

A Virtual Pipeline for the Evaluation of Optical Distortion in Panelled Glass

Guillermo Casas ^a, Adrian Betanzos ^b, Louis Moreau ^c

- a CIMNE, Spain, gcasas@cimne.upc.edu
- b VIA, United States, adrian@via.inc
- c MOGLEX Corporation, Canada, louis@moglex.com

Abstract

Architectural panelled glass lacks an objective method for assessing optical distortion: surface topography captures only reflected flatness (weakly correlated with the through-glass distortion observers report) and visual inspection is neither reproducible nor deployable on installed panels. We address this gap with an image-based methodology that measures distortion as observers perceive it: the apparent displacement of a background pattern seen through the glass, reported as a two-channel field of sphere-equivalent optical power and cylinder-equivalent astigmatism. The methodology adapts Fourier Transform Profilometry and phase-measuring deflectometry to architectural glass, and introduces a reference-free Direct Method that recovers the undistorted grid via a RANSAC homography fit, removing the need for a glass-less calibration and enabling deployment both in-factory and in-situ. We design and validate the Distortion Analyzer within a virtual pipeline of ray-traced simulations with known ground truth. Validation against the Pipeline's independent analytical forward model confirms spatial agreement, convergence under grid refinement, and rotational stability. Physical validation and perceptual calibration remain as next steps.

Keywords

Optical distortion, Architectural glass, Phase-measuring deflectometry, Fourier Transform Profilometry, Visual quality assessment

Article Information

- Digital Object Identifier (DOI): [10.47982/cgc.10.687](https://doi.org/10.47982/cgc.10.687)
- Published by [Challenging Glass](#), on behalf of the author(s), at [Stichting OpenAccess](#).
- Published as part of the peer-reviewed [Challenging Glass Conference Proceedings](#), Volume 10, June 2026, [10.47982/cgc.10](https://doi.org/10.47982/cgc.10)
- Editors: Christian Louter, Freek Bos & Jan Belis
- This work is licensed under a [Creative Commons Attribution 4.0 International](#) (CC BY 4.0) license.
- Copyright © 2026 with the author(s)

1. Introduction

1.1. The measurement gap

The architectural glass industry is shifting toward larger thicker and often laminated glass. Heat-treated laminates with coated glass and frit are prone to create visible lensing effects, particularly under oblique viewing conditions. See-through distortion is now a leading visual-quality concern in contemporary façades, and with it the need for objective evaluation.

That need is largely unmet, a gap consolidated in the recent comprehensive review of Henriksen et al. (2026), which catalogues eight distinct optical-distortion phenomena in architectural glass, surveys the existing standards landscape (ASTM C1652/C1652M, EN 572-2, EN 1863-1, EN 12150-1, EN ISO 12543-6, EN ISO 1279-1), and concludes that the most promising path is to extend the monolithic-transmission millidiopter measurement to laminated glass, IGUs, and curved glass. Existing standards and inspection practices fall into two groups, both addressing related but distinct quantities. Flat-glass product specifications — ASTM C1036, ASTM C1048, EN 572, EN 12150-1/-2, EN 14179-1 — set quality classes and acceptance vocabulary (associated distortion, ream / string / line distortion, vision interference angle in C1036) without a perceptually-anchored distortion field. Distortion-specific test methods reduce the field to a single scalar: ASTM C1651 converts peak-to-valley roll-wave geometry to optical distortion in millidiopters; ASTM C1652/C1652M scores grid photographs against a fixed acceptance pattern.

In the automotive industry, all OEM windshields undergo safety testing that adheres to UN/ECE Regulation No. 43, ISO 3538, or ANSI/SAE Z26.1 standards. These regulations specify laboratory setups, while companies like SYNERGX and ISRA offer automated machines that provide the same measurement results in minutes of arc. These scanners are not currently available for use in architectural applications.

None of these methods provides a perception-aligned, geometry-aware metric usable on installed multi-layer assemblies (IGUs, laminates, coated systems), where the net effect depends on the combined behaviour of all layers. Visual-inspection methods — zebra-board tests, and computer-vision line detection (Hough transform; Duda and Hart 1972) applied to digitised zebra-board images — capture the perceived effect but only under fixed factory geometries and with operator dependence, making them neither reproducible nor applicable to installed panels. The same gap is documented from the practitioner side in the air-traffic-control-tower glazing review of Moreau (2024) and in our GPD2025 precursor (Moreau, Betanzos and Casas Gonzalez 2025), which compares architectural (ASTM C1036, C1652) and automotive (UN ECE R43) test methods on the same laminate and explicitly flags the projected-grid transmission test as a candidate for automation and on-site deployment. The absence of an objective, perception-aligned, geometry-aware method that works *both* at the factory line, with multi-channel rather than scalar reporting, *and* in-situ on installed panels with no glass-less reference capture, underlies many post-installation disputes, where subjective assessments must be reconciled without a shared quantitative reference.

This work addresses that gap with a single image-based methodology applicable in both contexts. Distortion is measured through the apparent displacement of a background pattern seen through the glass, and reported as a two-channel field of sphere-equivalent optical power and cylinder-equivalent astigmatism whose two channels respond to genuinely different distortion modes (area change vs. shape change), giving factory-line operators the diagnostic axes their scalar systems lack. A reference-free Direct Method recovers the undistorted grid from a single image, removing the need for a glass-

less calibration and making in-situ deployment on installed panels practical with the same protocol used at the factory. The methodology is developed within a Virtual Pipeline that pairs physically-based ray-traced renderings with an independent analytical forward model on the same glass geometry, supplying controlled scenarios with known ground truth.

The analysis assumes paraxial small-angle conditions, appropriate to the small deviations typical of architectural glass; recovering per-pane surface geometry rather than the assembly's net optical effect further depends on barometric and thermal state and is left as future work. Stress-induced optical anisotropy in toughened glass (e.g. birefringence; Illguth et al. 2015) is out of scope.

1.2. Optics and protocol

Optical distortion in architectural glass arises primarily from refraction. As light passes through a pane, any variation in thickness, surface orientation, or refractive index alters the direction of transmitted rays; in nearly-flat glass the cumulative effect is a continuous field of weak local lenses. The result is a spatially varying displacement of the background scene, the displacement field $\delta(x, y)$ that the paper measures.

What observers perceive is this displacement, not the glass geometry. Because architectural backgrounds lie far from the façade, the dominant effect is the angular deflection of transmitted rays, translating into lateral shifts of distant features. Shift magnitude depends strongly on viewing configuration (viewing angle, glass-to-background and viewer-to-glass distances), which is why distortion that looks mild under factory geometry can become prominent once installed.

Modern glazing is typically multi-layered (IGUs, laminates, coated systems), and an observer sees the net angular deflection after light has traversed the full stack. We therefore model the assembly as a single equivalent thin optical surface reproducing this net deflection $\Delta\theta_{\text{equiv}}(x, y)$. In the paraxial regime, this is well founded in geometrical optics (Hecht 2017): it collapses layer-spacing, coating, and environmental complexity into one perceptually meaningful descriptor while preserving the dominant visual effect (Figure 1).




Fig. 1: From multi-layer reality to the equivalent single surface. (a) In an IGU, light from background point B is refracted at every glass-air interface, with local angular deflection $\Delta\theta$. (b) The equivalent net deflection $\Delta\theta_{\text{equiv}}$ is reproduced by a single equivalent thin surface, justifying the abstraction used throughout.

The equivalent surface admits a natural two-channel readout. Local convergence or divergence of δ corresponds to a positive or negative lens, identifying the sphere-equivalent optical power as a measure of local area change (after ophthalmic optics; Atchison and Smith 2000; full-trace convention, defined in §2.2.3). Anisotropic deformation, where one axis is stretched while the perpendicular is compressed, identifies the cylinder-equivalent astigmatism as a measure of local shape change. Together they form an interpretable two-axis decomposition of the distortion field, the basis for the protocol in §2.

