated at θ and at zero yaw.

The declared covariance of each already-differenced four-channel range row is

$$\Sigma_r = \sigma_r^2 C, \qquad \sigma_r = 0.30 \text{ nm}, \tag{13}$$

where C has unit diagonal, within-pair correlation 0.35, and cross-pair correlation 0.08. This is the covariance of Equation (7), not of either raw acquisition. If $LL^{\mathsf{T}} = \Sigma_r$, $H = \partial \Delta\boldsymbol{\tau}/\partial \boldsymbol{p}$, and D_p is the diagonal parameter-scale matrix, the scaled whitened Jacobian and information matrix are

$$J = L^{-1} H D_p, \qquad F = J^{\mathsf{T}} J. \tag{14}$$

Every column and singular value of J is dimensionless. Rank uses the numerical threshold

$$\tau_{\text{svd}} = \sigma_{\max}(J)\max(N_{\text{obs}}, N_p)\epsilon_{\text{mach}}. \tag{15}$$

For full rank, $\kappa_2 = \sigma_{\max}/\sigma_{\min}$, and the local correlation matrix is obtained from F^{-1}. Exact numerical rank is reported separately from practical conditioning. A maximum $|\rho|$ close to one is treated as evidence of a narrow parameter valley even when Equation (15) declares full rank.

The analytic Jacobian was checked against complex-step and five-point differences. A separately coded three-coordinate ray–plane generator was also compared with the inverse formula after disabling all generator departures. The generator uses three-component vectors but the executed stage motion remains planar; this check reduces shared-code error risk without constituting an independent physical model or an out-of-plane validation.

2.5. Finite Equal-Budget Yaw-Set Enumeration

The spatial calibration path is a bidirectional raster spanning ± 140 mm in both axes. Eight nonzero yaw candidates are defined:

$$\Theta_c = \{-2.4, -1.8, -1.2, -0.6, 0.6, 1.2, 1.8, 2.4\} \text{ mrad}. \tag{16}$$

Every four-level subset receives the same budget of $N_{\text{cal}} = 1200$ positions, divided as evenly as possible among its yaw levels. The finite search therefore evaluates $\binom{8}{4} = 70$ sets; it does not optimize the raster, dwell time, velocity, or continuous yaw.

Robustness to geometry selection is evaluated with 24 parameter draws. Each draw adds independent Gaussian perturbations with a standard deviation 0.30 times the declared parameter scale to the design reference, and clips the result inside the physical bounds. For candidate $\mathcal{T}$, the recorded quantities are the worst rank, worst $\log\det F$, worst $\sigma_{\min}$, largest condition number, and largest parameter correlation across those draws. The predeclared compromise score is

$$S(\mathcal{T}) = \mathcal{N}\Big[\min_{\omega} \log\det F\Big] + \mathcal{N}\Big[\log\min_{\omega} \sigma_{\min}\Big] - \mathcal{N}\Big[\log\max_{\omega} \kappa_2\Big], \tag{17}$$

where $\mathcal{N}$ maps each finite candidate metric linearly to $[0, 1]$. The selected design maximizes S after rejecting rank-deficient candidates. This is a finite-grid compromise, not a proof of global D-, E-, or minimax optimality.

2.6. Synthetic Generator, Data Split, and Scenario Inputs

The inverse uses the planar analytic model. The generator independently codes the direct three-coordinate ray–plane intersection and then introduces fixed channel-level departures that are absent from the inverse: origin offsets of at most 0.23 µm, pointing-trim offsets of at most 0.75 µrad, and channel-dependent quadratic mirror-form coefficients of magnitude 4.5–7.0 $\times 10^{-8}$ m^{-1}. For a tangent hit coordinate s_h, the added one-way form term is $k_h s_h^2/2$.

Each yaw block receives a common Gaussian reference jitter plus a smaller samplewise component. A prescribed sinusoidal air-wavelength record encodes the redundant stored phase representation; it is neither decoded by the estimator nor estimated. Channel drift is a sinusoidal equivalent-range perturbation. Correlated Gaussian noise and sparse Laplace contamination are then added. The hidden true yaw and individual generator components are retained in the CSV files for audit, but fitting reads only commanded (X, Y, θ) and the four observed range columns. Table 2 summarizes the frozen protocol.

Table 2. Frozen numerical protocol. All magnitudes are assumptions of the synthetic generator.

Quantity	Baseline Value	Interpretation
Vacuum wavelength/refractive index	632.8 nm/1.0002718	Prescribed phase conversion
Paired-range standard deviation	0.30 nm	After differencing in Equation (7)
Within-pair/cross-pair correlation	0.35/0.08	Four-channel covariance
Block/local yaw jitter	25/2.5 nrad	Hidden generator yaw departure
Drift amplitude	0.18 nm	Channel-dependent sinusoid
Contamination	0.25% at 3.0 nm scale	Independent Laplace samples
Design/main calibration/validation rows	1200/1000	Raster/Lissajous
Rows per algorithm-MC/stress fit	600/500	Separately generated calibration rows
Profile rows	400	Deterministic one-in-three main-calibration subset
Algorithm/stress repetitions	100/20	Independent synthetic repetition sets
Primary random seed	20,260,821	Named derivative seeds used for substudies

The calibration yaw levels are selected only from Equation (16). The held-out set uses $(-2.1, -0.9, 0.9, 2.1)$ mrad and a Lissajous spatial path with 1000 positions. These yaw values do not intersect the selected calibration set, and the validation file is generated with

a separate seed. The same frozen validation realization is used for paired comparison of all calibration replicates. Consequently, variation across solver boxes describes calibration-estimate variability conditional on that realization; it is not a full deployment distribution.

2.7. Solver Configurations and Recovery Metrics

Four complete solver configurations are compared:

1. Scaled Gauss–Newton (GN) with backtracking and no bounds;
2. Standard unbounded L–M with quadratic loss;
3. Bounded trust-region reflective least squares (TRF) with quadratic loss;
4. CW–RLM, a robust multistart L–M configuration with a smooth bound map, Huber iteratively reweighted least squares (IRLS), and two seeded starts per Monte Carlo fit.

The first three use one common perturbed start in each replicate; CW–RLM uses the design-reference start plus one seeded interior start. This is a comparison of complete configurations, not an ablation. Differences cannot be assigned uniquely to damping, bounds, robust loss, or multistart.

GN stops when the maximum scaled step is below 10^{-8} or after 60 iterations. L–M and TRF use SciPy's `least_squares` with analytic derivatives and `ftol` = `xtol` = `gtol` = 10^{-12}. CW–RLM uses Huber threshold 1.5 in whitened-residual units, at most 10 IRLS outer iterations, 150 function evaluations per inner solve, and convergence when the scaled physical step and weight change are simultaneously small or the relative robust-objective change is below 2×10^{-6}. The number of evaluations reported for CW–RLM is summed over starts and IRLS cycles and therefore is not directly equivalent to one single-start solve.

For estimated parameter p_i with scale d_i, the normalized error is

$$e_i = \frac{\widehat{p}_i - p_i^{\mathrm{true}}}{d_i}, \qquad \|\boldsymbol{e}_p\|_2 = \left(\sum_{i \in \mathrm{P6}} e_i^2 \right)^{1/2}. \tag{18}$$

The held-out pooled residual is

$$\mathrm{RMSE}_{\mathrm{val}} = \left[\frac{1}{4N_{\mathrm{val}}} \sum_{h=1}^{4} \sum_{k=1}^{N_{\mathrm{val}}} \left(\Delta\tau_{hk}^{\mathrm{obs}} - \Delta\tau_h(\boldsymbol{q}_k; \widehat{\boldsymbol{p}}) \right)^2 \right]^{1/2}. \tag{19}$$

One hundred calibration replicates compare the finite-solution rate, convergence flag, median held-out RMSE, normalized parameter error, and evaluations. Central empirical 95% spans describe the repeated-fit distribution; local Wald inclusion proportions are archived for transparency but are not treated as calibrated confidence coverage because only 100 replicates were used, and the covariance does not implement a robust sandwich or bootstrap.

