Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Decoupled Camera Intrinsics Calibration via a Binary Hexagonal Galois Field Pattern

Gennadiy Kis

gennadiykiss@gmail.com

Oleksandr Kis

State University of Information and Communication Technologies

Research Article

Keywords:

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11345316/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Decoupled Camera Intrinsics Calibration via a Binary Hexagonal Galois Field Pattern

Gennadiy Kis[1*] and Oleksandr Kis[2]

[1]Independent Researcher, Kyiv, Ukraine.
[2]Information Systems and Technologies, State University of Information and Communication Technologies, Solomyanska Str. 7, Kyiv, 03110, Ukraine.

*Corresponding author(s). E-mail(s): gennadiykiss@gmail.com;
Contributing authors: o.kis@duikt.edu.ua;

Keywords: camera calibration, Binary Hexagonal Galois Pattern, m-sequences, algebraic pattern decoding, plumb-line calibration

Traditional calibration target registration degrades significantly under severe motion blur, perspective slant, high lens distortion, and partial target occlusions. This paper introduces the Binary Hexagonal Galois Pattern (B-HGP) approach: a single-frame pattern registration method that maps one-dimensional Galois-field M-sequences onto a two-dimensional hexagonal lattice. The resulting pattern encodes absolute spatial coordinates locally, supports simultaneous algebraic error and erasure recovery, and tolerates these severe geometric and optical distortions. The research systematically investigates the pattern's statistical, geometric, and error-resilience properties. Furthermore, we derive the exact Galois Linear Feedback Shift Register (LFSR) initial state configuration that guarantees a zero phase shift between the forward and backward sequence tracking. Stress tests generated in Blender under perspective foreshortening demonstrate that the algebraic decoding pipeline keeps tracking error in sub-pixel range while maintaining absolute registration precision and a high recall rate of the pattern nodes. Finally, a decoupled plumb-line calibration approach that isolates lens distortion from intrinsic parameter estimation is presented as a practical application of the pattern. The open-source implementation and benchmarks are available at: https://github.com/gkis-conda/robust_calibration_pattern.

1 Introduction

The rapid development of autonomous vehicles [1, 2], augmented reality glasses [3], and omnidirectional imaging systems has made the fast and robust calibration of multi-sensor devices critical. These complex platforms heavily rely on visual sensor fusion for precise self-positioning. A fundamental requirement for successful fusion is the accurate estimation of sensor intrinsic and extrinsic parameters, which resolve strict temporal and spatial synchronization via known multi-sensor signals.

The calibration of multi-sensor setups, particularly Visual-Inertial Navigation Systems (VINS), faces severe challenges when resolving coupled spatiotemporal parameters. For instance, to establish spatial synchronization between multiple cameras via baseline estimation, a graph of scene views sharing the same spatial location must be constructed. The graph edges reflect the visual intersections of the camera fields of view. In practice, these views intersect only partially and asynchronously. Generally, time is a critical parameter in the estimation process; because there are no hardware guarantees regarding the exact moment a camera frame is captured, the sensor trajectories must be aligned in the temporal domain to achieve precise calibration.

Registering a dedicated calibration pattern offers a direct approach for identifying these spatiotemporal correspondences across multiple video streams. Even advanced hand-eye calibration methodologies require robust visual trajectory estimation, which introduces independent tracking complexities and cannot replace a pure optical pattern baseline.

The structural design of the registration pattern is directly tied to the success of constructing the view graph and establishing its edges. In unconstrained environments, pattern-based registration frequently suffers from perspective distortion, cluttered backgrounds, partial occlusions, and varying illumination conditions. Because camera fields of view often capture only cropped or rotated segments of the target, the pattern must remain entirely resilient against these geometric transformations. Generic targets easily fail under these conditions due to orientation and indexing ambiguities. For example, a conventional symmetric circular target from the OpenCV library [4] lacks robustness against 180° rotation ambiguities due to its structural symmetry [5]. While an asymmetric circle grid with an alternating row-shift layout seems to resolve this ambiguity under full visibility, it immediately fails as soon as the camera field of view captures only a cropped segment of the lattice.

To address these spatial limitations and environmental distortions, this paper introduces a registration approach based on algebraic coding theory. The proposed pattern inherently handles pattern crops by embedding error-correcting codes directly onto a regular hexagonal lattice. The pattern encodes three independent binary sequences into the lattice grid. The visible binary state at each lattice node is evaluated as a deterministic combination of the channel bits taken from each sequence at that specific lattice coordinate. Consequently, local crops of the printed target can be decoded to absolute lattice coordinates without supportive structural anchors, such as center tags or perimeter marks. This design enables single-frame localized decoding, significantly improving pattern tracking robustness and view graph generation in complex imaging configurations.

2 Motivation and the B-HGP Architecture

2.1 Related Approaches and Functional Limitations

Circle grids, tag-based fiducials (e.g., AprilTag [6], ArUco [7], WhyCode [8]), and chessboard or ChArUco [9] patterns remain industry standards for camera calibration due to their straightforward feature detection and low computational decoding overhead. However, these classical methodologies exhibit severe operational constraints in real-world environments. Furthermore, maintaining the strictly controlled operational conditions they require—such as uniform illumination, glare control, and shadow-free setups—demands substantial labor and engineering effort.

Localization Bias Induced by Pattern Design

Camera calibration requires establishing an exact mapping between physical coordinates on a rigid target canvas and discrete pixel coordinates on an image sensor. However, printed geometric features can introduce systematic localization biases purely through their interaction with perspective projection, optical defocus, and finite sensor resolution. Because modern calibration patterns fundamentally rely on two primary primitives—circles and line-intersection saddle points—it is critical to examine their individual geometric resilience under deformation.

Circle grids are highly valued in metrology because fitting a continuous ellipse to a pixel cluster yields sub-pixel precision through statistical averaging. However, a fundamental geometric bias emerges under oblique viewpoints: a physical circular feature projects onto the sensor as an ellipse whose 2D geometric center does not coincide with the projected image of the true 3D circle center [10]. This spatial shift mislocates the true calibration anchor, directly corrupting estimated focal lengths and radial distortion parameters. While applying concentric rings can provide sufficient projective constraints to analytically recover the true projected center via intersecting common polar lines [11], this approach is computationally demanding and highly fragile under image degradation. Under severe blur or low spatial resolution, subtle pixel noise drastically corrupts the estimated conic coefficients, breaking numerical stability and reintroducing localization drift.

To eliminate perspective center-shift bias, state-of-the-art frameworks heavily favor saddle points, such as checkerboard corner intersections. Under any planar homography, straight lines project strictly as straight lines; consequently, the mathematical intersection of two lines remains perfectly invariant under perspective transformation. Nevertheless, line-intersection features are exceptionally sensitive to optical defocus and illumination gradients [12, 13]. When a corner is blurred, high-intensity quadrants bleed into low-intensity quadrants. Consequently, gradient-tracking and

intensity-thresholding algorithms systematically pull the calculated corner toward the darker or lighter region, introducing an environment-dependent localization bias.