1.3. Objectives and organisation

This paper develops an image-based methodology for assessing optical distortion in architectural panelled glass within a Virtual Pipeline supplying controlled test cases with known ground truth. It pursues four objectives:

1. Methodology. Develop a perception-aligned methodology that measures distortion through the apparent displacement of a background pattern seen through the glass and is deployable in both deployment contexts the industry needs: at the factory line, where a glass-less reference can be captured but multi-channel diagnostic reporting is currently not standard; and in-situ on installed panels, where no reference capture is possible. Single-camera capture on the same printed grid suffices in both cases. The present implementation characterises the detrended local component (ripples and localised defects); bulk curvature and residual perspective are deliberately removed, left for future multi-view or independently-calibrated extensions (§2.2.3).
2. Virtual Pipeline. Build a Virtual Pipeline pairing physically-based ray-traced renderings of deformed glass with a parallel analytical forward model on the same distortion definition, producing ground-truth-paired test cases that let the Distortion Analyzer be designed, tuned, and validated against an independent prediction without physical experiments at every iteration.
3. Validation within the Virtual Pipeline. Validate the Distortion Analyzer against the analytical forward model on three criteria: spatial agreement, convergence under grid refinement, and stability under in-plane grid rotation. Hereafter, validation denotes this restricted sense; physical validation of the Virtual Pipeline itself is future work (§5).
4. Candidate industrial metrics. Propose physically interpretable distortion descriptors for industrial specification: sphere-equivalent optical power (area change) and cylinder-equivalent astigmatism (shape change), with diagnostic value that single-scalar systems lack. Promotion to deployed acceptance metrics depends on the physical validation and perceptual calibration identified as future work (§5).

Section 2 describes the Virtual Pipeline and Distortion Analyzer; Section 3 reports the three validation experiments; Section 4 discusses how the two-channel readout maps onto industrial use; Section 5 summarises the contributions and outlines future work.

2. The Virtual Pipeline

The Virtual Pipeline is the experimental framework this paper develops. It takes a programmatic Glass Distortion Definition (the geometry and deformation of a chosen glass) and processes it down two parallel branches whose outputs are compared (Figure 2). The theoretical branch feeds the definition into an Optical Theoretical Model, an analytical forward model of refraction through the glass, producing the distortion metrics it predicts. The synthetic-measurement branch passes the same definition through a physically-based ray-traced renderer (Blender Cycles, §2.1), generating the image

a real camera would capture; that image enters the Distortion Analyzer (§2.2), which produces the metrics measured from the image. The two image-formation routes — analytical Snell-law forward modelling and ray-traced rendering — share no common image-formation or demodulation pathway; for a like-for-like comparison, both outputs are then passed through the same post-projection evaluation stack (carrier-to-camera displacement recovery, detrending, smoothing, and metric computation), so the parity test in §3 is a virtual parity check on phase / displacement recovery rather than on the shared evaluation formulas themselves.

Two terms warrant disambiguation. “Virtual Pipeline” denotes this broader framework; “Distortion Analyzer” denotes the inner numerical method that takes images and returns metrics. From here on, the unqualified word “validation” means validation within the Virtual Pipeline: of the Distortion Analyzer against the Optical Theoretical Model, with both routes operating on synthetic inputs. Physical validation of the Virtual Pipeline itself is future work (§5).

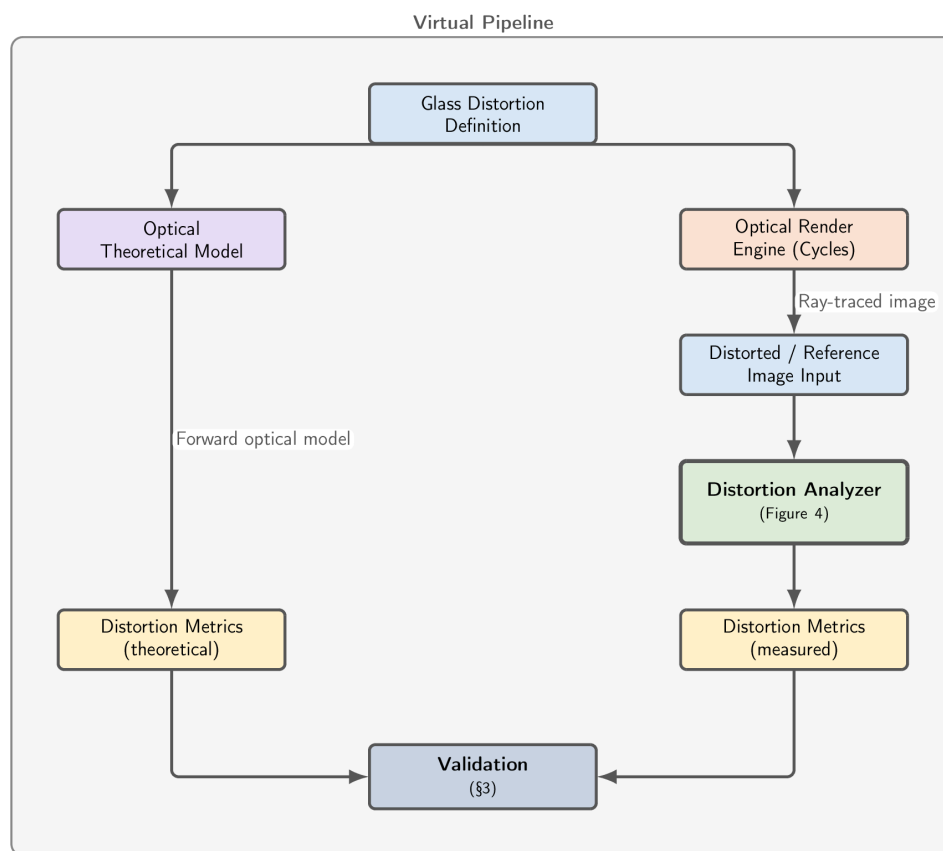


Fig. 2: The Virtual Pipeline. A single Glass Distortion Definition feeds two parallel branches: (left) an Optical Theoretical Model produces predicted distortion metrics; (right) an Optical Render Engine produces a Distorted / Reference Image Input that the Distortion Analyzer (Figure 4, detailed in §2.2) reduces to measured distortion metrics. The two image-formation routes are independent — analytical Snell-law forward modelling vs ray-traced rendering — and are then passed through a shared post-projection evaluation stack (A1, detrending, metric-anchored smoothing, A3) for a like-for-like comparison.

2.1. Generating distortion images

To develop and validate the Distortion Analyzer without the cost of physical experiments, the synthetic-measurement branch needs a controlled supply of physically realistic distorted grid images with known ground-truth geometry. Real 3D deformations are applied to a glass model, sinusoidal “roller waves” from thermal tempering and spherical “pillowing” from barometric pressure differentials in IGUs, and the scene is ray-traced in Blender’s Cycles path tracer (Figure 3), which enforces Snell’s law natively in a volumetric dielectric ($n=1.52$, used in §3.1). Distortions therefore arise from physical optics rather than 2D pixel warping.




Fig. 3: Synthetic setup in Blender. Top-left black object: camera. Orange-highlighted square: glass window. Furthest white square: background calibration grid. Geometry used in §3: glass-to-grid distance $L = 10$ m, glass thickness $d = 20$ mm, refractive index $n = 1.52$, narrow-tele camera intrinsics ($f = 300$ mm on a 36 mm sensor; horizontal FOV $\approx 6.9^\circ$).

2.2. Distortion Analyzer

The Distortion Analyzer (Figure 4) consumes a grayscale photograph of the printed line grid (with known pitch p_g) seen through the glass, the capture geometry that fixes physical scale (p_g , the glass-to-background distance L , and the mm-per-pixel of the analysis crop), a JSON-externalised set of algorithmic parameters, and, in the Comparative Method only, a separately captured glass-less reference. It returns a dense displacement field and the two-channel sphere/cylinder optical-power metrics. Block F handles phase recovery, Block A converts phase into physical channels, and Block I supplies the missing reference when needed.