Two-dimensional profile diagnostics are computed for (a_X, β_X) and (a_Y, β_Y). At each grid point, the other four P6 parameters are re-optimized under bounds on a deterministic one-in-three calibration subset. The displayed $\Delta\chi^2$ surfaces are local quadratic-objective diagnostics under the declared covariance, not physical confidence regions.

2.8. Stress Matrix and Reproducibility

Eight exploratory stress rows vary one calibration-side setting at a time: half and double readout-noise magnitude while retaining the baseline whitening matrix, zero and 75 nrad block yaw jitter, triple mirror-form mismatch, 1% contamination, and a common $+3\,\mu$rad offset in the two fixed mirror angles. Each row has 20 fits on 500 calibration rows and is scored on the same baseline validation realization. To keep this screening matrix finite, its CW–RLM cap is eight IRLS cycles and 120 evaluations per inner solve,

compared with ten and 150 in the primary comparison. Thus, the noise rows combine a generator-noise change with deliberate covariance/loss-scale misspecification, and the table does not represent stressed deployment data. With 20 trials, percentages are coarse descriptive frequencies.

All calculations use Python 3.13.5, NumPy 2.1.3, SciPy 1.15.3, pandas 2.2.3, and Matplotlib 3.10.0 on Windows 11 [36,37]. The executable study writes raw synthetic rows, all summary CSV files, generated LaTeX tables, vector/raster figures, seeds, software versions, and a SHA-256 manifest. Analytic derivatives are checked at randomly generated poses. Failed optimizer calls remain in denominators rather than being silently dropped. The author-held archive also retains the added controlled-test inputs, all fitted parameter vectors, component cross-moments, and scripts linking numerical outputs to manuscript tables. Appendix A specifies the additional protocol; access status is stated under Data Availability. OpenAI Codex was used to assist with study-design review, code review, visualization, literature organization, and manuscript drafting. The authors retain responsibility for the equations, citations, scripts, outputs, interpretations, disclosure, and submitted work.

2.9. Separating Yaw, Readout, and Parameter-Estimation Contributions

The frozen held-out observation is decomposed without refitting. Let $f(\boldsymbol{p}, \theta)$ denote the inverse paired-range operator at the held-out positions, θ_a the hidden actual yaw, θ_c the commanded yaw, and $g(\boldsymbol{p}^*, \theta_a)$ the noiseless generator, including channel and form departures. Here, $\boldsymbol{p}^*$ and $\widehat{\boldsymbol{p}}$ are the scenario truth and the fitted CW–RLM geometry. Then,

$$\begin{aligned} \boldsymbol{r} &= \boldsymbol{e}_p + \boldsymbol{e}_\theta + \boldsymbol{e}_g + \boldsymbol{\epsilon} + \boldsymbol{d} + \boldsymbol{o}, \\ \boldsymbol{e}_p &= f(\boldsymbol{p}^*, \theta_c) - f(\widehat{\boldsymbol{p}}, \theta_c), \\ \boldsymbol{e}_\theta &= f(\boldsymbol{p}^*, \theta_a) - f(\boldsymbol{p}^*, \theta_c), \\ \boldsymbol{e}_g &= g(\boldsymbol{p}^*, \theta_a) - f(\boldsymbol{p}^*, \theta_a), \end{aligned} \tag{20}$$

where $\boldsymbol{\epsilon}$, $\boldsymbol{d}$, and $\boldsymbol{o}$ are readout noise, drift, and contamination. This telescoping decomposition is additive at the sample level. Component RMS values are not additive, because cross-moments may be nonzero. Subtracting $\boldsymbol{e}_\theta$ is an oracle diagnostic with fitted parameters held fixed, not a deployable yaw correction.

To isolate the effect on parameter estimation, a second test switches calibration readout noise and calibration yaw-reference error independently in a 2×2 design. Each setting has 100 fits using paired random draws. The validation operator has neither noise nor yaw error, so its residual measures prediction error caused by the estimated geometry alone. Other generator departures are disabled in this test. Appendix A specifies the solver, sample counts, and random seeds.

3. Results

3.1. Derivative and Cross-Implementation Verification

The analytic paired-range Jacobian agreed with complex-step differentiation to a maximum relative discrepancy of 1.54×10^{-13}, and the double-precision five-point check differed by 8.99×10^{-7}; both were below their predeclared tolerances. After channel departures and mirror form were disabled, the separately coded generator and inverse agreed over 128 random poses, with a maximum absolute difference of 1.11×10^{-16} m and an RMS difference of 2.46×10^{-17} m. The covariance in Equation (13) was positive definite. These checks verify implementation agreement and units; they do not validate either model against physical optics or hardware.

3.2. Rank, Conditioning, and Subset Choice

Figure 2 shows the scaled singular spectra and the effect of fixing the two mirror angles. The zero-yaw design has rank zero because Equation (7) vanishes by construction. A single +1.2 mrad level reaches P6 rank 6 but has condition number 5.91×10^4 and a maximum correlation of 0.999. Adding signed yaw levels substantially increases the weakest singular value.

For the selected four-yaw design, P8 has exact numerical rank 8 but $\kappa_2 = 9.421 \times 10^5$ and there is maximum local correlation of 0.9994. Figure 3a shows near-perfect a_X–α_X and a_Y–α_Y coupling. This is a local numerical rank result, not a proof of global structural identifiability or stable eight-parameter recovery. Fixing α_X and α_Y reduces the P6 condition number to 2519, but the remaining maximum correlation is still 0.9536. Table 3 quantifies the equal-budget designs.

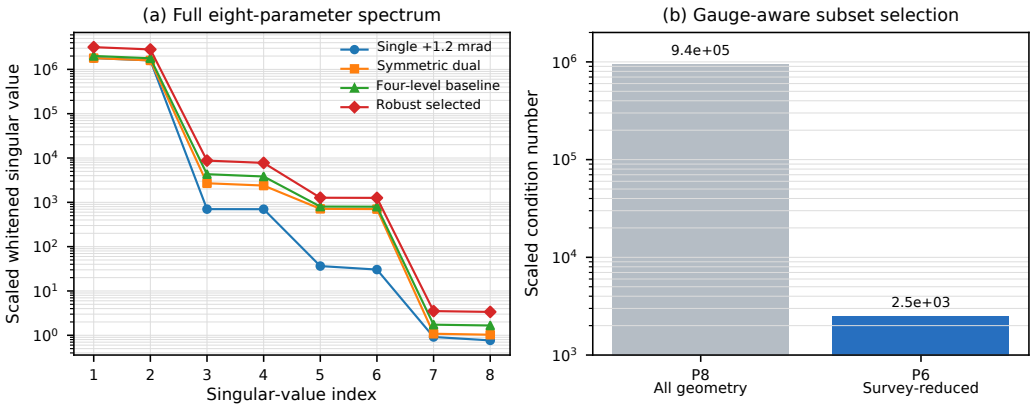

Figure 2. Local scaled-Jacobian diagnostics. (**a**) Singular-value spectra of the full P8 geometry for four representative yaw designs. (**b**) Condition number for P8 and the P6 subset after synthetically fixing both mirror angles. Full numerical rank in P8 does not imply stable recovery because near-degenerate directions remain.