Conversely, circular primitives exhibit a unique advantage under extreme wide-angle deformations. For sensors adhering to wide-angle or fisheye projection models, local mappings closely approximate conformal transformations, natively preserving circular boundary geometries across the viewport [14, 15]. Even under intense radial distortion where analytical boundaries undergo non-linear stretching, local contour parameters degrade gracefully into regular ellipses. This geometric resilience allows precise sub-pixel centroid extraction under severe angular skews and peripheral lens distortion, where traditional saddle-point profiles suffer severe structural degradation.

Because both primitive geometries exhibit distinct, mutually exclusive resilience profiles against perspective center-shifts, illumination bleeding, and wide-angle conformal warping, achieving optimal tracking robustness requires strategically incorporating both geometric primitives into the target layout.

Pattern Registration Operational Limits

The fundamental trade-offs between circular and saddle-point primitives analyzed above directly account for why existing pattern registration pipelines fail in non-ideal conditions:

- **Spatial Scale Constraints:** When the distance between camera and target increases, small bit-cells of matrix tags fall below the sensor's resolution limit. These payload tags fail to decode long before larger geometric anchors lose visibility, severely restricting the operational range of the calibration setup.
- **Occlusions and Spatial Payload Overhead:** Fiducial markers require intact bit arrays for identification. Allocating surface area to payload data rather than dense metrology anchors leaves them highly vulnerable: even at full resolution, partial occlusions or frame cropping corrupt the tag array, rendering surrounding features unindexed.
- **Blur Sensitivity:** Standard checkerboard corner extraction degrades rapidly under motion or out-of-focus blur. When spatial image gradients suffer severe degradation [12, 13], localized operators fail to isolate clean saddle-point signatures, restricting calibration to static, highly controlled multi-frame routines.
- **Localization Bias:** Systems relying exclusively on circular primitives suffer from perspective localization bias under tilted views. Without complex corrective pipelines or implicit conic estimation [10, 11], this geometric shift systematically corrupts the estimated intrinsic focal length and lens distortion parameters.
- **Peripheral Fisheye Degradation:** Ultra-wide and fisheye lenses introduce catastrophic radial barrel distortion that warps straight checkerboard lines into high-curvature arcs near frame boundaries [16, 17]. This non-linear deformation destroys localized saddle-point signatures, causing feature extraction dropouts precisely near the image periphery.

2.2 The Binary Hexagonal Galois Pattern (B-HGP) Framework

To systematically address the spatial localization and payload constraints of traditional targets outlined above, we introduce a planar, pure key-point based pattern by embedding algebraic sequence identification directly into the local spatial lattice. The decoding pipeline executes GF(2) algebraic decoding to recover from arbitrary detection errors and erasures via unbiased, sub-pixel anchors built on compact geometric features.

The prefix **B** designates the current implementation as a binary code, chosen for its optimal contrast and mathematical simplicity on standard digital sensors. The letter **H** defines the Hexagonal topology of the framework, dictating that the control points are positioned strictly at the nodes of a regular hexagonal lattice to maximize the spatial packing density of anchors. Crucially, the configuration of the underlying finite field $\mathrm{GF}(2^5)$ exhibits an exact algebraic bijection with the geometric properties of this lattice. While any primitive polynomial naturally possesses a reciprocal twin, the extension polynomial $x^{31} - 1$ over GF(2) yields a total of exactly six primitive polynomials of degree five, which partition into three mutually reciprocal pairs. This complete algebraic ensemble maps perfectly onto the three primary spatial axes (U, V, W) of the physical hexagonal lattice, where each pair uniquely governs the forward and reverse sequence tracking paths along a single coordinate axis without structural redundancy. This underlying algebraic framework is inherently scalable; the two-dimensional spatial superposition principles generalize directly to non-binary fields — such as

q-ary maximum-length sequences over higher Galois fields GF(p^m) — by expanding the geometric primitives to support multi-shape, temporal, or color-encoded modulations.

In strict accordance with the algebraic foundations governing this architecture, this pattern class is formally designated in tribute to Évariste Galois solely as an expression of profound scientific respect for his pioneering formulation of finite fields and groups. By physically mapping abstract finite field operations onto the geometry of a two-dimensional optical target, the B-HGP framework demonstrates a practical application of algebra to robotic perception.

Furthermore, while framed here in the context of passive planar calibration targets, the B-HGP framework is uniquely suited for active structured light projection systems. In active 3D scanning, projected patterns suffer from severe non-linear perspective warping and localized erasures caused by object occlusions, ambient reflections, and complex surface geometries. Because the B-HGP decoder relies on phase-decoupled syndrome minimization, it can robustly decode projected spatial coordinates even across discontinuous surfaces and sharp depth steps.

3 Mathematical background

3.1 Hexagonal Lattice and Axial Coordinates

Every graphical spot of the pattern encodes a binary value and forms a two-dimensional hexagonal sampling lattice that serves as a blueprint binary template. In recent years, interest in non-Cartesian coordinate systems has grown. In particular, barycentric-style coordinates and the related integer cube or axial coordinates for hexagonal lattices are useful because their structures are preserved better under affine spatial transformations, and because the triangular cell is the minimal rigid two-dimensional element.

Triangular and hexagonal grids outperform square grids for several reasons:

- **Local rigidity and uniqueness of triangulation.** Triangles are rigid primitives with no internal degrees of freedom. Consequently, there is a unique triangulation in general position that is much more robust under homography lattice distortion and has robust graph under Delaunay triangulation. By contrast, a regular square grid contains many co-circular nodes (four vertices on the same circumcircle), so the Delaunay triangulation becomes unstable and ambiguous, as each quad admits two equally valid diagonals. This degeneracy makes a square layout sensitive to arbitrarily small perturbations: vanishingly small noise can flip diagonals and change the local topology, which complicates square-grid restoration.
- **Node degree, connectivity and redundancy.** A hexagonal lattice has vertex degree 6 (each node has six neighbors), whereas a square lattice has degree 4. More neighbors provide more independent local constraints. When nodes are missing due to occlusion or detection failure, nodes in a hexagonal lattice still have more neighboring constraints to support localization and reconstruction. Indeed, a square lattice node has additionally 4 diagonal elements at distance only $\sqrt{2}$ making high ambiguity under distortion.
- **Isotropy and directional balance.** The hexagonal lattice is more isotropic, showing sixfold symmetry versus fourfold for a square grid. For calibration and matching tasks, isotropy reduces directional bias: pattern features respond similarly for arbitrary camera orientations, which improves the robustness of detectors. Under camera roll rotation, a unique axis aligned with the image horizontal direction can always be selected.
- **Packing, quantization, and sampling advantages.** Circles packing in two dimensions is optimal for the hexagonal lattice and minimizes quantization error for uniform sampling. The packing densities are: hexagonal packing $\pi/(2\sqrt{3}) \approx 0.9069$ versus square packing $\pi/4 \approx 0.7854$. For equal node density, a hexagonal lattice yields a smaller covering radius and tighter spatial coverage, which benefits feature localization and interpolation.
- **Lattice Initialization and Growth:** Establishing a unique, unambiguous initial tracking seed requires collecting 4 true points for a square grid compared to only 3 points for a triangular layout. This increases the probability of successful seed initialization in highly degraded frames, which is particularly important for the performance of randomized algorithms like RANSAC. Furthermore, expanding a square lattice requires 2 points per step, whereas a triangular lattice requires only 1, making spatial growth substantially simpler.