The two routes differ in how $F2'$'s reference is sourced. Comparative runs $F \rightarrow A$ directly against the glass-less photograph and serves as the reference-anchored comparator within the Virtual Pipeline, but presumes that reference can be captured. Direct runs $F \rightarrow I \rightarrow F \rightarrow A$: F' 's first pass demodulates the distorted image against a synthetic flat carrier, yielding the unwrapped phase Block I needs; Block I infers a projective homography H_{opt} from that phase; F' 's second pass uses H_{opt} to synthesise a perspective-warped reference internally, and A processes its output. Block I therefore manufactures the missing reference from the distorted image alone, removing the dependence on a separate capture and making in-situ deployment on installed panels practical with the same protocol used at the factory line.

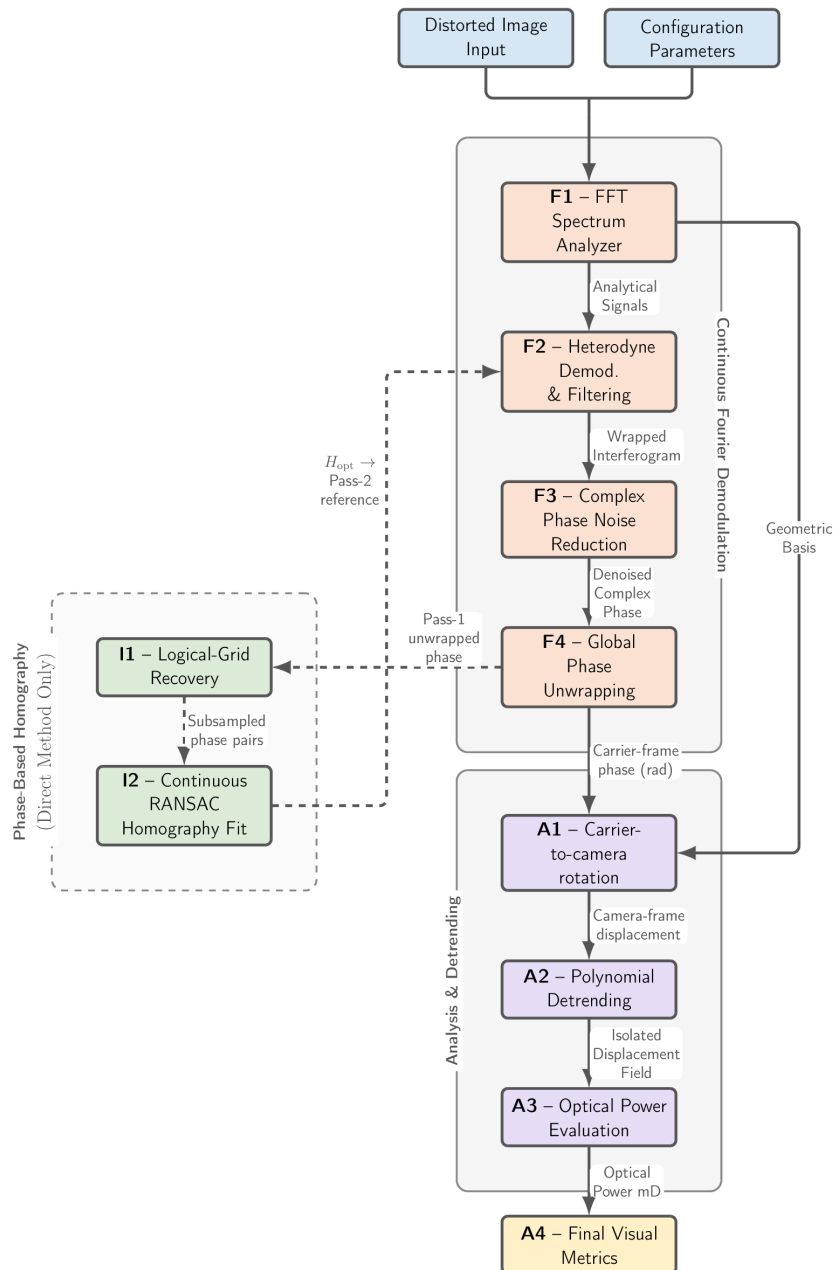


Fig. 4: Architecture of the Distortion Analyzer. The Comparative Method follows the central F1–F4 chain once, with F2’s reference supplied by a real glass-less photograph, then proceeds through A1–A4. The Direct Method (Block I, dashed branch) executes F2–F4 twice: pass 1 heterodynes against a synthetic flat carrier built from F1’s basis and feeds Block I, which infers H_{opt} ; pass 2 reuses F2–F4 with H_{opt} to synthesise a perspective-warped reference, and the pass-2 output of F4 feeds A1.

The pipeline-wide symbols used throughout §2.2 and §3 are summarised in Table 1.

Table 1: Notation. Pipeline-wide symbols used throughout §2.2 and §3.

| Symbol | Meaning | Units |
|---|---|----------------------|
| $\delta(\mathbf{x})$ | Displacement field (apparent shift of background features through the glass) | mm |
| $\delta_*(\mathbf{x})$ | Detrended residual (A2 output) | mm |
| $\Delta\theta_{\text{equiv}}(\mathbf{x})$ | Net angular deflection of the equivalent thin surface | rad |
| L | Glass-to-background distance | m |
| p_g | Physical grid pitch (printed) | mm |
| $\mathbf{f}_i, \mathbf{k}_i$ | Carrier i spatial frequency / wavevector ($\mathbf{k}_i = 2\pi\mathbf{f}_i$) | cycles/px, rad/px |
| $\hat{\mathbf{k}}_i$ | Unit vector $\mathbf{k}_i / \ \mathbf{k}_i\ $ | dimensionless |
| λ_i^{px} | Carrier wavelength in pixels, $1 / \ \mathbf{f}_i\ $ | px |
| H_{opt} | Inferred projective homography (logical grid \rightarrow camera pixels) | dimensionless |
| $M = [\hat{\mathbf{k}}_1 \mid \hat{\mathbf{k}}_2]$ | Unit-carrier basis matrix used for the A1 dual-basis inverse $(M^T)^{-1}$ | dimensionless |
| $\varepsilon = \text{sym}(\nabla\delta_*)$ | Virtual optical strain tensor | dimensionless |
| N | Total polynomial degree of the A2 detrend ($i + j \leq N$) | dimensionless |
| $\phi_{\text{eff}}, \phi_{\text{shear}}, \phi_{\text{total}}$ | Sphere-equivalent / cylinder-equivalent / Frobenius-norm channels (§2.2.3) | mD |
| ϕ_* | Stand-in for any of the three ϕ channels above | mD |
| σ | Smoothing kernel scale (per-stage subscripted; reporting scale §2.2.2) | px |
| $\sigma_{\text{target,mm}}$ | Physical reporting smoothing scale (§2.2.2); local pixel kernel is $\sigma_{\text{local}} = \sigma_{\text{target,mm}} / (\text{mm/px})$ | mm |

- **Coordinate frame:** $\mathbf{x} = (x, y)$ are camera-frame coordinates on the background-grid plane after A1, in mm; δ is the observed-minus-reference apparent feature displacement of the same coordinate frame.
- **Continuous Fourier demodulation (Block F):** Block F recovers, at every pixel, the displacement of the grid pattern caused by the glass. A Fourier route is natural because the printed grid is periodic: glass distortion appears as a local phase shift of two sinusoidal carriers along \mathbf{k}_1 and \mathbf{k}_2 , which Fourier methods extract densely and with sub-pixel precision. The technique adapts Fourier Transform Profilometry (FTP; Takeda et al. 1982; Su and Chen 2001) and the closely related phase-measuring deflectometry (Knauer et al. 2004; Huang et al. 2018) to architectural glass. Four stages carry it out.
 - **F1, FFT Spectrum Analyzer:** Locates the two principal grid wavevectors \mathbf{k}_1 and \mathbf{k}_2 in the image’s 2D Fourier spectrum and extracts a complex analytical signal for each. Perspective foreshortening and lens-induced refraction make the carriers non-orthogonal and unequal in magnitude, which is what the rest of the Distortion Analyzer measures. The strongest peak away from DC fixes \mathbf{k}_1 ; \mathbf{k}_2 is the strongest peak more than 45° from \mathbf{k}_1 at lower spatial frequency than the cross-modulation term $\mathbf{k}_1 + \mathbf{k}_2$. Both peaks are refined to sub-bin precision by parabolic interpolation (Jacobsen and Kootsookos 2007), without which a fractional-bin error would propagate as a residual phase ramp masquerading as a uniform tilt; each is then isolated by a 2D Gaussian bandpass and inverse-FFT’d.
 - **F2, Heterodyne Demodulation:** Each analytical signal has the form $A(x, y) e^{i(\mathbf{k}_i \cdot \mathbf{x} + \phi_{\text{dist},i}(\mathbf{x}))}$, a large linear carrier phase plus the small distortion term we want. Pointwise multiplication by

the complex conjugate of a reference signal at the same carrier subtracts the reference phase, cancels $\mathbf{k}_i \cdot \mathbf{x}$, and leaves the wrapped distortion phase $\varphi_{\text{dist},i}(\mathbf{x}) \in [-\pi, \pi]$. The reference is the analytical signal of an actual glass-less photograph in the Comparative Method; the Direct Method runs F2 twice, against a synthetic flat carrier in pass 1 and against a perspective-warped carrier from H_{opt} in pass 2 (§2.2.4).