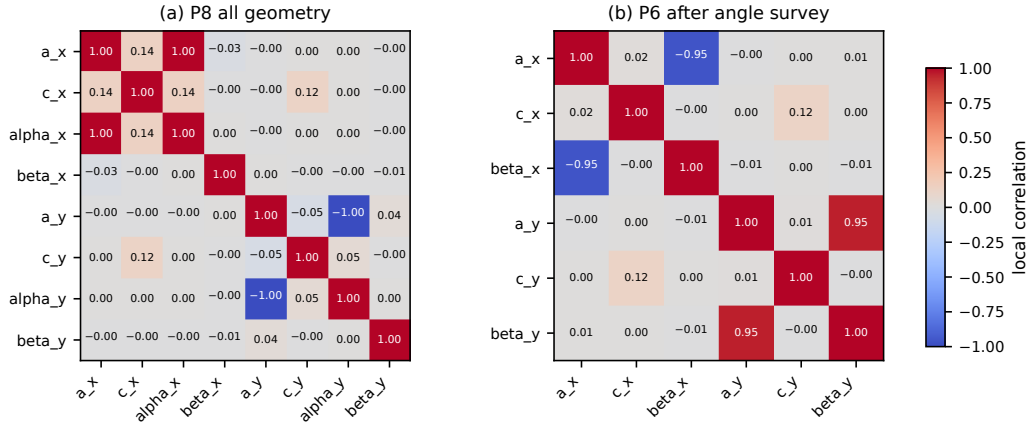

Figure 3. Local parameter-correlation matrices under the selected yaw set. (**a**) P8 retains near-perfect coupling of each Abbe arm with its mirror angle. (**b**) P6 removes those survey-fixed angles, but strong a_j–β_j coupling remains. Correlations derive from the inverse scaled information matrix and are conditional on the declared geometry and covariance.

Table 3. P6 local identifiability for equal-budget paired-yaw designs. $\sigma_{\min}$, κ_2, and ρ use the dimensionless whitened Jacobian.

Trajectory Design	Rank	$\sigma_{\min}$	κ_2	Max. $\lvert\rho\rvert$
Zero yaw	0/6	0.0	∞	–
Single +1.2 mrad	6/6	30.4	5.91×10^4	0.999
Asymmetric dual	6/6	786.9	2.55×10^3	0.824
Symmetric dual	6/6	703.8	2.55×10^3	0.956
Four-level baseline	6/6	797.3	2.52×10^3	0.954
Robust selected	6/6	1260.7	2.52×10^3	0.954

The two profile surfaces in Figure 4 confirm that the remaining high correlations correspond to elongated objective valleys rather than to a harmless covariance-table artifact. Their minima are displaced slightly from the scenario truth because the inverse uses commanded rather than hidden true yaw and also omits the smaller channel and form departures. Consequently, P6 supports conditional recovery under the synthetic survey assumption, but it does not establish unique physical geometry in the presence of broader model discrepancy.

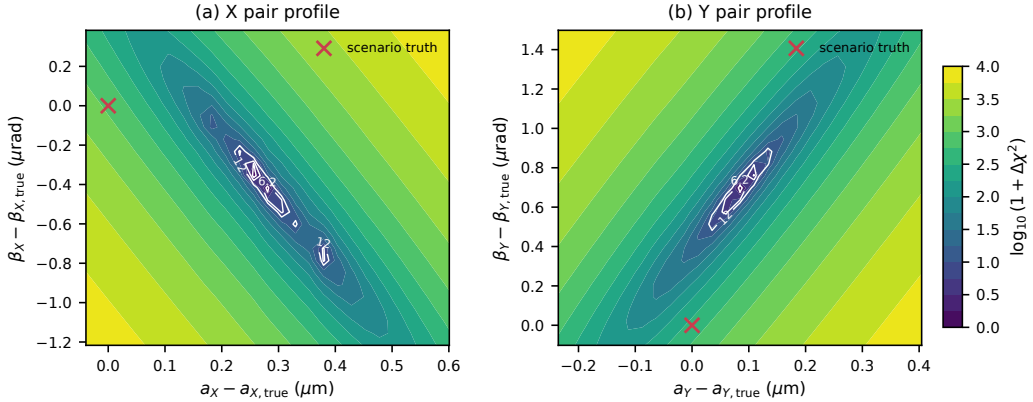

Figure 4. Local two-parameter profile diagnostics after re-optimizing the other four P6 parameters. **(a)** (a_X, β_X); **(b)** (a_Y, β_Y). Axes are offsets from the scenario truth. White contours show selected $\Delta\chi^2$ levels; colour is $\log_{10}(1 + \Delta\chi^2)$. These elongated synthetic objective surfaces are not metrological confidence regions.

3.3. Finite Yaw-Set Selection

The selected set was $\{-2.4, -1.8, 1.8, 2.4\}$ mrad. Figure 5 shows all finite candidates in worst-case information space and the eight highest compromise scores. The extremes enter the selected set because they increase the amplitude of yaw-dependent geometric signatures, while the interior ± 1.8 mrad levels supply additional signed diversity. Robust D and E rankings were checked separately in the saved champion table; the reported selection is the score in Equation (17). Because the spatial raster was fixed and only 70 yaw sets were enumerated, the result should be interpreted as identifiability-guided yaw-set selection within Equation (16), not as a globally optimal calibration trajectory.

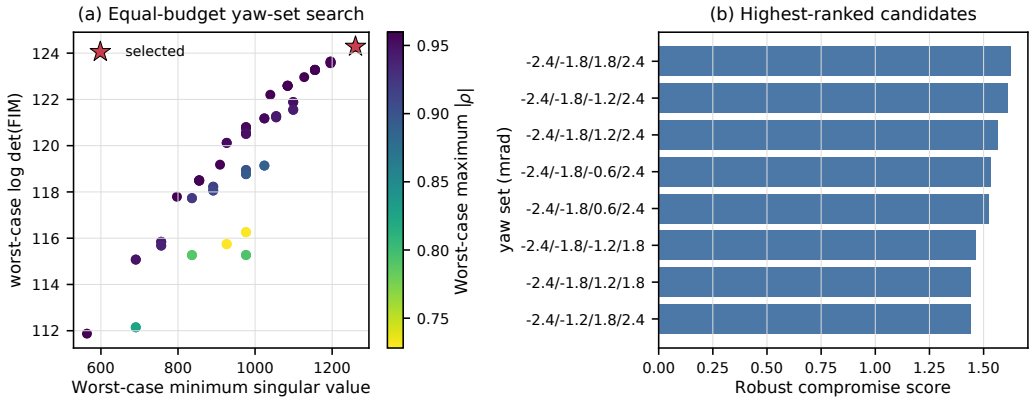

Figure 5. Finite equal-budget yaw-set enumeration over 24 prior geometry draws. (**a**) Worst-case minimum singular value versus worst-case log determinant, coloured by the largest local parameter correlation. The star marks the selected compromise. (**b**) The eight highest-scoring four-level yaw sets; labels are in mrad.

3.4. Calibration and Held-Out Partition

Figure 6 demonstrates the immutable split. Calibration uses a raster and the selected yaw levels, whereas validation uses a Lissajous path and four unseen yaw values. In the executable protocol, the yaw-set and subset calculations receive no validation data, and fitting is completed before held-out scoring. The common held-out realization allows paired solver comparison but does not sample validation-to-validation variability.

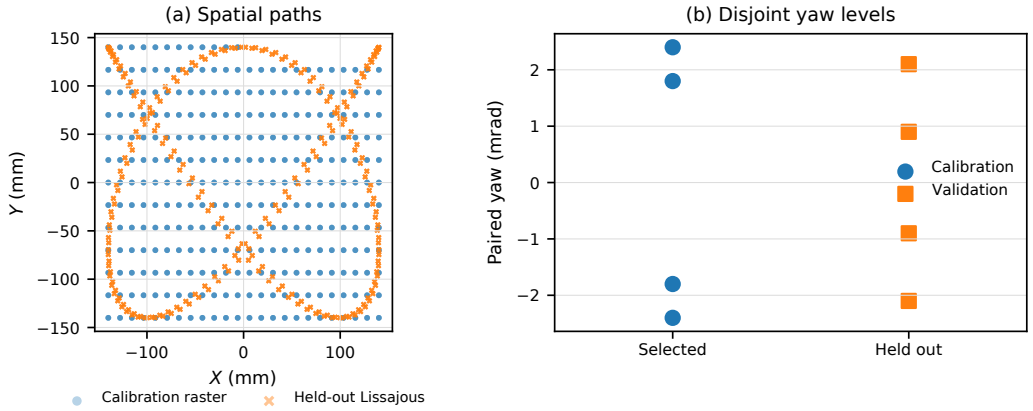

Figure 6. Synthetic data partition. (**a**) Calibration raster and held-out Lissajous spatial paths. (**b**) Calibration yaw levels $(-2.4, -1.8, 1.8, 2.4)$ mrad and disjoint validation levels $(-2.1, -0.9, 0.9, 2.1)$ mrad.