Hexagonal Pattern Layout

The convergence of these two independent factors—the structural and spatial advantages of the hexagonal grid detailed above, alongside the intrinsic 6-fold symmetry of GF(2^5)—uniquely dictates the choice of this pattern geometry.

We represent lattice nodes using integer axial (or cube) coordinates (u, v, w) subject to the invariant

$$u + v + w = 0, \qquad w = -u - v \tag{1}$$

which are the standard integer coordinates for a hex lattice (not to be confused with normalized triangle barycentrics $\alpha + \beta + \gamma = 1$).

Coordinate Transformation Between Axial and Matrix Layouts.

To systematically organize the coordinate tracks, the system utilizes a direct, reversible bijection between the axial coordinates $(u, v)^\top$ of the hexagonal lattice and the row-column indices $(r, c)^\top$ of a standard rectangular matrix storage buffer. Because physical matrix indices are restricted to finite, non-negative integers, the axial indices are bounded to remain structurally compliant with the addressing constraints of the underlying layout. The forward and inverse transformations operate simultaneously over these coordinate spaces via the unified system:

$$\begin{pmatrix} r \\ c \end{pmatrix} = \begin{pmatrix} v \\ u + \lfloor \frac{v}{2} \rfloor \end{pmatrix}, \qquad \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} c - \lfloor \frac{r}{2} \rfloor \\ r \end{pmatrix} \tag{2}$$

The exact, reversible bijection in Eq. (2) guarantees that the topological grid connectivity remains fully preserved and unskewed across the entire plane.

Linear Conversion to Euclidean Blueprint Coordinates.

The mapping from the continuous two-axis axial coordinate space $(u, v)^\top$ to the orthogonal Euclidean blueprint plane $(x, y)^\top$ is strictly linear, operating via the forward affine transformation matrix:

$$\begin{pmatrix} x \\ y \end{pmatrix} = s \cdot \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \tag{3}$$

where $s \in \mathbb{R}^+$ represents the physical lattice step size. In the physical canvas domain, s defines the absolute Euclidean distance between any two adjacent, first-order nodes of the isometric hexagonal grid.

Substituting the discrete index mappings from Eq. (2) directly into the affine projection matrix of Eq. (3) reveals the explicit geometric transition from matrix indices to the continuous spatial grid:

$$\begin{pmatrix} x \\ y \end{pmatrix} = s \cdot \begin{pmatrix} 1 & \frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} \end{pmatrix} \begin{pmatrix} c - \lfloor \frac{r}{2} \rfloor \\ r \end{pmatrix} = s \cdot \begin{pmatrix} c + \frac{r}{2} - \lfloor \frac{r}{2} \rfloor \\ \frac{\sqrt{3}}{2} r \end{pmatrix} = s \cdot \begin{pmatrix} c + \frac{1}{2}(r \bmod 2) \\ \frac{\sqrt{3}}{2} r \end{pmatrix} \tag{4}$$

Eq. (4) mathematically demonstrates how the row-column profile can be used for direct computation of physical pattern coordinates.

Lattice Centroid Position.

To eliminate systematic geometric biases during lens distortion estimation, the physical origin $(0, 0)^\top$ of the world coordinate system must be positioned precisely at the geometric center of the printed target. For an engineered pattern containing a storage matrix layout of total height R rows and width C columns, calculating the grid center as a simple midpoint of discrete indices introduces a fractional phase shift due to the alternating row-staggering offset. To guarantee absolute spatial symmetry and zero-bias calibration tracking, we implement continuous matrix-space index centering constants, c_0 and r_0, defined directly by the array boundaries:

$$c_0 = \frac{1 - 2C}{4}, \qquad r_0 = \frac{1 - R}{2} \tag{5}$$

By appending these continuous matrix-space adjustments directly into the grouped index terms of Eq. (4), the final forward mapping to the centered absolute Euclidean blueprint coordinates is

expressed as:

$$\begin{pmatrix} x \\ y \end{pmatrix} = s \cdot \begin{pmatrix} c + c_0 + \frac{1}{2}(r \bmod 2) \\ \frac{\sqrt{3}}{2}(r + r_0) \end{pmatrix} \tag{6}$$

This explicit coordinate normalization completely compensates for the non-linear alternating row shift invariants.

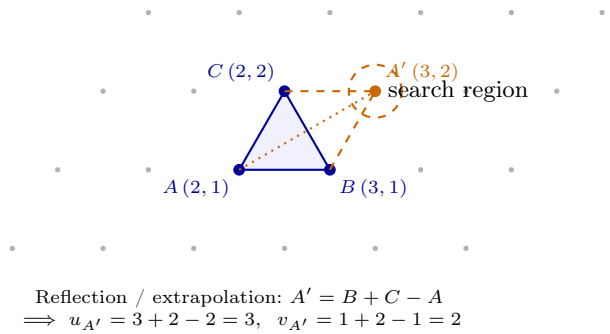


Fig. 1: Axial and Cartesian coordinate extrapolation on the hex lattice. Given a verified baseline triplex (A, B, C), the reflection/extrapolation across the edge $\vec{BC}$ is computed as $A' = B + C - A$. The dashed circle around A' defines the local spatial image window queried via the KD-Tree to bind the prediction to pixel sensor data.

Affine Properties of the Hexagonal Lattice.

Under perspective transformations, the structural topology of a hexagonal lattice remains invariant: while the continuous Cartesian coordinates $\mathbf{p} = (x, y)^\top$ undergo severe geometric distortion, the discrete axial coordinates $\mathbf{u} = (u, v, w)^\top$ are preserved within the projective mapping. The axial tracking indices serve as the localized axial weights of the reference frame. By substituting the dependent third-axis variable $w = -u - v$, the continuous Cartesian coordinate $\mathbf{p}$ of any node under lattice distortion on the image sensor simplifies to the following affine baseline:

$$\mathbf{p} \approx u\mathbf{p}_A + v\mathbf{p}_B + w\mathbf{p}_C = u(\mathbf{p}_A - \mathbf{p}_C) + v(\mathbf{p}_B - \mathbf{p}_C) \tag{7}$$

Accounting for this local affine approximation within the grid, the position of any adjacent node is governed by the equality of parallel lattice vectors. For any elementary triangle, the adjacent node coordinate $\mathbf{u}_{A'}$ forming the local rhombus $ABA'C$ is resolved via the vector identity $\vec{AC} = \vec{BA'} \implies \mathbf{u}_C - \mathbf{u}_A = \mathbf{u}_{A'} - \mathbf{u}_B$. This arithmetic structure maps across both the topological matrix and the continuous physical image plane:

$$\mathbf{u}_{A'} = \mathbf{u}_B + \mathbf{u}_C - \mathbf{u}_A, \ \mathbf{p}_{A'} \approx \mathbf{p}_B + \mathbf{p}_C - \mathbf{p}_A \tag{8}$$

The expression holds as an invariant over the ideal axial coordinate space to yield a new topological node $(u', v', w')^\top$, while providing a localized affine approximation on the physical sensor plane to guide the nearest-neighbor search.