- **F3, Complex Phase Noise Reduction:** The complex signal leaving F2 is high-amplitude on printed grid lines and near-zero in the interstitial cells, where the wrapped phase is dominated by noise. Smoothing the real and imaginary parts of the signal independently (rather than the wrapped phase itself, whose values $+\pi$ and $-\pi$ represent the same direction yet average to 0) effectively inherits weak-region phase from the well-conditioned grid lines around it. The kernel width $\sigma_{\text{ch},i} = \lambda_i^{\text{px}}/3$ matches the cutoff to the carrier scale, attenuating residual carrier and harmonic artefacts left by F2 while preserving the slowly-varying distortion phase.
- **F4, Global Phase Unwrapping:** Converts the wrapped phase into a continuous, single-valued displacement map by walking the pixel graph and integrating $\pm 2\pi$ jumps. We use quality-guided flood-fill (Herráez et al. 2002; van der Walt et al. 2014), seeded at the highest-reliability pixel and expanding along reliability-ranked paths so that error-prone regions are reached last. The unwrap requires the true pixel-to-pixel phase difference to remain below π , the Itoh (1982) sampling condition. With δ defined as the observed-minus-reference apparent feature displacement (in pixels at this stage, before A1 converts to mm), the carrier-frame relation is $\varphi_{\text{dist},i} = -\mathbf{k}_i \cdot \delta$, so the Itoh condition caps $|\Delta(\widehat{\mathbf{k}}_i \cdot \delta)| < \lambda_i^{\text{px}}/2$ per pixel: a per-pixel displacement-gradient bound, not on total displacement. Finer grids (smaller λ_i^{px}) therefore lower the unambiguous per-pixel displacement-gradient ceiling, traded off against finer spatial sampling — a tradeoff exercised in §3.2.
- **Analysis and detrending (Block A):** Block A turns the recovered displacement field into physically interpretable distortion metrics. We first describe the detrending operation, which is applied at various stages and then the four distinct pipeline stages: a basis change from grid to camera coordinates (A1), removal of slow non-distortion trends (A2), the sphere/cylinder optical-power channels (A3), and visualisation (A4):
 - **Metric-anchored smoothing:** After A1 produces the camera-frame displacement field and A2 detrends it, and before A3 differentiates the result, the detrended field is smoothed once. The reporting bandwidth is conceptually specified by a physical scale $\sigma_{\text{target,mm}}$, with each system computing $\sigma_{\text{local}} = \sigma_{\text{target,mm}}/(\text{local mm per pixel})$: the optical pipeline derives mm-per-pixel from the FFT-extracted carrier wavelength against the known pitch, the theoretical evaluation grid from physical pane width over grid resolution. Statistics derived from the same physical low-pass filter are camera-FOV-invariant, a prerequisite for the §3.1 parity comparison and for cross-installation reporting. In the present renderings the pixel equivalent $\sigma_{\text{local}} = 40$ px was supplied to the analyzer directly for the shared camera geometry; for any other installation σ_{local} should be recomputed from the chosen $\sigma_{\text{target,mm}}$ and the local mm-per-pixel.
 - **A1, Carrier-to-camera rotation:** Assembles the two scalar phase maps from F4, each giving δ projected onto one carrier direction $\widehat{\mathbf{k}}_i$, into a single 2D field $\delta(x, y)$ in camera-frame millimetres. Two foreshortening effects must be handled cleanly. The carriers are non-orthogonal in the image, so the maps are oblique projections of δ requiring the dual-basis inverse rather than a plain rotation. They also carry different mm-per-pixel scales, and entangling those with the basis change would re-emerge in A3 as spurious anisotropy in the strain tensor and corrupt ϕ_{eff} . A1 therefore strips the per-channel scale, applies the dual-basis

inverse $(M^T)^{-1}$ with $M = [\hat{\mathbf{k}}_1 \mid \hat{\mathbf{k}}_2]$, and restores an isotropic mm-per-pixel scale $p_g \sqrt{\|\mathbf{f}_1\| \|\mathbf{f}_2\|}$, the geometric mean of the two per-channel scales (treated here as a reporting convention; the long-form companion paper supplies the perspective derivation).

- **A2, Polynomial Detrending:** Removes slow image-wide trends from δ before A3 differentiates it, by fitting and subtracting a least-squares polynomial of total degree N (i.e. monomials with $i + j \leq N$, applied componentwise to δ_x and δ_y). The choice is geometric: $N = 2$ captures tilt and mild quadratic bowing (residual perspective), $N = 4$ also handles pillowing-style domes whose curvature varies across the pane. Higher degrees are avoided because the polynomial basis at degree N resolves wavelengths down to roughly W/N , with W the fitted field extent, and on fine grids the actual ripples are short enough that the polynomial would fit them rather than merely the trend. The operational consequence is that the present Distortion Analyzer characterises only the short-wavelength component of distortion, ripples and localised defects on top of the global pane shape. This is a deliberate scope choice: the slow component conflates two genuinely different things, residual perspective foreshortening (a measurement artefact) and bulk pane curvature (real glass deformation), which a single image cannot distinguish without independent perspective calibration or multi-view capture. Figure 5 illustrates A2 on δ_x for a strongly-domed case before and after subtracting a degree-4 polynomial. Disentangling the two slow components, and reporting bulk curvature as its own channel, is identified as future work.
- **A3, Optical Power Evaluation is the headline metric stage:** Centralised finite differences of the detrended δ_* form the symmetric virtual optical strain tensor $\varepsilon = \text{sym}(\nabla\delta_*)$; symmetrising drops the antisymmetric part, which represents rigid local rotation and changes neither shape nor area. Three invariants of ε are extracted and divided by the glass-to-background distance L :
 - Sphere-equivalent optical power, the local divergence of δ_* , an area-change field, the optical signature of a spherical lens (the headline metric):

$$\phi_{eff} = \frac{\text{tr}(\varepsilon)}{L} = \frac{\varepsilon_{xx} + \varepsilon_{yy}}{L}.$$

- Cylinder-equivalent astigmatism, twice the maximum shear, a shape-change field at constant area:

$$\phi_{shear} = \frac{2}{L} \sqrt{\left(\frac{\varepsilon_{xx} - \varepsilon_{yy}}{2}\right)^2 + \varepsilon_{xy}^2}.$$

ϕ_{shear} is non-negative by construction; the principal-axis orientation, available from the eigenvectors of ε , is not used in the scalar summaries reported here.