3.5. Solver-Configuration Comparison

All four configurations produced finite estimates in all 100 baseline recalibrations and were flagged as converged under their own declared stopping criteria. Their held-out RMSE distributions were nearly indistinguishable, with medians around 3.70 nm (Figure 7 and Table 4). CW–RLM had a lower median normalized parameter-error norm (0.006 versus approximately 0.010) but used a median of 113 residual evaluations, compared with 5, 9, and 19 for GN, standard L–M, and TRF. The evaluation counts are not strictly equivalent because CW–RLM aggregates starts and IRLS cycles.

This result does not demonstrate that L–M, Huber loss, bounds, or multistart is individually superior. It shows only that the complete CW–RLM configuration traded more computation for lower median parameter error under the declared contamination,

starts, and bounds. The fixed validation mismatch floor masks most parameter-estimation differences in the RMSE metric.

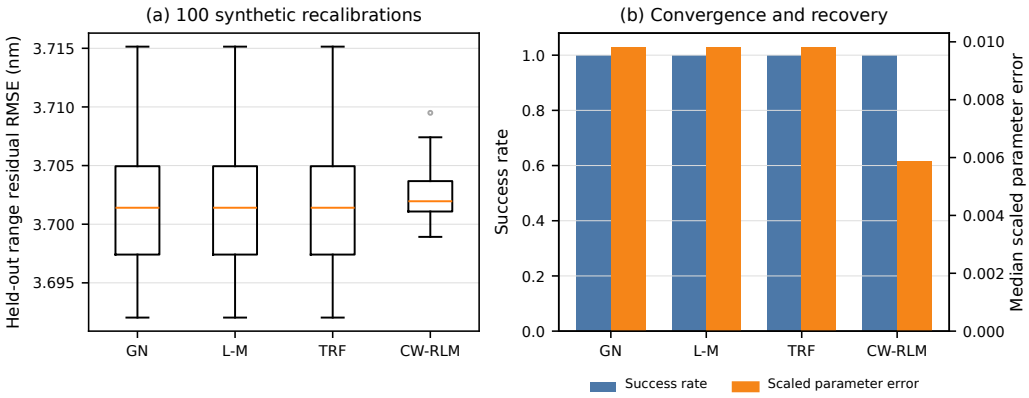

Figure 7. Comparison of complete solver configurations over 100 synthetic recalibrations. (**a**) Held-out pooled range-residual RMSE on one frozen validation realization. (**b**) Convergence fraction and median normalized parameter-error norm. GN: Gauss–Newton; L–M: Levenberg–Marquardt; TRF: trust-region reflective; CW–RLM: robust multistart bound-transformed L–M.

Table 4. Baseline solver-configuration summary over 100 synthetic recalibrations. Evaluation counts have different meanings for single-start and multistart/IRLS methods.

Method	Success (%)	Median RMSE (nm)	Median $\|\|e_p\|\|_2$	Median Eval.
Gauss–Newton	100.0	3.701	0.010	5
Standard L–M	100.0	3.701	0.010	9
Bounded TRF	100.0	3.701	0.010	19
CW–RLM	100.0	3.702	0.006	113

3.6. Parameter Recovery

Figure 8 and Table 5 summarize CW–RLM parameter estimates. Abbe-arm biases were below 0.08 µm, whereas longitudinal-offset biases were approximately −0.012 mm and −0.010 mm for the two families. The pointing biases were less than 0.06 µrad, with empirical standard deviations around 0.3 µrad. These are repeated-fit properties under synthetic truth, exact mirror-angle fixing, a fixed model discrepancy, and the stated noise model. They are not instrument calibration uncertainties.

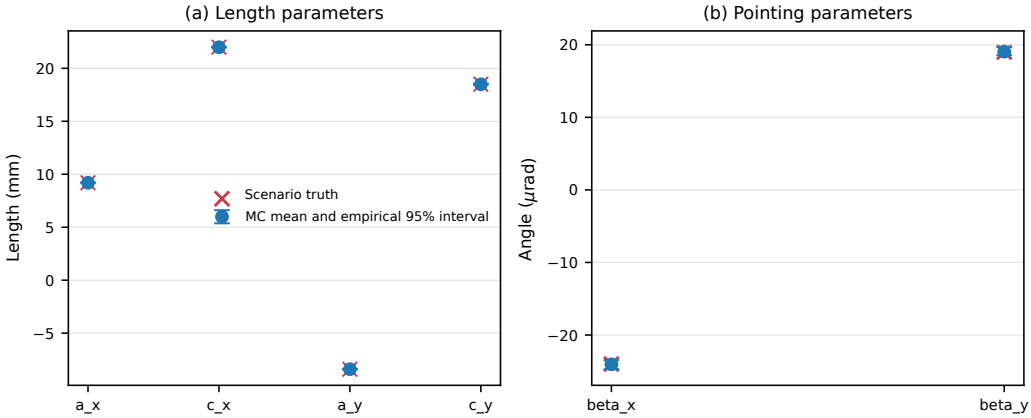

Figure 8. CW–RLM recovery across successful baseline recalibrations. (**a**) Signed Abbe-arm and longitudinal-offset parameters; (**b**) beam-pointing parameters. Circles show Monte Carlo means, error bars show central empirical 95% spans, and crosses mark scenario truth.

Table 5. CW–RLM synthetic recovery. The central empirical span describes the distribution of successful repeated fits; it is not a physical confidence interval.

Method	Success (%)	Median RMSE (nm)	Median $\|e_p\|_2$	Median Eval.
Gauss–Newton	100.0	3.701	0.010	5
Standard L–M	100.0	3.701	0.010	9
Bounded TRF	100.0	3.701	0.010	19
CW–RLM	100.0	3.702	0.006	113

3.7. Held-Out Numerical Prediction

Table 6 separates the frozen held-out residual into the terms of Equation (20). The maximum samplewise closure error is 1.65×10^{-24} m. Yaw-reference mismatch contributes 3.670 nm RMS, compared with 0.130 nm from the fitted-geometry difference and 0.300 nm from readout noise. Oracle subtraction of the yaw term lowers the residual from 3.702 to 0.433 nm without changing the fitted parameters. Thus, the original held-out metric is mainly a test of the assumed yaw reference, not a sensitive discriminator of the estimators. The component cross-moments are retained with the numerical outputs.

Table 6. Additive audit of the frozen held-out CW–RLM residual. The first six rows are samplewise components, not independent uncertainty contributions. The last row uses hidden synthetic yaw and is diagnostic only.

Component or Residual	RMS (nm)
Fitted-geometry difference	0.1297
Yaw-reference mismatch	3.67
Channel and mirror-form departures	0.1499
Readout noise	0.3001
Drift	0.1273
Sparse contamination	0.2463
Complete residual	3.702
Residual after oracle yaw-term subtraction	0.4333

The controlled test removes that validation-side floor. With clean validation and all other departures disabled, the median fitted-geometry prediction error is 0.0133 nm for calibration readout noise alone and 0.114 nm for calibration yaw error alone (Table A1). Enabling both gives 0.114 nm; the zero-perturbation control is at numerical roundoff. These results distinguish direct yaw-reference mismatch during validation from parameter-estimation error induced by yaw error during calibration. The magnitudes are conditional on the specified 0.30 nm readout noise and 25/2.5 nrad block/local yaw perturbations.