Moreover, knowing only an edge of a primary triangle, we can construct the expansion of a hexagonal sub-grid. Within a globally counter-clockwise (CCW) oriented lattice topology, where the positive signed area condition for the edge vectors satisfies $\det(\vec{AB}, \vec{AC}) > 0$, the directed edge axial vector $\vec{BC} = (\Delta u, \Delta v)^\top$ encapsulates the local orientation of the adjacent lattice cells. Let $\mathcal{G}_6 = \{\mathbf{e}_i\}_{i=0}^{5}$ denote the cyclic group of six unit direction vectors defining the primary axes of the hexagonal lattice. Crucially, because every lattice edge is inherently an element of this directional cyclic group, the edge vector maps directly to a discrete group index k such that $\vec{BC} = \mathbf{e}_k$. As a direct consequence of this group property, the relative spatial displacement targeting the new exterior vertex is resolved by the prediction formula:

$$\mathbf{u}_{A'} = \mathbf{u}_B + \mathbf{r}, \quad \text{where } \mathbf{r} = \mathbf{e}_{(k-1) \pmod 6} \tag{9}$$

This boundary formulation remains valid within the topological index domain only. To safely account for local skew distortion, the continuous Cartesian prediction in Eq. (8) should be used for tracking within the pixel domain.

Nearest-Neighbor Search and Ambiguity Reduction.

To isolate the nearest candidate point to the obtained spatial prediction, the framework exploits the geometric capabilities of the hexagonal lattice, which minimizes neighbor-search ambiguity under distortion. When subjected to Lowe's Nearest-Neighbor Distance Ratio (NNDR) filter ($d_1/d_2 < 0.8$) [18], which identifies regions of nearest-neighbor tracking uncertainty based on the distances to the first two closest candidates d_1 and d_2, the choice of grid geometry becomes critical. Lowe demonstrates that the NNDR test at a maximum threshold of 0.8 eliminates 90% of false matches while discarding less than 5% of correct assignments [18]. Indeed, this ratio test triggers and rejects a candidate point whenever the prediction coordinates fall inside a cell's geometric uncertainty zone, where the distance buffer between competing neighbor anchors shrinks below the confidence margin. Stochastic geometry and digital quantization research [19, 20] state that a square lattice exhibits a severe uncertainty zone covering approximately 24% of the entire cell area. In contrast, the symmetry of the hexagonal Voronoi geometry compresses this uncertainty zone to merely 7% of the cell area. By reducing this critical failure space by a factor of 3.4×, the hexagonal framework inherently minimizes the probability of triggering Lowe's rejection threshold. This geometric reduction prevents valid tracking nodes from being falsely discarded under severe perspective slants, anchoring the tracking prediction directly within unambiguous, highly stable geometric attraction basins.

Grid Expansion Efficiency.

The framework leverages local affine properties to predict the coordinates of adjacent nodes linearly. This vector extrapolation highlights a major topological advantage of the hexagonal configuration over traditional square grids under severe perspective slant and lens distortion. When expanding the mesh to construct an adjacent cell, a triangular grid requires only a single new node point to close the cell geometry, whereas a square grid requires tracking two distinct points. Consequently, assuming an individual node detection probability of p, the successful completion probability of a new cell scales quadratically as p^2 for a square layout, compared to the linear scaling p of the triangular grid. Furthermore, while growing a square grid requires executing full homography estimations across 4-point configurations to predict adjacent points [21], a triangular grid evaluates the new coordinate via a simple vector reflection across the shared boundary edge, requiring only three verified anchors and significantly reducing processing overhead.

Topological Transformations and Rotational Permutations.

Due to the rotational isotropy of the hexagonal lattice, its topological properties remain identical regardless of spatial orientation. This fundamental invariance allows the reconstruction algorithm to operate along an arbitrary coordinate axis chosen solely for processing convenience, such as alignment with the horizontal image frame, without needing prior knowledge of the target's global orientation. During grid expansion, the extracted axial coordinates align with the camera pixel grid up to a discrete 60° rotation step. Hexagonal rotations by multiples of 60° reduce to exact integer permutations of the three-axis vector $(u, v, w)^\top$. A single counter-clockwise (CCW) rotational step maps these parameters according to the sign-inverting permutation relation:

$$\begin{pmatrix} u' \\ v' \\ w' \end{pmatrix} = \begin{pmatrix} -w \\ -u \\ -v \end{pmatrix}. \tag{10}$$

Repeated sequential applications of this operator cycle through the six symmetric orientations of the hexagonal lattice without altering the graph topology.

The practical need for these exact rotations emerges during the global template alignment and decoding phases. The reconstruction pipeline initially extracts non-oriented, disconnected hexagonal islands across the image plane. Under arbitrary camera roll distortion, the local coordinate frames of these isolated sub-graphs are rolled relative to the global template. However, due to the sixfold symmetry of the lattice, the decoding stage can either rotate the extracted topological matrix or evaluate the permutation cycles of Eq. (10) to identify the optimal orientation with the minimal

error footprint. Once the true orientation axis is recovered, any isolated island can be seamlessly mapped to align with the global blueprint layout. Because the underlying connectivity network of the reconstructed island remains unbroken, the system achieves instant axis re-alignment and coordinate re-assignment with zero numerical drift or lattice deformation.

3.2 Physical Pattern Key-Point Geometry

To encode absolute coordinate data according to Eq. (3) onto a physical tracking target, the system utilizes a two-component hexagonal mosaic consisting of geometrically polar primitives: ideal circles and equilateral triangles. In this encoding layout, isolated triangles represent binary ones, while circles represent binary zeros. These specific shapes are chosen because their geometric invariants are maximally separated; a triangle represents the minimal-sided closed polygon, whereas a circle can be viewed as the limiting case of a maximal-sided polygon, structurally minimizing the risk of inter-class aliasing under heavy optical distortion.

The vertices of the extracted triangles can serve as calibration anchors. By utilizing boundary-minimization algorithms along their straight edges, triangle vertices provide sub-pixel projection accuracy, yielding a spatially homogeneous random coverage with anchors. Furthermore, each resolved triangle independently generates three additional distinct point matches instead of the single center coordinate provided by conventional dot or checkerboard grids, effectively multiplying the spatial constraints available for downstream optimization frameworks. Alternatively, detected shape centroids can be used as a fast approximation for the lattice reconstruction.

Polygon Extraction and Prefiltration.

The input image captured by the camera sensor first undergoes localized adaptive Gaussian thresholding to compensate for non-uniform illumination, lens vignetting, or spatial shadows. The resulting binary matrix is processed to extract a set of closed boundary loops, where each detected contour is mathematically defined as a discrete planar polygon.