- Total local distortion magnitude, a basis-invariant Frobenius-norm summary reported in §3.1:

$$\phi_{total} = \frac{\|\varepsilon\|_F}{L}.$$

In what follows, ϕ_* stands for any of the three when a statement applies to all. Units. δ_* is in millimetres and L in metres, so the formulas evaluate to inverse metres (diopters); reported values and colour bars are scaled to mD. Sphere-equivalent convention. ϕ_{eff} is the full

divergence $\text{tr}(\varepsilon)/L$, not the half-trace used for the ophthalmic spherical-equivalent diopter. With δ_m defined as the observed-minus-reference apparent displacement on the background-grid plane (in metres for this paragraph), and using the standard ray-angle convention for a weak positive spherical lens, $\Delta\theta = -P \mathbf{r}$, we take $\delta_m = L_m \Delta\theta$. Hence $\nabla \cdot \delta_m = -2PL_m$, so $\phi_{\text{eff}} = -2P$ and $P_{\text{sph,eq}} = -\phi_{\text{eff}}/2$. Reversing the displacement convention reverses the reported sign. The reported mD values follow dimensionally from differentiating δ_{mm} with respect to mm coordinates and dividing by L_m , with $\delta_m = 10^{-3} \delta_{\text{mm}}$. What ϕ_{eff} measures. Gradients in A3 are taken with respect to the camera-frame coordinates of the background grid plane, so ϕ_{eff} reports the perceived through-system optical power at the grid plane rather than the intrinsic curvature of the glass surfaces; the long-form paper expands the geometric chain that relates the two.

- **A4, Final Visual Metrics:** standardised heatmaps, threshold masks, and statistical summaries.




Fig. 5: A2 on a strongly-domed illustrative case. (a) camera-frame δ_x before detrending; (b) the same field after subtracting a degree-4 polynomial surface, leaving localised diagonal ripple. The metrics of A3 are computed on (b), so what they characterise is deviation from the smooth global trend.

- **Phase-based homography inference (Block I, Direct Method).** In the Direct Method, F2 needs a clean reference carrier and no glass-less photograph provides one. Block I manufactures it from the distorted image alone, exploiting a separation the Fourier route makes geometric: distortion wobbles the local phase of each carrier but leaves the global projective foreshortening (how grid lines converge under camera tilt) substantially intact.
 - I1, Logical-Grid Recovery. Turns F4's pass-1 unwrapped phase into a dense set of (pixel, logical-grid) point pairs. Because pass 1 heterodyned against a synthetic flat carrier built from F1's basis, F4's pass-1 output carries the residual phase only, projective foreshortening plus local distortion wobble, with the periodic carrier already mathematically subtracted. Adding the synthetic-carrier phase back at every pixel and normalising by 2π recovers $(U, V)(\mathbf{x}) = \frac{1}{2\pi} (\mathbf{k}_1 \cdot \mathbf{x} + \varphi_1^{\text{unwr}}, \mathbf{k}_2 \cdot \mathbf{x} + \varphi_2^{\text{unwr}})$, the logical grid coordinate that an undistorted pinhole capture would assign. The field is subsampled at one point per carrier period.
 - I2, Continuous RANSAC Homography Fit. Runs RANSAC homography estimation (Fischler and Bolles 1981) on those pairs to recover the 3×3 projective homography H_{opt} from logical grid coordinates to camera pixels, the most general transformation a pinhole camera applies to a flat target. The threshold is expressed as a fraction of the carrier pitch, so pairs sitting on the global projective plane fall within it while points displaced by the optical wave fall outside. The inliers therefore define the perspective baseline, and the wave-displaced outliers are precisely the signal the Direct Method aims to measure. H_{opt} then returns to F2, which inverts it to

back-project every camera pixel into its ideal logical-grid coordinates and synthesises a perspective-warped reference matching the actual image's foreshortening; the pass-2 heterodyne cancels the carrier far more cleanly than pass 1's flat reference. The cost of reference-free inference shows up here: RANSAC sees phase points that already contain the slow component of the optical wave, so H_{opt} absorbs a small fraction of that wave alongside the genuine perspective foreshortening. On biconvex this manifests as a roughly 5–9% method-dependent magnitude offset between Direct and Comparative on $|\overline{\phi_{\text{eff}}}|$ (§3.2). The Comparative Method incurs no such absorption, which is why §3 uses it as the reference-anchored comparator.

3. Numerical experiments

This section validates the Distortion Analyzer within the Virtual Pipeline of §2. By virtual experiment we mean a controlled comparison in which the Analyzer is checked against the Pipeline's independent analytical forward model on inputs whose ground truth is known by construction. The two routes share only the glass-distortion definition; no intermediate signal and no common image-formation or demodulation pathway up to the shared post-projection evaluation stack, so spatial agreement constitutes a virtual parity check on phase / displacement recovery rather than internal consistency-checking. Real illumination, sensor noise, printing tolerances, and reflections are not yet exercised, and physical validation of the Virtual Pipeline itself remains future work. The image-processing stages are made deliberately harder by framing the physical edge of the glass pane in the middle of every captured image. §3.1 reports a parity check against the analytical model, §3.2 a sweep over grid density, and §3.3 stability under in-plane rotation of the calibration grid. Unless stated otherwise, results use the Direct Method; the Comparative Method serves as the reference-anchored comparator within the Virtual Pipeline.

3.1. Numerical parity verification

For The Distortion Analyzer is compared, on a controlled-physics Blender scene with a programmatic 1D sinusoidal displacement on a parallel-faced glass slab, against an independent analytical forward model. The model is built from the rendered topography $Z(x, y)$, refractive index n , the standard geometrical-optics treatment of slab refraction (Hecht 2017, §4.4 *Refraction* and §5.5 *Prisms*; vectorial Snell's law in the form of Spencer and Murty 1962, linearized here in the surface-slope perturbation), and the same pinhole camera projection (focal length, sensor, pose) Cycles used to render. The two image-formation routes — analytical forward modelling and ray-traced rendering — are independent up to the comparison stage; both outputs are then passed through the same post-projection evaluation stack, so the parity test below validates the image-based recovery against an independently generated optical prediction rather than the shared evaluation formulas themselves.

The reference solution couples three first-order components from Snell's law on a parallel-faced slab of thickness d :

- **Slab lateral shift** (Hecht 2017, §4.4). A ray at incidence α to the slab normal exits parallel to entry, shifted by $s(\alpha) = d \sin(\alpha - \beta) / \cos \beta$ with $\sin \beta = \sin \alpha / n$ and d the physical slab thickness. For surface tilt ∇Z , the perturbation $\mathbf{u}_{\text{slab}} = (\partial s / \partial \alpha) \nabla Z$ acquires distinct longitudinal and transverse prefactors at oblique incidence ($\partial s / \partial \alpha$ and $d(1 - \cos \alpha / (n \cos \beta))$), both reducing to the textbook $d(1 - 1/n) \nabla Z$ as $\alpha \rightarrow 0$.
- **Thin-prism deflection** from a non-conformal back face (Hecht 2017, §5.5; Spencer and Murty 1962). When the back-face tilt differs from the front by $\Delta(\nabla Z) \approx \mathbf{H}_Z \Delta \mathbf{h}$ — with $\mathbf{H}_Z = \nabla \nabla Z$ the Hessian of Z and $\Delta \mathbf{h} = d \tan \beta \hat{\mathbf{t}}$ the in-glass lateral travel — the exit ray deflects by $\Delta \boldsymbol{\theta} = -K_0 \Delta(\nabla Z)$, where $K_0 = n \cos \beta - \cos \alpha$ is the oblique-Snell amplification factor (reducing to $n - 1$ at $\alpha = 0$); the negative sign is consistent with the displacement convention of §2.2.3. Propagation L in air contributes $\mathbf{u}_{\text{trav}} = L \Delta \boldsymbol{\theta}$, with the appropriate $\sec \alpha / \sec^2 \alpha$ projection factors onto the background-grid plane (full derivation in the long-form paper).
- **Pinhole camera projection** through the same camera matrix (focal length, sensor, pose) Cycles used to render.

Both displacement components are linear in the surface derivatives, so the prediction extends to any harmonic superposition. The rendered topography is captured to $< 0.2\%$ residual by three Fourier harmonics ($k, 2k, 3k$; the higher harmonics arise from the Blender Displace texture-sampling interaction), and the prediction is their linear superposition. The linearization is verified against full vector-Snell ray-tracing at sampled incidences from 0° to 45° , to better than 10^{-6} .