The choice to study geometry is also supported by a controlled omission test on the same held-out poses. Setting both Abbe arms to zero without refitting changes the paired-range prediction by 14.23 µm RMS; setting the common pointing angles or mirror angles to zero changes it by 10.08 nm or 1.052 nm RMS, respectively. These are isolated model-effect scales relative to the declared truth, not a measured error budget. Geometry is consequential before correction in this scenario, whereas angular-reference error dominates the residual after fitting. No dominance over omitted optical effects is inferred.

Figure 9 shows representative channels. Their paired ranges span hundreds of micrometers, while the residuals describe conditional numerical prediction at the supplied poses.

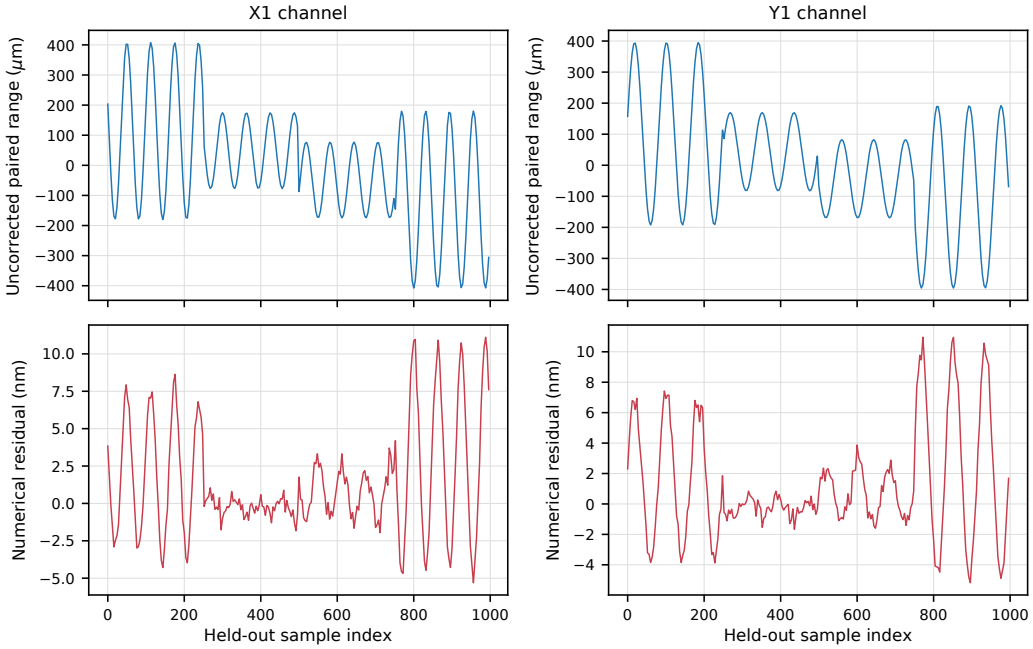

Figure 9. Held-out synthetic prediction for representative X_1 and Y_1 channels. The upper panels show observed one-way-equivalent paired ranges; the lower panels show residuals after evaluating the fitted reduced model. The validation path and yaw levels were not used in yaw-set selection or fitting.

3.8. Exploratory Calibration-Side Stress Tests

Table 7 and Figure 10 report the 20-fit exploratory scenarios. The median conditional held-out RMSE changes little because the same yaw-reference-dominated validation realization is used in every row. Convergence frequency is more sensitive: it falls to 65% under doubled readout noise with baseline weighting and to 30% for 75 nrad block yaw jitter. The no-yaw-jitter case converges in all 20 trials. Triple form mismatch and 1% contamination retain 95% convergence, while the +3 µrad mirror-angle survey offset gives 90%.

Because each percentage has denominator 20, these frequencies are exploratory and have broad binomial uncertainty. The stress-study baseline is an independently seeded 20-fit batch, not a subset of the 100-replicate comparison in Table 4; its 90% frequency therefore does not contradict the separate 100% baseline result. Moreover, the stress is applied to calibration generation only; the common validation file remains at baseline. The rows diagnose failure modes of the estimator configuration, not environmental robustness of a physical interferometer.

Table 7. Exploratory calibration-side stress summary. Every row contains 20 synthetic fits evaluated on the same baseline held-out realization; success percentages are coarse descriptive frequencies.

Scenario	Success (%)	Median RMSE (nm)	90th pct. (nm)
Baseline	90.0	3.702	3.708
Half readout noise	95.0	3.702	3.705
Double readout noise	65.0	3.703	3.708
No yaw-reference jitter	100.0	3.700	3.701
75 nrad yaw jitter	30.0	3.703	3.713
Triple mirror-form mismatch	95.0	3.703	3.707
One-percent contamination	95.0	3.702	3.704
Mirror-angle survey +3 µrad	90.0	3.704	3.709

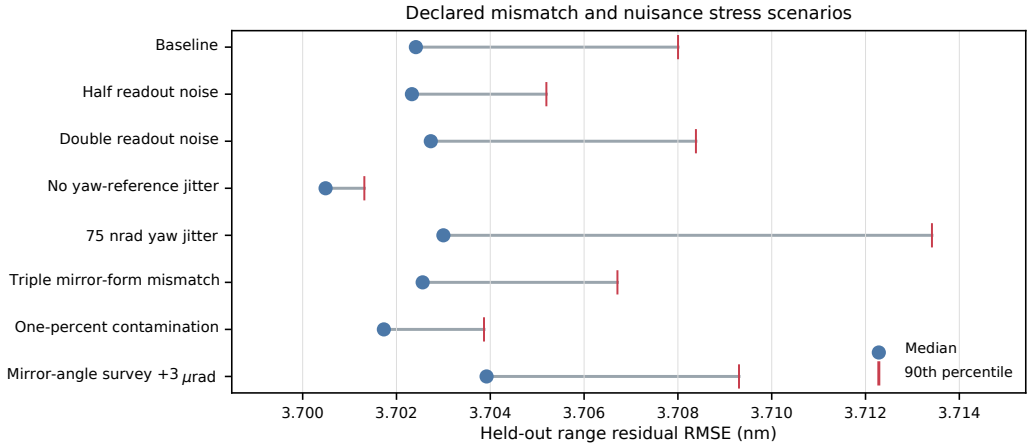

Figure 10. Conditional held-out RMSE after calibration-side stress. Points are medians and red ticks are 90th percentiles over 20 fits. The magnified horizontal scale reveals small differences around a common yaw-reference/model-discrepancy floor.

4. Discussion

4.1. What the Identifiability Results Establish

The central result is that exact numerical rank and useful calibration are not equivalent. The selected P8 Jacobian is rank complete, yet its nearly unit parameter correlations and 10^6-scale condition number expose narrow directions in which materially different geometries have similar first-order signatures. Strong damping could still return a finite parameter vector, but that would not make the eight geometric quantities independently credible. Fixing two mirror angles removes the most severe directions and improves the condition number by more than two orders of magnitude, at the explicit cost of assuming survey information. Information-matrix nonsingularity is a local property of the declared parameterization [28]; consequently, the P6 rank and condition are conditional on the fixed-angle assumptions, which must not be credited to trajectory excitation.

The remaining P6 a_j–β_j correlations have a geometric interpretation. Changing a transverse ray offset modifies the yaw-dependent lever arm, while a pointing change perturbs the ray–mirror denominator and hit position. Over a small yaw domain, these effects can produce similar signatures. Signed and multi-level yaw increases the weakest singular values, but Figure 4 shows that an elongated valley remains. A future physical calibration should therefore include an independent beam-angle survey, a broader validated motion set, or an additional observable sensitive to beam direction. Merely increasing the L–M damping parameter would not supply missing information.