To eliminate extreme high-frequency sensor noise and macro-environmental background artifacts before executing algebraic analysis, the polygon set is passed through a broad area gate constraint. Any extracted polygon whose pixel area fails the boundary condition $A_{\min} < \text{Area} < A_{\max}$ is immediately discarded, ensuring that the downstream pipeline only evaluates structurally relevant candidates.

A major threat to classification integrity arises from clipping constraints. When a tracking marker partially intersects the edge of the camera sensor, its physical geometry is truncated by the image frame. A clipped shape experiences a sudden drop in its area-to-perimeter ratio, causing misclassifications while introducing a severe projective center bias. To eliminate this projection artifact, the pipeline enforces a strict spatial viewport barrier. Any candidate contour whose extreme coordinates touch the boundary pixels of the image matrix within a minimal safety margin is instantly eliminated.

Stage 1: Convex Polygon Filtration.

Under any perspective distortion, polygon convexity is strictly preserved. Following the initial area gate, the surviving polygons must be cleared of complex binarization noise that distorts polygon boundaries and their topological properties. To prevent this degradation, the framework incorporates a topological pre-filtering gate by evaluating the structural solidity γ, defined as the ratio of the polygon's mass μ_L (or Lebesgue area) to the mass of its convex hull (μ_H):

$$\gamma = \frac{\mu_L}{\mu_H} \tag{11}$$

This ratio is positive and strictly bounded by 1 for convex shapes, degrading non-linearly for non-convex or star-like structures. We can mathematically justify the operational threshold by analyzing the geometric topology of two distinct object classes:

- **Calibration Primitives:** In the continuous domain, both ideal circles and triangles are strictly convex, compact sets yielding a theoretical solidity of $\gamma = 1$. When rasterized onto a discrete pixel matrix, spatial aliasing and light optical blur introduce minor edge irregularities, causing the actual solidity to experience a slight decay. As established in digital geometry frameworks [22], the

pixelation effect causes the empirical convex hull μ_H to overestimate the true bounded Lebesgue area μ_L, dropping the continuous ideal solidity to ≈ 0.9 even for uncorrupted shapes.

- **Fractal Binarization Noise:** Conversely, random high-frequency thresholding noise, sensor artifacts, and shadow gradients manifest as porous or highly non-convex pixel clusters behaving as discrete fractal aggregates. Geometric analysis proves that due to their branching structures, their calculated convex hull μ_H consistently overestimates their true mass μ_L, trapping their empirical solidity distribution below a statistical ceiling. In accordance with stochastic geometry models and discrete contour quantization frameworks [20, 22], this spatial scattering bound establishes the formal random structure class centroid at $\gamma_{\text{noise}} \approx 0.55$.

By establishing the lower-bound convexity constraint at $\gamma_{\min} \approx 0.8$, the system permits up to a 20% planar area deformation on the boundary edges of true markers due to camera blur, while instantly annihilating high-Hausdorff-dimension fractal noise assets. This process isolates true, topologically closed geometric polygons before executing final shape classification.

Stage 2: Shape Classification.

Once the input polygons are cleared of 2D-porous noise, the remaining valid, solid polygons undergo shape identification. Standard detectors relying on discrete vertex counting or local gradient thresholding fail under motion blur. To achieve robust scale and noise invariance, the proposed pipeline utilizes the circularity metric:

$$\text{Circularity} = \frac{4\pi \cdot \text{Area}}{\text{Perimeter}^2} \tag{12}$$

For an ideal circle, this value evaluates strictly to 1.0, whereas for an ideal equilateral triangle, it reduces analytically to exactly $\frac{\pi}{3\sqrt{3}} \approx 0.605$. This deterministic separation allows the shape-identity problem to be formulated via a Maximal-Margin Linear Discriminant Classifier layout. A candidate node is classified strictly based on its proximity to these centroids: if the absolute distance to the triangular reference is smaller than the distance to the circular reference, it is assigned to the respective class. Conversely, if the circularity value falls beyond a maximum margin distance from both centroids, the node is flagged as unclassified. Perspective slant and optical blur distort the contours and shift the empirical circularity values inward, but because the circularity margin is wide, the attraction basins of both classes remain linearly separable even under catastrophic imaging conditions.

3.3 Maximum-length Sequences: Key Properties

M-sequences (maximal-length LFSR sequences) are used here as axis codes because of their sharp autocorrelation, low mutual correlation, and compact algebraic description. These properties make M-sequences highly suitable for spatial encoding and for robust identification from short local fragments. By analogy, Gold codes built from M-sequences are used in CDMA and GPS systems to separate and synchronize multiple transmitters sharing the same frequency channel.

Definition and Primitive Polynomials.

A maximal-length linear-feedback shift register (LFSR) of degree n over GF(2) generates an M-sequence b_t with a period of $P = 2^n - 1$, where the sequence time index $t \in \mathbb{Z}_P$. We use the following polynomial notation throughout:

$$f(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0, \tag{13}$$

with coefficients $a_j \in \text{GF}(2)$ and $a_0 = 1$ for primitive polynomials. Classically, the structural definition of the generated sequence is governed by the recurrence relation in Fibonacci form:

$$b_{k+n} = \bigoplus_{j=0}^{n-1} a_j \, b_{k+j}. \tag{14}$$

The number of distinct primitive polynomials of degree n equals $\varphi(2^n-1)/n$, where φ denotes Euler's totient function. Crucially, $n = 5$ is the minimal degree that yields exactly six distinct primitive polynomials ($\varphi(31)/5 = 6$), which naturally form three reciprocal (time-reversed) pairs — the precise number required to map independently onto the three spatial axes (U, V, W) of the hexagonal layout

without structural redundancy. In the proposed framework, this degree $n = 5$ setting achieves an optimal balance between spatial footprint and code period ($P = 31$). Due to the exclusion of the trivial all-zero register state, an M-sequence exhibits a nearly perfect statistical balance over its full orbit, containing exactly $\frac{P+1}{2}$ ones and $\frac{P-1}{2}$ zeros.

To minimize execution latency and ensure computational efficiency, the system utilizes a Galois LFSR configuration with a right-shift architecture where the output stream is extracted from the least significant bit (LSB). This choice is motivated by algorithmic simplicity, as the Galois state update can be executed via a single, parallel bitwise operation on a standard CPU register:

```
output ← state & 1
state ← state ≫ 1
if output = 1:
  state ← state ⊕ (poly ≫ 1)
```

Although the internal state trajectories differ between the Fibonacci definition and the active Galois implementation, the generated output bitstream remains strictly identical.

Time Reversal and Reciprocal Polynomials.

The time reversal of an M-sequence is fundamentally associated with its reciprocal connection polynomial: if $f(x)$ is primitive, its reciprocal polynomial generates the time-reversed sequence up to a specific cyclic shift. On the hexagonal lattice, this algebraic symmetry matches the six principal lattice directions (three axes with two opposite orientations each), requiring the decoder to handle both forward and reversed readouts. By selecting a proper initial LFSR state, it is possible to satisfy the sequence reversibility condition $b_{-k} = b'_k$, where b and b' represent the sequences generated by the corresponding forward and reciprocal polynomials.