The same evaluation stack is then applied on both sides so the displayed ϕ_* maps reflect the analyzer-detectable component: per-carrier complex smoothing of the analytical signals before unwrapping ($\sigma_{\text{ch},i} = \lambda_i^{\text{px}} / 3$), A1 carrier-to-camera displacement recovery, polynomial detrending ($N = 0$ here, since the wave is the signal), the metric-anchored pre-strain Gaussian (§2.2.2), and A3's finite-difference metric evaluation. Figure 6 shows the agreement in spatial structure, both routes resolving the wave's fundamental and second-harmonic banding at the correct wavelengths with a uniform amplitude offset. The matching analytical magnitudes for ϕ_{eff} and ϕ_{shear} (both 7.57 mD) follow from the geometry: a 1D wave produces $\varepsilon_{xx} \neq 0$ with $\varepsilon_{yy} = \varepsilon_{xy} = 0$, so the two channels carry the same magnitude by construction.




Fig. 6: Top row, reference and distorted captures of the wavy slab, with the analysis region overlaid in red. Middle row, Analyzer (Direct Method, left) vs analytical forward-model (right) ϕ_{eff} . Bottom row, same comparison for ϕ_{shear} . Panel titles report $|\overline{\phi_*}|$ over the displayed valid analysis field. Wavy-slab controlled-physics case: 1D sinusoidal $Z(x, y)$, $n = 1.52$, $d = 20$ mm, $L = 10$ m; per row the two panels share the same colour limits.

Table 2: Distortion Analyzer-to-analytical magnitude ratios (wavy-slab case, Direct Method). The three ratios are equal by geometry: a 1D wave produces $\varepsilon_{xx} \neq 0$ with $\varepsilon_{yy} = \varepsilon_{xy} = 0$, so all three scalars carry the same magnitude on both sides.

Procrustes alignment of recovered (u_x, u_y) to analytical (u_x, u_y) : rotation $\theta = +0.10^\circ$, isotropic scale $s = 1.05$.

| Channel | Analyzer / Analytical | Notes |
|-----------------------|-----------------------|---|
| ϕ_{eff} | 0.963 | sphere-equivalent full-trace channel, $\phi_{\text{eff}} = \text{tr}(\varepsilon)/L$ |
| ϕ_{shear} | 0.963 | cylinder-equivalent astigmatism, $\phi_{\text{shear}} = (2/L)\sqrt{((\varepsilon_{xx} - \varepsilon_{yy})/2)^2 + \varepsilon_{xy}^2}$ |
| ϕ_{total} | 0.963 | basis-invariant Frobenius-norm summary, $\phi_{\text{total}} = \ \varepsilon\ _F/L$ |

All three scalars agree with the analytical prediction within approximately 4% (ϕ_{eff} at 3.7%), and the recovered displacement frame is aligned with the analytical to within 0.1° . This is achieved with an anisotropic-Gaussian bandpass around each FFT carrier, with the narrow axis oriented towards the bilinear cross-term $\pm(\mathbf{k}_1 + \mathbf{k}_2)$ to suppress its leakage while the wide perpendicular axis preserves

the wave sidebands. Without this directional shaping the cross-term leaks into the heterodyne phase of the k_2 channel, producing a uniform approximately 50° rotation between the recovered and analytical frames (ϕ_{shear} and ϕ_{total} insensitive to this frame rotation; ϕ_{eff} projected by $\cos\theta$). The Analyzer therefore correctly recovers the spatial structure of the strain field on all three scalars; the residual approximately 4% absolute-calibration offset is consistent with the reporting smoothing scale ($\sigma_{\text{local}} = 40 \text{ px}$ in this rendering; see §2.2.2) and is treated as a common reporting-scale uncertainty in the relative comparisons of §3.2 and §3.3.

3.2. Grid size dependence

The second virtual experiment tests whether the recovered optical power is stable across the grid pitches an inspector might reasonably print, since available wall space may dictate a coarser grid than is ideal. The test case is a biconvex glass lens imaged from a slanted viewpoint at $L = 10 \text{ m}$ against three pitches: coarse (96.9 mm), medium (64.6 mm), and fine (48.4 mm); the polynomial detrend is set to $N = 2$ here, removing residual perspective and bulk lens curvature so the reported ϕ_* are the analyzer-detectable short-wavelength residuals. With the Direct Method, $|\overline{\phi_{\text{eff}}}|$ holds to within $\pm 4\%$ of its mean across the three pitches (8.42 \rightarrow 8.82 \rightarrow 8.97 mD), with the Comparative Method tracking it (7.84 \rightarrow 8.06 \rightarrow 8.51 mD); the two methods agree within approximately 5–9% on every grid (coarse 6.9%, medium 8.6%, fine 5.1%, with relative difference defined as $|a - b|/\max(a, b)$). The Direct Method therefore delivers magnitudes consistent with the reference-anchored Comparative Method without an unobstructed reference, the regime real installations impose. Absolute calibration is established separately by §3.1.




Fig. 7: Rendered-input mosaic across coarse / medium / fine grid pitches.

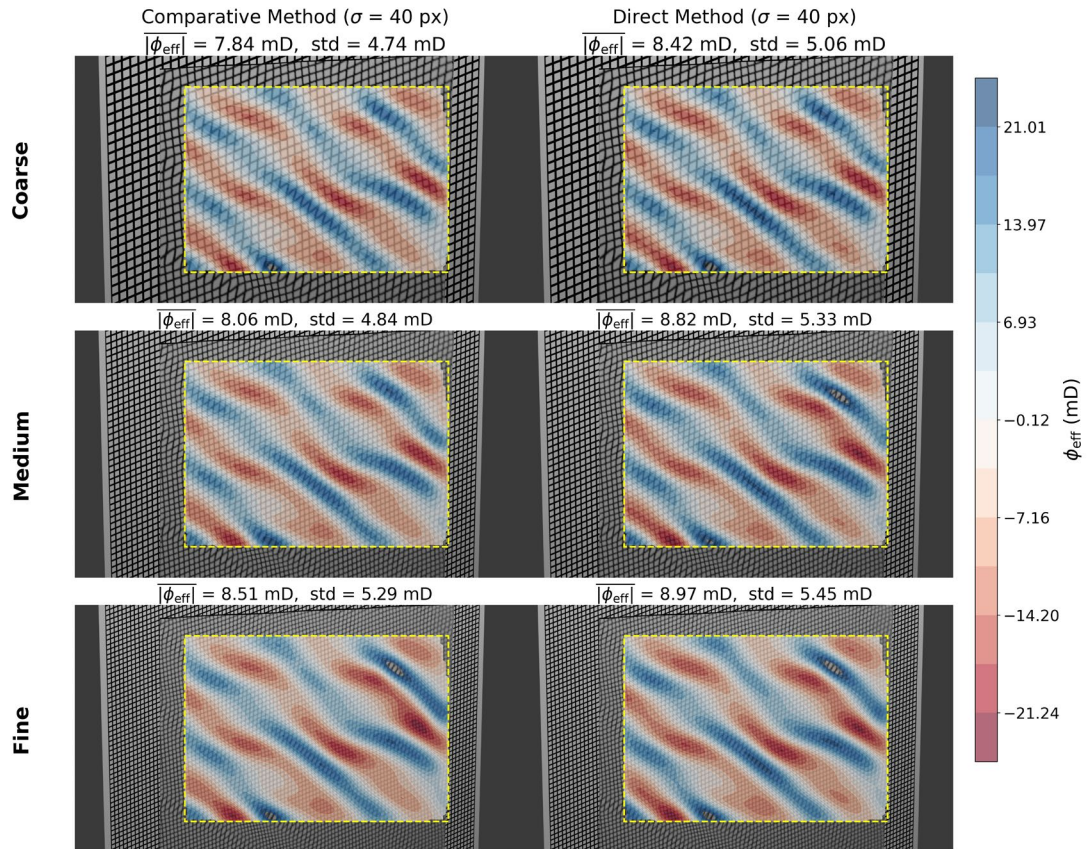



Fig. 8: Power mosaic illustrating grid-size dependence (Comparative Method on the left, Direct Method on the right, slanted viewpoint). Panel titles report $\overline{|\phi_{\text{eff}}|}$ and the standard deviation of signed ϕ_{eff} over the displayed valid analysis field.