4.2. Interpretation of the Finite Design Search

The selected extreme-and-near-extreme signed yaw set is plausible because Abbe terms grow with angular excitation and sign reversal changes odd coupling signatures. Nevertheless, the search is conditional on a fixed raster, equal allocation, selected bounds, one covariance model, and 24 prior draws. D-optimality rewards information volume, while E-type criteria protect the weakest direction; these objectives need not agree for a constrained nonlinear optical system [29–31]. The normalized compromise score prevents one metric from numerically dominating this particular finite table, but its normalization depends on the candidate collection. The correct claim is therefore improved local information within a declared 70-set enumeration. Continuous yaw scheduling, transition dynamics, beam-footprint limits, collision constraints, and apparatus settling were not optimized.

The paired-yaw construction also imposes a demanding acquisition assumption: nonzero- and zero-yaw samples must correspond to the same (X, Y). In physical data,

interpolation, timing, drift, and reference-pose error would affect the difference covariance. Equation (7) treats that paired range as already available and assigns it a covariance directly. Before hardware use, the protocol would need synchronized reference metrology or a trajectory that repeatedly returns to common spatial points. Without it, cancellation of range datums and common-mode errors cannot be presumed.

4.3. Estimator Results and Their Limits

The solver comparison is intentionally reported as a configuration comparison. GN, standard L–M, and TRF share one perturbed start, while CW–RLM combines a reference start, another seeded start, bounds, a nonlinear bound map, Huber weighting, and IRLS. Its lower normalized parameter error may result from any combination of these choices. A causal algorithm study would require a common start bank and ablations for scaling, bounds, robust loss, and multistart. Standard L–M is a baseline rather than an innovation [33,34].

The residual decomposition explains why improved parameter estimates barely change the original held-out RMSE. The direct validation-side yaw term masks smaller geometric differences, while the clean-validation factorial isolates their calibration-side origin. These two tests answer different questions and should not be combined into a single claimed accuracy. In practice, both the angular reference used to collect calibration data and the reference used to assess prediction would need independent characterization.

The empirical standard deviations and central 95% spans in Table 5 describe 100 repeated calibration fits under one generator family and one frozen validation realization. They are neither confidence intervals for physical parameter values nor expanded measurement uncertainties. A metrological evaluation would require uncertainties and correlations for surveyed geometry, yaw and translation references, wavelength, air refractive index, and model discrepancy, propagated to a defined displacement measurand [38–40]. Fixed mirror angles would then have uncertain inputs rather than exact values. Archived local Wald inclusion proportions are diagnostic only: the covariance does not propagate robust weights through a sandwich or bootstrap calculation, and successful-fit conditioning is not unconditional coverage. More repetitions can reduce Monte Carlo sampling error, but cannot supply missing physical inputs or traceability.

4.4. Optical Effects Outside the Executable Model

The geometric paired range is only one part of an interferometric measurement. Geometric and non-geometric tilt-to-length coupling need not have the same sign or magnitude; phase also depends on wavefront overlap and detector geometry [41–43]. Table 8 identifies the main omitted mechanisms and the additional information needed to evaluate them. None are assigned an arbitrary numerical amplitude to claim a complete error budget.

An omitted signal component aligned with columns of the calibration Jacobian can be absorbed into a biased geometry estimate; an orthogonal component remains in the residual. Thus, a small fit residual cannot by itself rule out these mechanisms. The independently coded generator remains planar, with zero vertical motion, pitch, roll, and beam elevation. Its channel/form departures test selected model mismatch, not the omitted effects in Table 8. Similarly, the prescribed wavelength record is only a phase-encoding input and is not evidence of wavelength compensation.

Table 8. Physical mechanisms outside the executed model and the resulting applicability limits. These effects were not simulated in the additional controlled tests.

Mechanism	Why It Matters	Required Extension or Check
Beam walk-off and finite apertures	Motion changes the footprint and overlap; clipping can alter phase and signal availability.	Trace the return path, beam radii, and detector/mirror apertures; restrict the yaw set to unclipped operation.
Polarization leakage	Parasitic interference and frequency mixing can create periodic phase error [14,17].	Characterize polarization and parasitic paths; test phase nonlinearity over multiple fringes.
Diffraction and wavefront aberration	A detected phase is a field-overlap quantity, not solely a geometric ray length [43].	Use wave-optical propagation and measured wavefront/detector information.
Pitch, roll, and vertical motion	Additional lever arms and mirror rotations modify the path and couple into the in-plane readout.	Extend the pose and Jacobian to three dimensions, with independently observed nuisance motions.
Air-path and wavelength variation	Uncertain scale and spatially varying optical path affect inferred displacement [16,44,45].	Measure environmental and wavelength inputs and propagate their uncertainty.
Pairing and temporal drift	Nonzero- and zero-yaw acquisitions may not share the same position or time-dependent offsets.	Synchronize reference metrology and characterize the covariance of the actual differencing procedure.

4.5. Transfer Across Layouts and Interferometer Structures

The transferable procedure is to derive the actual observation operator, remove representation gauges, scale its Jacobian, test rank and conditioning over feasible motions, and validate fitted parameters on independent observations. The paired-yaw formula and selected yaw values themselves are topology dependent. They require matched translations, identifiable mirror/channel associations, stable geometry, and an adequate angular reference. A survey-fixed parameter must be known sufficiently well for the intended measurand.

Appendix B varies the two beam spacings and the stand-off within the declared four-channel architecture. All five cases retain P8 rank 8 but condition numbers between 6.75×10^5 and 1.08×10^6. P6 remains near 2.52×10^3 after both mirror angles are fixed. This supports persistence of the local rank–conditioning distinction under those changes; it does not establish generality across optical architectures. Grating, multipass, heterodyne phase-mixing, or non-planar arrangements require their own phase/path operator and a repeated information/design analysis. Neither the present parameter subset nor the four selected yaw levels should be transferred without that analysis.

4.6. Requirements for Physical Validation

An experimental implementation should first survey beam origins, directions, mirror normals, and finite footprints, and measure phase nonlinearity over several fringes. The paired trajectories should then be acquired with synchronized independent translation and yaw references, such as reference interferometers and an autocollimator with suitable characterized uncertainty. Their errors should be propagated rather than replaced by commanded pose. Calibration and validation should use separate trajectories and sessions, including different speeds and days. Recording wavelength, environmental conditions, and beam overlap would help distinguish geometric bias from optical and temporal effects. Finally, the evaluation should report parameter variability, functional-point error, repeatability, and measurement uncertainty separately. This sequence tests the specific survey and yaw-reference requirements exposed by the numerical analysis.

5. Conclusions

The paired-yaw analysis identifies two distinct limits to geometric calibration: weak separation of Abbe arms from mirror angles, and a validation floor imposed by the yaw reference. Full local numerical rank did not make the eight-parameter geometry stable. Fixing both mirror angles reduced the condition number from 9.42×10^5 to 2519, while a finite equal-budget search selected $\{-2.4, -1.8, 1.8, 2.4\}$ mrad within the prescribed grid. The additive residual audit and controlled noise/yaw tests explain why solver configurations with different parameter errors can produce similar held-out residuals. These model-specific results provide a numerical screen for survey requirements and calibration motions. Their experimental use requires an independently characterized pose reference, optical-effect checks, and uncertainty propagation to a defined measurand.

Author Contributions: Conceptualization, C.W. and H.D.; methodology, C.W. and H.D.; software, C.W.; validation, C.W.; formal analysis, C.W.; investigation, C.W.; resources, H.D.; data curation, C.W.; writing—original draft preparation, C.W.; writing—review and editing, C.W. and H.D.; visualization, C.W.; supervision, H.D.; project administration, H.D.; funding acquisition, H.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Guangxi Natural Science Foundation (Young Scientists Program), grant number 2025GXNSFBA069395, under the project "Research on Periodic Nonlinear Error Compensation Method for Nanoscale Grating Interferometer".