For any odd prime degree n of a primitive shift register, the total number of non-trivial, time-reversal invariant phase anchor states N in the M-sequence is given by:

$$N = 2^{\frac{n-1}{2}} - 1. \tag{15}$$

For a Galois right-shift configuration of degree n, a state vector $\mathbf{s} = (s_{n-1}, \ldots, s_1, s_0)^T$ acts as a spatial inversion checkpoint if and only if it satisfies the following structural palindrome requirements:

1. The LSB must be zero: $s_0 = 0$.
2. The remaining bit sub-vector $(s_{n-1}, \ldots, s_1)$ of length $(n-1)$ must form a symmetric binary palindrome around its central element.

In this framework, we utilize the initial state 12 (binary `01100`) to obtain a spatially reversible axis phase across all six primitive polynomials of degree 5. For the complete derivation of these topological properties, see Appendix A.

Differential Sequences.

The core algebraic trait of any M-sequence is the shift-and-add property, which states that the bitwise XOR combination of a sequence with any cyclically shifted copy of itself yields the same sequence at a unique, non-linear phase offset τ:

$$b_{t+\tau} = b_t \oplus b_{t+\delta}, \tag{16}$$

where all indices are evaluated modulo P. Such transforms act as strict linear mappings over GF(2) onto the same sequence space.

Discrete Logarithm Extraction and Phase-to-Lattice Mapping.

Every n-bit window is unique across the full sequence period, guaranteeing a bijective mapping between the extracted sub-sequence phase and the coordinate space of the calibration target. Algebraically, this phase tracking operation reduces to a discrete logarithm problem over the finite field GF(2^n). Evaluating this discrete logarithm in a Galois layout introduces internal state scrambling, meaning that consecutive bits from the output stream do not align with raw register triggers.

To resolve this phase offset efficiently, the framework avoids heavy matrix inversions by formulating the discrete logarithm as a highly optimized substring lookup. At system initialization, the

canonical reference sequence $\mathbf{b}_{\text{ref}}$ of maximum length $P = 2^n - 1$ is cyclically extended by n bits to construct a reference array of length $P + n$ elements. During tracking, the decoder extracts the first n clean bits of the recovered sequence as a search pattern. The absolute cyclic phase shift $\phi \in \{0, \ldots, P-1\}$ is then resolved by executing a substring search leveraging the Boyer-Moore-Horspool algorithm. This choice minimizes computational latency with a linear time complexity of $O(P + 2n)$.

The resolved phase parameters map directly to the row and column positions $(r, c)^\top \in \mathbb{Z}^2$ of the physical hexagonal lattice island. Specifically, the global row coordinate r is identical to the vertical phase component ϕ_v, while the column coordinate c incorporates a hexagonal stride shift:

$$r = \phi_v,\ c = \left(\phi_u + \left\lfloor \frac{r}{2} \right\rfloor\right) \pmod{P}. \tag{17}$$

This relation follows directly from Eq. (2) after applying modulo P and accounting for the valid range of (r, c) indices.

3.4 Hexagonal Lattice Synthesis and M-Sequence Distillation

Two-Dimensional Spatial Hexagonal M-Sequence Superposition.

To encode absolute spatial coordinates onto a continuous two-dimensional plane, we map a hexagonal lattice utilizing a three-axis coordinate system (u, v, w). We select three independent, primitive maximum-length sequences (M-sequences), denoted as U, V, and W, and assign them to each respective spatial axis. The binary state $S(u, v)$ synthesized at any discrete lattice node (u, v) on the sensor canvas is generated via a linear GF(2) superposition incorporating the time-reversal duality of the reciprocal configurations:

$$S(u,v) = U_{u \bmod P} \oplus V_{v \bmod P} \oplus W_{(-u-v) \bmod P} = U_{u \bmod P} \oplus V_{v \bmod P} \oplus W'_{(u+v) \bmod P}. \tag{18}$$

Due to the strict linearity of the Galois field, a local linear operator exists such that its application to the synthesized spatial grid cleanly extracts the underlying one-dimensional axis streams.

Statistical Balance of the Superposition.

While an isolated one-dimensional M-sequence of degree $n = 5$ and period $P = 31$ exhibits a slight structural imbalance—containing $\frac{P+1}{2} = 16$ ones and $\frac{P-1}{2} = 15$ zeros—the linear spatial superposition across three axes shifts this probability distribution closer to an ideal uniform state.

Let $p = 16/31 \approx 0.5161$ be the individual probability of a binary one within a sequence timeline. The linear combination at any discrete hexagonal node $S(u, v)$ evaluates to a binary one if and only if an odd number of active ones are sampled across the three independent axis channels. Applying the properties of independent Bernoulli trials, this condition is modeled exactly via the binomial distribution profile:

$$\Pr(S(u,v) = 1) = \binom{3}{1} p(1-p)^2 + \binom{3}{3} p^3. \tag{19}$$

Substituting the exact rational probabilities into the algebraic expression yields:

$$\Pr(S(u,v) = 1) = 3\left(\frac{16}{31}\right)\left(\frac{15}{31}\right)^2 + \left(\frac{16}{31}\right)^3 = \frac{14896}{29791} \approx 0.50002. \tag{20}$$

This evaluation highlights a key structural advantage of the spatial design: the linear superposition of three correlation-orthogonal binary streams does not distort the baseline statistical distribution of the sequence. Instead, it refines the balance, making the distribution of ones and zeros uniform across the continuous physical lattice. Consequently, the physical canvas behaves as a balanced spatial pattern field. Because binary ones are represented by explicit triangles whose vertices provide three distinct sub-pixel projection points, this uniform spatial entropy guarantees bias-free calibration point coverage under arbitrary imaging conditions.

Sequence Distillation via Rhombus Kernels.

To reconstruct the absolute phase along an individual coordinate axis, the intertwined two-dimensional lattice pattern undergoes an algebraic inverse distillation process. This transformation

is executed using a sliding window spanning two adjacent lattice rows, incurring a minimal 1-bit boundary truncation. By evaluating compact 4-pixel XOR rhombus kernels, the contributions from the two unwanted coordinate axes are identically canceled over GF(2), yielding a clean, isolated 1D differential observable. The distillation operators evaluate a sliding column index c across the two-row patch for the horizontal and 60°-slanted directions, respectively.