3.3. Grid size dependence

The third virtual experiment exercises the inspection-geometry variable any field deployment introduces beyond grid pitch: the in-plane orientation of the calibration grid. The test case is a separate flat-slab scene with mild localised distortion (different glass from §3.2, hence smaller absolute magnitudes), imaged under a near-orthographic view at $L = 10$ m against the fine 48.4 mm pitch (polynomial detrend $N = 2$ as in §3.2). Because the 2D FFT explicitly isolates $(\mathbf{k}_1, \mathbf{k}_2)$ regardless of their orientation, the demodulated phase fields are stable under in-plane rotation: the same panel with the background grid axis-aligned and then rotated about its own normal returns matching $\overline{|\phi_{\text{eff}}|}$ to within approximately 5% (1.06 vs 1.01 mD, Figure 9). In the grid-registered display, the recovered ϕ_{eff} pattern rotates rigidly with the grid, as it must, and the residual sits inside the magnitude variability of the §3.2 sweep. A small unavoidable wobble in how an inspector aligns the calibration grid therefore remains at the few-percent level in the reported metric, and this points toward a multi-position acceptance protocol that averages over the residual geometric variability rather than a single in-situ snapshot.



abs_mean ratio (min/max) = 0.952 relative difference = 4.8%

Fig. 9: Grid rotation invariance check (Direct Method ϕ_{eff}). The same panel is imaged under a near-orthographic view with the background grid axis-aligned (top) and rotated about its own normal (bottom); panel titles report $\overline{|\phi_{\text{eff}}|}$ and the standard deviation of signed ϕ_{eff} over the displayed valid analysis field, and the integrated-magnitude ratio is annotated below.

4. Practical use of the two-channel readout

§3.1–3.3 thus mark the boundary of the present virtual experiment: spatial agreement with an independent optical-theory model, stability under grid refinement, and stability under in-plane rotation. Real illumination, sensor noise, printing tolerances, reflections, and ambient-light contamination are the natural next-phase questions, now made tractable by the same synthetic environment that supplied the present validation.

§3 established that, within the Virtual Pipeline, the Distortion Analyzer agrees with the analytical forward model and produces numerically stable optical-power fields. The headline sphere-equivalent (full-trace) channel ϕ_{eff} , defined in §2.2.3, is

$$\phi_{\text{eff}} = \frac{\nabla \cdot \delta_*}{L},$$

with the trace and unit conventions defined in §2.2.3. The channel is locally additive (the divergence is the first-order local area-change term), and the $1/L$ scaling removes the first-order dependence on background distance under the assumed paraxial geometry. The measured field remains protocol- and viewpoint-dependent through detrending scope, ROI selection, smoothing scale, and viewing angle; a fully panel-intrinsic metric is future work (§5). Combined with the cylinder-equivalent companion ϕ_{shear} , the Distortion Analyzer emits a two-channel readout that maps onto three concrete industrial

applications spanning the full life cycle of a glass pane: factory-line acceptance and root-cause routing during manufacture, design-time and contractual specification at procurement, and in-situ dispute resolution after installation. The same Distortion Analyzer is used in all three; only the reference branch differs (Comparative at the factory, Direct in the field).

4.1. Factory-line acceptance with multi-channel reporting

Routing a manufacturing root cause to a specific line component requires reporting which kind of distortion has failed, not merely that the pane has failed. Existing factory-line systems often prioritise high-throughput inspection and scalar acceptance summaries on monolithic flat glass; mD-reporting scanners adapted from automotive (e.g. SYNERGX WINDX) are reviewed against architectural standards in our GPD2025 precursor (Moreau, Betanzos and Casas Gonzalez 2025), which demonstrates a measurable correlation between architectural grid-based and automotive scanner readings on the same laminate. The two-channel decomposition introduced here is instead intended to separate area-change and shape-change signatures: $\phi_{\text{eff}} = \text{tr}(\varepsilon)/L$ tracks local area change and goes silent on any deformation that stretches one grid axis as much as it compresses the orthogonal one, while ϕ_{shear} tracks shape change and is non-zero precisely there. ϕ_{shear} excursions may therefore suggest furnace-roller wear or roller-bed misalignment, while broadly isotropic ϕ_{eff} patterns are more consistent with quench non-uniformity or asymmetric thermal loading. The factory line is also the natural setting for the Comparative Method, since a glass-less reference capture is straightforward to obtain there; extending §3.2's Direct-vs-Comparative comparison to the cylinder channel, the two methods agree on ϕ_{shear} at the same 5–10% level seen on ϕ_{eff} (Figure 10, coarse-grid case).

4.2. Design-time and contractual specification

ϕ_{eff} provides a physically-grounded candidate metric family for procurement and acceptance criteria. Combined with a standardised capture protocol (in-plane grid-orientation robustness exercised in §3.3; the full protocol identified in §5 as deployment future work), it offers a defensible objective basis for architect-fabricator agreement, in place of qualitative contract language, once acceptance thresholds are set by physical validation. The same two-channel readout used at the factory line carries through to the contractual document: a pane that passes a ϕ_{shear} -bounded acceptance criterion at manufacture is reported on the same scalar axis the architect-fabricator contract is written on.

4.3. In-situ dispute resolution

Once installed, a glass-less reference capture is no longer practicable; the Direct Method takes over. A single image of a printed grid placed behind the installed lite, requiring no panel removal and no separately captured baseline, recovers an objective, reproducible record of the as-installed optical state via the Direct Method's ϕ_{eff} map. Because the recovered fields use the same channels as the factory acceptance and the contractual document, the comparison across the life cycle is direct: a matched synthetic capture under the same nominal viewing geometry can help distinguish whether a failed acceptance criterion is consistent with a mechanical defect of the glass or with geometry-induced amplification of the observed distortion, a distinction presently almost impossible to settle in the field.

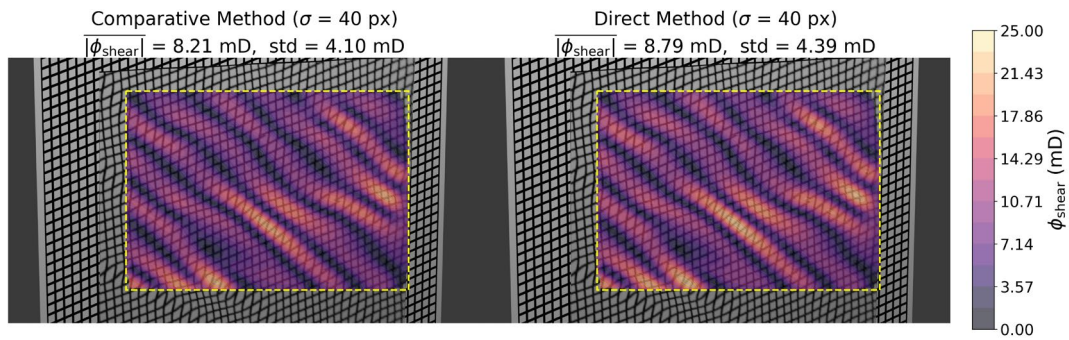


Fig. 10: Shear / astigmatism panel on the coarse-grid biconvex case ($\phi_{\text{shear}} = 2 \cdot \text{max shear}$, the cylinder-equivalent). Left: Comparative Method. Right: Direct Method. Panel titles report $|\overline{\phi_{\text{shear}}}|$ and the standard deviation of signed ϕ_{shear} over the displayed valid analysis field. The diagonal shape-distortion banding is recovered consistently between the two methods, confirming that for a panel with appreciable astigmatic content the cylinder channel carries genuinely independent information rather than a redundant rescaling of ϕ_{eff} .

A note on tolerance. The percentage disagreements reported across §3.1–§3.3 and above sit in the 5–10% range. This would be substantial in classical dimensional metrology, but ϕ_{eff} and ϕ_{shear} are intended as visual-perception channels, so the relevant tolerance is perceptual rather than dimensional. Whether a 5–10% pane-average difference is visually distinguishable awaits empirical perceptual calibration (e.g. a 1-to-5 severity scale), identified as future work (§5) and treated in detail in the long-form companion paper.

5. Conclusion

Architectural glass distortion is currently evaluated by manufacturers through surface topography measurements, which capture the flatness of the surface or mechanical factors such as roller waves and edge effects. Our eye perception of this parameter works only in reflection.