Institutional Review Board Statement: Not applicable. This mathematical and computational study involved neither humans nor animals.

Informed Consent Statement: Not applicable.

Data Availability Statement: All observations were generated numerically; no physical measurement data were used. The source code, frozen random seeds, input configurations, and supporting numerical data are retained in an author-held archive and are not publicly available. The authors authorize confidential access to this archive by the handling editors and assigned reviewers for evaluation of this manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

CW–RLM	Covariance-weighted robust multistart bound-transformed L–M
FIM	Fisher information matrix
GN	Gauss–Newton
IRLS	Iteratively reweighted least squares
L–M	Levenberg–Marquardt
P6/P8	Six-/eight-parameter geometry subset
RMS/RMSE	Root mean square/root mean squared error
TRF	Trust-region reflective

Appendix A. Controlled Calibration Perturbation Test

The four settings switch correlated readout noise ($\sigma_r = 0$ or 0.30 nm) and yaw-reference error (zero or 25 nrad block plus 2.5 nrad local standard deviation) independently. Each fit uses 600 raster positions and the selected four yaw levels. Repeat $i = 0, \ldots, 99$ uses NumPy seed $20261102 + 1000 + i$; the same generated noise and yaw arrays are reused across the four settings. The zero-perturbation setting repeats an identical deterministic control and does not provide 100 independent uncertainty samples.

Channel departures, mirror form, drift, and outliers are disabled. P6 is fitted with bounded TRF, analytic derivatives, a fixed design-reference start, tolerances 10^{-12}, and at

most 400 function evaluations. The two mirror angles remain fixed at truth. Prediction is scored against the exact inverse operator at 1000 clean Lissajous positions and the four unseen validation yaw levels, with no validation noise or yaw perturbation. All 400 fits returned finite estimates and a convergence flag. Table A1 reports empirical quantiles, not physical confidence bounds. This controlled TRF test is separate from the original CW–RLM configuration comparison.

Table A1. Clean-validation prediction error due to the fitted geometry in the controlled 2×2 test. "On" denotes the magnitudes specified above; each setting has 100 fits. The off/off row is a repeated deterministic roundoff control.

Readout Noise	Yaw Error	Median (nm)	2.5th Pct. (nm)	97.5th Pct. (nm)
Off	Off	3.482×10^{-8}	3.482×10^{-8}	3.482×10^{-8}
On	Off	0.01326	0.006254	0.02034
Off	On	0.1141	0.0233	0.254
On	On	0.1136	0.0224	0.2547

Appendix B. Local Screen of Five Four-Channel Layouts

The X-pair and Y-pair nominal spacings are 90 and 80 mm, and the nominal stand-off is 250 mm. The screen halves or doubles both spacings, halves only the Y spacing, or increases stand-off to 350 mm. All other geometry parameters, the 1200-position raster, yaw levels, scales, and covariance remain unchanged. Information metrics are evaluated at the scenario truth; the finite yaw search and parameter recovery are not repeated. Table A2 therefore tests local conditioning within one architecture, not optical feasibility or calibration performance in other architectures. Finite apertures and return-beam overlap are not checked.

Table A2. Local information metrics for five declared layouts. All cases have P8 rank 8 and P6 rank 6. Beam spacings and stand-off are in millimeters. The P6 condition changes little at the displayed precision.

Layout	X Spacing	Y Spacing	Stand-off	P8 κ_2	P6 κ_2
Declared	90	80	250	9.421×10^5	2519
Half spacings	45	40	250	1.079×10^6	2519
Double spacings	180	160	250	6.75×10^5	2519
Unequal spacings	90	40	250	1.078×10^6	2519
Longer stand-off	90	80	350	9.421×10^5	2519

References

1. Holmes, M.; Hocken, R.; Trumper, D.L. The long-range scanning stage: a novel platform for scanned-probe microscopy. *Precis. Eng.* **2000**, *24*, 191–209. https://doi.org/10.1016/S0141-6359(99)00044-6.
2. Fan, K.C.; Chen, M.J. A 6-degree-of-freedom measurement system for the accuracy of X-Y stages. *Precis. Eng.* **2000**, *24*, 15–23. https://doi.org/10.1016/S0141-6359(99)00021-5.
3. Gao, W.; Arai, Y.; Shibuya, A.; Kiyono, S.; Park, C.H. Measurement of multi-degree-of-freedom error motions of a precision linear air-bearing stage. *Precis. Eng.* **2006**, *30*, 96–103. https://doi.org/10.1016/j.precisioneng.2005.06.003.
4. Dirckx, J.J.J.; van Elburg, H.J.; Decraemer, W.F.; Buytaert, J.A.N.; Melkebeek, J.A. Performance and testing of a four channel high-resolution heterodyne interferometer. *Opt. Lasers Eng.* **2009**, *47*, 488–494. https://doi.org/10.1016/j.optlaseng.2008.05.011.
5. Hsieh, H.L.; Pan, S.W. Three-degree-of-freedom displacement measurement using grating-based heterodyne interferometry. *Appl. Opt.* **2013**, *52*, 6840–6848. https://doi.org/10.1364/AO.52.006840.
6. Gillmer, S.R.; Smith, R.C.G.; Woody, S.C.; Ellis, J.D. Compact fiber-coupled three degree-of-freedom displacement interferometry for nanopositioning stage calibration. *Meas. Sci. Technol.* **2014**, *25*, 075205. https://doi.org/10.1088/0957-0233/25/7/075205.