U-Axis Distillation: The operator samples a slanted rhombus kernel across a staggered coordinate step, computing the local XOR sum:

$$dU_c = S(0, c+1) \oplus S(0, c+2) \oplus S(1, c) \oplus S(1, c+1). \tag{21}$$

This specific geometric footprint completely cancels the V and W' interference pathways, distilling a pure one-dimensional forward differential sequence:

$$dU_c = U_c \oplus U_{c+2}. \tag{22}$$

W-Axis Distillation: The operator samples a compact, tight 2×2 spatial diamond block:

$$dW_c = S(0, c) \oplus S(0, c+1) \oplus S(1, c) \oplus S(1, c+1). \tag{23}$$

This block geometry annihilates both the U and V coordinate interferences, leaving a clean, unmasked differential output bound to the third axis:

$$dW_c = W_c \oplus W_{c-2}. \tag{24}$$

Crucially, due to the commutativity of addition over GF(2), this expression can be rewritten identically as a forward-stepping relation. Furthermore, because the coordinate w decrements as the column index c increases according to Eq. (1) and Eq. (2), this equivalence maps the spatial stream directly to a positive-direction step along the axis of the dual (reciprocal) sequence:

$$dW'_c = W'_c \oplus W'_{c+2}. \tag{25}$$

Note that the dW stream is extracted from the direction immediately adjacent to dU in counterclockwise (CCW) order, while the primary dU stream can be matched to any arbitrary direction to compensate for camera rotation and the inherent hexagonal symmetry of the pattern. Using Eq. (16), we can directly decode the differential sequence while accounting for an additional phase shift τ. The extracted subsequences are matched against the polynomial ensemble (see Section 3.6). Once the polynomials are determined, the decoder evaluates the exact phases for both detected subsequences, fully localizing the sliding window within the global pattern lattice.

Induced Erasure and Error Structure.

A practical consequence of pixel distillation is the mapping of camera sensor dropouts and physical occlusions. When any node within the 4-pixel rhomboid footprint encounters an obscured value, the underlying bits cannot be distilled, resulting in an explicit erasure. On the other hand, every pattern node participates in two sequential rhombus distillation steps along both axes, which imposes the following structural properties:

- **Pairing Effect:** Any missed or incorrectly detected spatial node generates correlated, paired erasures or bit-flip errors in the distilled differential stream.
- **Syndrome Co-Location:** Due to the overlapping geometry of the sliding window, both distilled bit streams (dU and dW) produce identical syndrome configurations at co-located bit positions.

Consequently, the sequence decoder does not receive an ideal continuous stream; instead, it processes a short fragment containing bit-paired erasures alongside correlated bit-paired errors. This specific error topology validates the necessity of the syndrome minimization decoder, which natively solves the phase resolution problem under arbitrary erasure and paired-error distributions.

3.5 Statistical Analysis of Sequence Alignment

To establish a mathematically rigorous framework for spatial sequence identification under optical noise, let the reference axis M-sequence be mapped to bipolar symbols $s_t \in \{\pm 1\}$, and let $o_t \in \{\pm 1\}$

denote the observed spatial fragment of length L. We assume each observed bit is inverted with an error probability ε.

Let e represent the actual number of bit errors within the observed window L. The empirical error frequency e/L acts as an unbiased sample estimate of ε, with expected error count $\mathbb{E}[e] = L\varepsilon$. The normalized empirical cross-correlation threshold $\delta \in [-1, 1]$ is directly linked to the observed error count, matching the normalized sum of true bit matches:

$$\delta = \frac{1}{L}\sum_{k=1}^{L} o_k s_k = 1 - \frac{2e}{L}, \mathbb{E}[\delta] = 1 - \frac{2\mathbb{E}[e]}{L} = 1 - 2\varepsilon. \tag{26}$$

Conversely, correlating the observed fragment against an uncorrelated background or random spatial noise maps directly to a one-dimensional random walk with standard deviation $\sigma = \sqrt{L}$. Normalizing the expected signal value by the fluctuation amplitude yields the operational Signal-to-Noise Ratio (SNR) defining the detection pipeline:

$$\mathrm{SNR} = \frac{L(1-2\varepsilon)}{\sqrt{L}} = (1-2\varepsilon)\sqrt{L} = \delta\sqrt{L}. \tag{27}$$

Equation (27) indicates that even under severe optical channel degradation (high ε), a successful match is guaranteed by collecting a longer local sequence fragment L.

Stochastic Axis Mismatch Probability Bounds.

To estimate the exact analytical boundary for L necessary to guarantee robust alignment, the framework evaluates the structural properties of M-sequences under worst-case noise configurations. For a period $P = 2^n - 1$, M-sequences possess an ideal two-valued periodic autocorrelation function $R(k)$:

$$R(k) = \sum_{t=0}^{P-1} s_t s_{(t+k) \pmod P} = \begin{cases} P, & k \equiv 0 \pmod P, \\ -1, & k \not\equiv 0 \pmod P. \end{cases} \tag{28}$$

While global autocorrelation governs long periods, practical spatial tracking operates on local truncated windows ($L \ll P$). Evaluating the search space across competing incorrect hypotheses, the single-axis familywise error rate (FER) for a false correlation event under background noise is bounded by Hoeffding's inequality:

$$\Pr(\text{axis mismatch}) \le (N_{\text{cand}} - 1)\exp\left(-\frac{\delta^2 L}{2}\right). \tag{29}$$

Crucially, in this spatial calibration architecture, $N_{\text{cand}} = 6$ represents the total spatial coordinate directions on the hexagonal lattice. Because exactly one direction corresponds to the true alignment, the number of competing false hypotheses is strictly less by one.

Spatially Correlated Multi-Axis Verification.

In practical radio-frequency communications, differential sequence demodulation ($d_t = o_t \oplus o_{t-1}$) is employed to ensure phase polarity invariance and eliminate the need for absolute carrier phase recovery, allowing the decoder to reliably distill signals. In a standard one-dimensional channel, a single bit-flip corrupts two adjacent differential decision symbols, effectively doubling the apparent bit error count within the decoded observation window [23]. In our two-dimensional spatial pattern, however, a single corrupted or missing physical symbol affects a 2D neighborhood and corrupts up to 6 adjacent spatial neighbors, simultaneously propagating correlated error patterns across all three coordinate directions during signal distillation. Modeling these spatially coupled bit errors as statistically independent within the Hoeffding tail evaluation yields a strict upper bound. This independent model guarantees that the actual operational reliability under fully correlated spatial noise will be higher.

The spatial decoder never accepts an isolated coordinate axis independently once a candidate syndrome is resolved. Instead, it evaluates unified, paired geometric hypotheses (Section 3.4). Let primary axis U be the candidate selected via minimum syndrome error weight W_U. Due to the rigid

2D calibration target geometry, the orientation of secondary axis W is deterministically fixed by axis U. Furthermore, because orthogonal sequences U and W are extracted from the same local image frame, physical surface defects (glares, blur, or scratches) induce spatially correlated error distributions (see Section 3.4). Channel errors on secondary track W manifest at identical or predictably shifted bit indices relative to track U. The decoder utilizes the secondary axis as a structural error-mask filter ($\mathbf{e}_W[t] = \mathbf{e}_U[t + \delta]$, where δ is a distillation kernel shift). Since orthogonal sequences exhibit low mutual cross-correlation and candidate choices are bound by the primary axis, the probability that both channels simultaneously misidentify the candidate and yield matching spatial error configurations can be assumed to be upper-bounded by the product of independent tail probabilities:

$$\Pr(\text{joint axis mismatch}) \leq (N_{\text{cand}} - 1) \left[\exp\left(-\frac{\delta^2 L}{2}\right)\right]^2 = (N_{\text{cand}} - 1) \exp\left(-\delta^2 L\right). \tag{30}$$

A candidate spatial orientation is confirmed if and only if the error configurations independently resolve on both channels and satisfy the error mask alignment constraint. If the error patterns fail to validate simultaneously, the joint hypothesis is instantly rejected, forcing the decoder into a safe no-decision state.