Our goal is to quantify the visual impression that occurs when looking through glass. The virtual pipeline provides us with a rapid development environment for evaluating optical see-through distortion.

The methodology measures distortion through the channel observers actually use: the apparent displacement of a background pattern seen through the glass. It reads this out as sphere-equivalent and cylinder-equivalent fields on an equivalent thin surface. In the paraxial regime, this surface collapses IGU, laminate, and coated-system complexity into one perceptually-meaningful descriptor. The readout characterises the detrended local component (ripples and localised defects). At the factory line, the Comparative Method exploits the available glass-less reference capture; in the field, the reference-free Direct Method enables single-image in-situ analysis on installed panels. The same Distortion Analyzer and the same two reported channels span both contexts, so an acceptance criterion set at the factory carries over directly to a post-installation dispute.

Within the Virtual Pipeline, the Distortion Analyzer agrees with an independent analytical forward model on the same glass-distortion definition: it agrees in spatial structure and remains stable across grid pitches and under in-plane rotation of the calibration grid. The two routes share no common image-formation or demodulation pathway, so this is a virtual parity check on phase / displacement recovery rather than internal consistency-checking; physical validation against real glass and

perceptual calibration of the metric channels are the natural next steps rather than claims resolved here.

The sphere/cylinder decomposition provides two physically interpretable channels as candidate descriptors for procurement, factory-line acceptance, and post-installation dispute resolution. Promotion to deployed acceptance metrics depends on the physical validation and perceptual calibration identified above; the underlying analysis does not change as those calibrations accumulate.

Several deployment pieces remain: physical validation against real glass with independent geometry and through-glass reference measurements, statistical characterisation of real-world defects, perceptual anchoring through visual studies, separation of bulk curvature from residual perspective so it can be reported as its own channel, and the move from viewpoint-dependent to panel-intrinsic metrics. The present results suggest that these extensions may be compatible with consumer-grade imaging and offline single-image analysis, but this remains to be demonstrated under field lighting, reflections, printing tolerances, and installation constraints; the remaining work is therefore primarily one of validation, calibration, protocol design, and field engineering.

References

- ANSI/SAE: ANSI/SAE Z26.1 American National Standard for Safety Glazing Materials for Glazing Motor Vehicles and Motor Vehicle Equipment Operating on Land Highways. SAE International, Warrendale, PA (1996)
- ASTM International: ASTM C1036-21 Standard Specification for Flat Glass. ASTM International, West Conshohocken, PA (2021). <https://doi.org/10.1520/C1036-21>
- ASTM International: ASTM C1048-18 Standard Specification for Heat-Strengthened and Fully Tempered Flat Glass. ASTM International, West Conshohocken, PA (2018)
- ASTM International: ASTM C1651-11(2018) Standard Test Method for Measurement of Roll Wave Optical Distortion in Heat-Treated Flat Glass. ASTM International, West Conshohocken, PA (2018). <https://doi.org/10.1520/C1651-11R18>
- ASTM International: ASTM C1652/C1652M-21 Standard Test Method for Measuring Optical Distortion in Flat Glass Products Using Digital Photography of Grids. ASTM International, West Conshohocken, PA (2020). https://doi.org/10.1520/C1652_C1652M-21
- Atchison, D.A., Smith, G.: Optics of the Human Eye. Butterworth-Heinemann, Oxford (2000)
- Duda, R.O., Hart, P.E.: Use of the Hough transformation to detect lines and curves in pictures. *Commun. ACM* 15(1), 11–15 (1972). <https://doi.org/10.1145/361237.361242>
- Efferz, L., Dix, S., Schuler, C., Kolling, S.: Optical anisotropy effects in laminated tempered glass. **Glass Struct. Eng.** 10(3) (2025). <https://doi.org/10.1007/s40940-024-00285-w>
- European Committee for Standardization: EN 572 Glass in Building — Basic Soda Lime Silicate Glass Products (Parts 1–9). CEN, Brussels (2012, with amendments e.g. 2016 for Part 1).
- European Committee for Standardization: EN 12150-1:2015 Glass in Building — Thermally Toughened Soda Lime Silicate Safety Glass. CEN, Brussels (2015)
- European Committee for Standardization: EN 12150-2:2004 Glass in Building — Thermally Toughened Soda Lime Silicate Safety Glass — Part 2: Evaluation of Conformity / Product Standard. CEN, Brussels (2004)
- European Committee for Standardization: EN 14179-1:2016 Glass in Building — Heat-Soaked Thermally Toughened Soda Lime Silicate Safety Glass — Part 1: Definition and Description. CEN, Brussels (2016)
- Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24(6), 381–395 (1981). <https://doi.org/10.1145/358669.358692>
- Hecht, E.: Optics, 5th edn. Pearson, Boston, MA (2017)
- Henriksen, T., Stokes, E., Louter, C., Overend, M.: Optical distortions in architectural glass: review of categorization, evaluation and measurement methods. **Glass Struct. Eng.** (2026). <https://doi.org/10.1007/s40940-025-00314-2>

- Herráez, M.A., Burton, D.R., Lalor, M.J., Gdeisat, M.A.: Fast two-dimensional phase-unwrapping algorithm based on sorting by reliability following a noncontinuous path. *Appl. Opt.* 41(35), 7437–7444 (2002). <https://doi.org/10.1364/AO.41.007437>
- Huang, L., Idir, M., Zuo, C., Asundi, A.: Review of phase measuring deflectometry. *Opt. Lasers Eng.* 107, 247–257 (2018). <https://doi.org/10.1016/j.optlaseng.2018.03.026>
- Illguth, M., Schuler, C., Bucak, Ö.: The effect of optical anisotropies on building glass façades and its measurement methods. *Front. Archit. Res.* 4(2), 119–126 (2015). <https://doi.org/10.1016/j.foar.2015.01.004>
- International Organization for Standardization: ISO 3538:1997 Road Vehicles — Safety Glazing Materials — Test Methods for Optical Properties. ISO, Geneva (1997)
- Itoh, K.: Analysis of the phase unwrapping algorithm. *Appl. Opt.* 21(14), 2470 (1982). <https://doi.org/10.1364/AO.21.002470>
- Jacobsen, E., Kootsookos, P.: Fast, accurate frequency estimators. *IEEE Signal Process. Mag.* 24(3), 123–125 (2007). <https://doi.org/10.1109/MSP.2007.361611>
- Knauer, M.C., Kaminski, J., Häusler, G.: Phase measuring deflectometry: a new approach to measure specular free-form surfaces. *Proc. SPIE* 5457, 366–376 (2004). <https://doi.org/10.1117/12.545704>
- Moreau, L.: Optical and structural developments in air traffic control tower glazing. In: Louter, C., Bos, F., Belis, J. (eds.) **Challenging Glass Conference Proceedings**, vol. 9. TU Delft, Delft (2024). <https://doi.org/10.47982/cgc.9.624>
- Moreau, L., Betanzos, A., Casas Gonzalez, G.: Measuring see-through distortion. In: **Glass Performance Days 2025 Conference Proceedings**, Tampere (2025).
- Spencer, G.H., Murty, M.V.R.K.: General ray-tracing procedure. *J. Opt. Soc. Am.* 52(6), 672–678 (1962). <https://doi.org/10.1364/JOSA.52.000672>
- Su, X., Chen, W.: Fourier transform profilometry: a review. *Opt. Lasers Eng.* 35(5), 263–284 (2001). [https://doi.org/10.1016/S0143-8166\(01\)00023-9](https://doi.org/10.1016/S0143-8166(01)00023-9)
- Takeda, M., Ina, H., Kobayashi, S.: Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *J. Opt. Soc. Am.* 72(1), 156–160 (1982). <https://doi.org/10.1364/JOSA.72.000156>
- United Nations Economic Commission for Europe: UN Regulation No. 43 — Uniform Provisions Concerning the Approval of Safety Glazing Materials and Their Installation on Vehicles, Revision 4. UNECE, Geneva (2017)
- van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T.: scikit-image: image processing in Python. *PeerJ* 2, e453 (2014). <https://doi.org/10.7717/peerj.453>


Platinum Sponsor



Gold Sponsors



Silver Sponsors



Organisation