7. Yan, H.; Duan, H.Z.; Li, L.T.; Liang, Y.R.; Luo, J.; Yeh, H.C. A dual-heterodyne laser interferometer for simultaneous measurement of linear and angular displacements. *Rev. Sci. Instrum.* **2015**, *86*, 123102. https://doi.org/10.1063/1.4936771.
8. Yu, X.; Gillmer, S.R.; Woody, S.C.; Ellis, J.D. Development of a compact, fiber-coupled, six degree-of-freedom measurement system for precision linear stage metrology. *Rev. Sci. Instrum.* **2016**, *87*, 065109. https://doi.org/10.1063/1.4953335.
9. Cai, Y.; Lou, Z.; Ling, S.; Liao, B.S.; Fan, K.C. Development of a compact three-degree-of-freedom laser measurement system with self-wavelength correction for displacement feedback of a nanopositioning stage. *Appl. Sci.* **2018**, *8*, 2209. https://doi.org/10.3390/app8112209.
10. Chang, C.P.; Tu, T.C.; Huang, S.R.; Wang, Y.C.; Chang, S.C. Development of the heterodyne laser encoder system for the X–Y positioning stage. *Sensors* **2021**, *21*, 5775. https://doi.org/10.3390/s21175775.
11. Straube, G.; Fischer Calderón, J.S.; Ortlepp, I.; Füßl, R.; Manske, E. A heterodyne interferometer with separated beam paths for high-precision displacement and angular measurements. *Nanomanufacturing Metrol.* **2021**, *4*, 200–207. https://doi.org/10.1007/s41871-021-00101-x.
12. Bryan, J.B. The Abbé principle revisited: an updated interpretation. *Precis. Eng.* **1979**, *1*, 129–132. https://doi.org/10.1016/0141-6359(79)90037-0.
13. Estler, W.T. High-accuracy displacement interferometry in air. *Appl. Opt.* **1985**, *24*, 808–815. https://doi.org/10.1364/AO.24.000808.
14. Wu, C.M.; Lawall, J.; Deslattes, R.D. Heterodyne interferometer with subatomic periodic nonlinearity. *Appl. Opt.* **1999**, *38*, 4089–4094. https://doi.org/10.1364/AO.38.004089.
15. Lawall, J.; Kessler, E. Michelson interferometry with 10 pm accuracy. *Rev. Sci. Instrum.* **2000**, *71*, 2669–2676. https://doi.org/10.1063/1.1150715.
16. Ciddor, P.E. Refractive index of air: new equations for the visible and near infrared. *Appl. Opt.* **1996**, *35*, 1566–1573. https://doi.org/10.1364/AO.35.001566.
17. Cavagnero, G.; Mana, G.; Massa, E. Aberration effects in two-beam laser interferometers. *J. Opt. Soc. Am. A* **2006**, *23*, 1951–1959. https://doi.org/10.1364/JOSAA.23.001951.
18. Du, H.; Zhang, W.; Xiong, X.; Zeng, Q.; Wang, Y.; Zhang, Y.; Xu, S.; Li, H. Influence of installation error of grating interferometer on high-precision displacement measurement. *Opt. Eng.* **2021**, *60*, 045102. https://doi.org/10.1117/1.OE.60.4.045102.
19. Clark, L.; Shirinzadeh, B.; Tian, Y.; Oetomo, D. Laser-based sensing, measurement, and misalignment control of coupled linear and angular motion for ultrahigh precision movement. *IEEE/ASME Trans. Mechatronics* **2015**, *20*, 84–92. https://doi.org/10.1109/TMECH.2014.2301824.
20. Cai, Y.; Wang, L.; Liu, Y.; Li, C.; Fan, K.C. Accuracy improvement of linear stages using on-machine geometric error measurement system and error transformation model. *Opt. Express* **2022**, *30*, 7539–7550. https://doi.org/10.1364/OE.453111.
21. Lin, R.; Cheng, Z.; He, Y.; Li, J.; Li, R. On-line six-degree-of-freedom errors measurement and compensation for XY stages. *Opt. Precis. Eng.* **2023**, *31*, 1761–1773. https://doi.org/10.37188/OPE.20233112.1761.
22. Zhang, J.; Iwata, K.; Shibuya, A.; Kikuta, H.; Park, C.S. Method for calibrating system parameters of a multidirectional interferometers system. *Appl. Opt.* **2006**, *45*, 605–610. https://doi.org/10.1364/AO.45.000605.
23. Evans, C.J.; Hocken, R.J.; Estler, W.T. Self-calibration: reversal, redundancy, error separation, and "absolute testing". *CIRP Ann.* **1996**, *45*, 617–634. https://doi.org/10.1016/S0007-8506(07)60515-0.
24. Nitsche, J.; Franke, M.; Haverkamp, N.; Heißelmann, D. Six-degree-of-freedom pose estimation with micrometre/microradian accuracy based on laser multilateration. *J. Sens. Sens. Syst.* **2021**, *10*, 19–24. https://doi.org/10.5194/jsss-10-19-2021.
25. Ye, W.; Zhang, M.; Zhu, Y.; Wang, L.; Hu, J.; Li, X.; Hu, C. Real-time displacement calculation and offline geometric calibration of the grating interferometer system for ultra-precision wafer-stage measurement. *Precis. Eng.* **2019**, *60*, 413–420. https://doi.org/10.1016/j.precisioneng.2019.06.012.
26. Wen, T.; Hu, J.; Zhu, Y.; Hua, G. Multi-axis laser interferometer not affected by installation errors based on nonlinear computation. *Appl. Sci.* **2023**, *13*, 10887. https://doi.org/10.3390/app131910887.
27. Hsu, W.H.; Tzeng, Y.C.; Liu, C.S. A systematic modeling approach based on skew ray tracing for optical measurement systems using geometrical optics and its applications in a six-DOF geometrical error measurement system for linear stages. *Measurement* **2026**, *267*, 120465. https://doi.org/10.1016/j.measurement.2026.120465.
28. Rothenberg, T.J. Identification in parametric models. *Econometrica* **1971**, *39*, 577–591. https://doi.org/10.2307/1913267.
29. Kiefer, J.; Wolfowitz, J. The equivalence of two extremum problems. *Can. J. Math.* **1960**, *12*, 363–366. https://doi.org/10.4153/CJM-1960-030-4.
30. Jung, Y.; Lee, I. Optimal design of experiments for optimization-based model calibration using Fisher information matrix. *Reliab. Eng. & Syst. Saf.* **2021**, *216*, 107968. https://doi.org/10.1016/j.ress.2021.107968.
31. Wilson, A.D.; Schultz, J.A.; Murphey, T.D. Trajectory synthesis for Fisher information maximization. *IEEE Trans. Robot.* **2014**, *30*, 1358–1370. https://doi.org/10.1109/TRO.2014.2345918.

32. Levenberg, K. A method for the solution of certain non-linear problems in least squares. *Q. Appl. Math.* **1944**, *2*, 164–168. https://doi.org/10.1090/qam/10666.
33. Marquardt, D.W. An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Ind. Appl. Math.* **1963**, *11*, 431–441. https://doi.org/10.1137/0111030.
34. Moré, J.J. The Levenberg–Marquardt algorithm: implementation and theory. In *Numerical Analysis*; *Lecture Notes in Mathematics*; Watson, G.A., Ed.; Springer: Berlin\Heidelberg, Germany, 1978; Volume 630, pp. 105–116. https://doi.org/10.1007/BFb0067700.
35. Coleman, T.F.; Li, Y. An interior trust region approach for nonlinear minimization subject to bounds. *SIAM J. Optim.* **1996**, *6*, 418–445. https://doi.org/10.1137/0806023.
36. Harris, C.R.; Millman, K.J.; van der Walt, S.J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N.J.; et al. Array programming with NumPy. *Nature* **2020**, *585*, 357–362. https://doi.org/10.1038/s41586-020-2649-2.
37. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **2020**, *17*, 261–272. Correction in *Nat. Methods* **2020**, *17*, 352. https://doi.org/10.1038/s41592-019-0686-2.
38. Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—Guide to the Expression of Uncertainty in Measurement*; Technical Report JCGM 100:2008; Bureau International des Poids et Mesures: Sèvres, France, 2008. https://doi.org/10.59161/JCGM100-2008E.
39. Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—Supplement 1 to the Guide to the Expression of Uncertainty in Measurement—Propagation of Distributions Using a Monte Carlo Method*; Technical Report JCGM 101:2008; Bureau International des Poids et Mesures: Sèvres, France, 2008. https://doi.org/10.59161/JCGM101-2008.
40. Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—Supplement 2 to the Guide to the Expression of Uncertainty in Measurement—Extension to Any Number of Output Quantities*; Technical Report JCGM 102:2011; Bureau International des Poids et Mesures: Sèvres, France, 2011. https://doi.org/10.59161/JCGM102-2011.
41. Fitzsimons, E.D.; Bogenstahl, J.; Hough, J.; Killow, C.J.; Perreur-Lloyd, M.; Robertson, D.I.; Ward, H. Precision absolute positional measurement of laser beams. *Appl. Opt.* **2013**, *52*, 2527–2530. https://doi.org/10.1364/AO.52.002527.
42. Hartig, M.S.; Schuster, S.; Wanner, G. Geometric tilt-to-length coupling in precision interferometry: mechanisms and analytical descriptions. *J. Opt.* **2022**, *24*, 065601. https://doi.org/10.1088/2040-8986/ac675e.
43. Hartig, M.S.; Schuster, S.; Heinzel, G.; Wanner, G. Non-geometric tilt-to-length coupling in precision interferometry: mechanisms and analytical descriptions. *J. Opt.* **2023**, *25*, 055601. https://doi.org/10.1088/2040-8986/acc3ac.
44. Haitjema, H. Calibration of displacement laser interferometer systems for industrial metrology. *Sensors* **2019**, *19*, 4100. https://doi.org/10.3390/s19194100.
45. Liao, Y.; Zhang, Z.; Wang, N.; Chen, Z.; Hao, L.; Yang, X. Environmental compensation of laser interferometer based on particle swarm algorithm. *Appl. Opt.* **2022**, *61*, 3648–3655. https://doi.org/10.1364/AO.452854.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.