Numerical Evaluation and Operational Boundaries.

Table 1 compiles the analytical risk limits under a single paired bit error ($e = 2$, corresponding to $\delta = 1 - 4/L$) across tracking window lengths $L \in [8, 17]$. The probabilities are evaluated using the single-axis bound (Eq. (29)) and the joint orthogonal bound (Eq. (30)).

Table 1: False-Alarm Probability Bounds and Code Ambiguity Under a Paired Error ($e = 2$)

Length L (bits)	Threshold δ	SNR	$\Pr_{\text{single}}$	$\Pr_{\text{joint}}$	Ambiguous #	$\Pr_{\text{ambiguity}}$
8	0.500	1.414	1.00000	0.67668	36	0.4194
9	0.556	1.667	1.00000	0.40454	18	0.1935
10	0.600	1.897	0.82593	0.13643	3	0.0323
11	0.636	2.111	0.54652	0.03583	2	0.0215
12	0.667	2.309	0.34738	0.00804	1	0.0108
13	0.692	2.496	0.21523	0.00163	0	0.0000
14	**0.714**	**2.673**	**0.13068**	**0.00030**	**0**	**0.0000**
15	0.733	2.840	0.07820	0.00005	0	0.0000
16	0.750	3.000	0.04623	0.00001	0	0.0000
17	0.765	3.153	0.02703	$< 10^{-6}$	0	0.0000

The numerical breakdown defines two fundamental operational thresholds:

1. **Algebraic Uniqueness Threshold ($L \geq 13$):** For $L < 13$, truncated subsequences intersect across primitive polynomials at degree $n = 5$, yielding non-zero code ambiguities (Ambiguous # > 0). At $L = 8$, the algebraic ambiguity probability reaches 41.94%. At $L = 13$, codebook intersections collapse strictly to zero (Ambiguous # $= 0$), guaranteeing absolute algebraic uniqueness of uncorrupted sequence fragments.
2. **Statistical Reliability Threshold ($L \geq 14$):** Although $L = 13$ ensures algebraic uniqueness, an isolated single axis under $e = 2$ remains vulnerable to noise ($\Pr_{\text{single}} = 0.215$). The target reliability threshold of less than one false alignment per thousand calibrations ($\alpha = 10^{-3}$) is achieved on joint orthogonal channels at $L = 14$ bits ($\Pr_{\text{joint}} = 0.00030$).

Practical Linear Scaling and Early-Exit Decision Criteria.

To minimize computational latency, the streaming decoder accepts variable-length observation buffers L containing an arbitrary error count e, executing an Early-Exit decision rule as soon as Hoeffding's statistical bound condition is satisfied.

Equating the joint false-matching bound (Eq. (30)) to the required significance level α yields:

$$(N_{\text{cand}} - 1) \exp\left(-\delta^2 L\right) \leq \alpha \implies \delta^2 L \geq \ln\left(\frac{N_{\text{cand}} - 1}{\alpha}\right). \tag{31}$$

Substituting $\delta = 1 - \frac{2e}{L}$ and expanding quadratic terms gives $L - 4e + \frac{4e^2}{L} \geq \ln\left(\frac{N_{\text{cand}}-1}{\alpha}\right)$. Neglecting the small higher-order residual $4e^2/L$ for $L \gg e$ defines a practical, closed-form Linear Buffer Length Formula:

$$L_{\min}(e, \alpha) \approx 2 \cdot e + C(\alpha), \qquad \text{where } C(\alpha) = \left\lceil \ln\left(\frac{N_{\text{cand}} - 1}{\alpha}\right) \right\rceil. \tag{32}$$

For $N_{\text{cand}} = 6$ and $\alpha = 10^{-3}$, the baseline offset evaluates to $C(10^{-3}) = 9$ bits. Thus, the minimum required window length scales linearly with detected errors as $L_{\min}(e) = 2e + 9$.

Table 2: Minimum Buffer Length $L_{\min}$ for Early-Exit Decision Under Joint Verification

Error Count (e)	$\alpha = 1.0\%$ ($C = 7$)	$\alpha = 0.5\%$ ($C = 7$)	$\alpha = 0.1\%$ ($C = 9$)
0	7	7	9
1	9	9	11
2	11	11	13
3	13	13	15
4	15	15	17
5	17	17	19

Table 2 summarizes the resulting statistical thresholds across different error counts and significance levels.

In practice, a minimal operational window length of $L = 11$ is used, as the probability of algebraic ambiguity remains low and the comparative syndrome margin test reliably identifies the correct solution even below the $L_0 = 13$ combinatorial limit. The statistical requirement for $L \geq 11$ under $e = 2$ bit errors ($\alpha \leq 1.0\%$) directly corresponds to the fundamental physical capacity of the greedy syndrome decoder as explicitly derived in Section 3.6.

Final Decision Criteria.

The complete phase assignment and early termination pipeline is governed by two strict verification stages:

1. **Early-Exit Uniqueness Verification:** A candidate hypothesis with e errors extracted from buffer length L is declared statistically definitive and yields an immediate Early-Exit without further tracking expansion *if and only if* both conditions hold simultaneously:

$$L \geq \max(L_0, L_{\min}(e, \alpha)). \tag{33}$$

2. **Syndrome Distance Test:** Under noisy channels where candidates do not satisfy early exit, the decoder evaluates the syndrome weight gap (i.e., the difference in estimated error count) between the leading polynomial hypothesis (e_1) and the runner-up hypothesis (e_2):

$$\Delta e = e_2 - e_1 > \tau_{\text{gap}}, \tag{34}$$

where $\tau_{\text{gap}} > 0$ defines the required significance margin. If $\Delta e \leq \tau_{\text{gap}}$, both polynomial hypotheses are rejected due to classification ambiguity, and the entire decoded sequence is discarded to prevent coordinate aliasing.

3.6 The Syndrome Minimization Pipeline

Within the B-HGP framework, the first task is a polynomial identification via its syndrome minimization. In algebraic coding theory, syndrome decoding is typically confined to continuous blocks contaminated by independent white noise. Within our framework, however, the 2D spatial layout is subjected to erasures and local rhomboid distillation operators impose differential degradation (see Section 3.4).

The physical input to the decoding pipeline consists of a received binary fragment vector $\mathbf{b} \in \text{GF}(2)^L$ extracted over an observation window of length L, accompanied by a co-located binary

observation mask vector $\mathbf{m} \in \mathrm{GF}(2)^L$. The components of this mask indicate the completeness of the spatial stream: $\mathcal{U} = \{j \in \mathbb{Z}_L \mid m_j = 0\}$ flags unknown erasure gaps (such as sensor blowouts or lens occlusions) containing uninitialized random noise with an *a priori* error probability of $P(\text{error}) = 0.5$, while $\mathcal{K} = \{j \in \mathbb{Z}_L \mid m_j = 1\}$ marks known, readable data positions where physical channel errors occur sparsely ($P(\text{error}) \ll 0.5$, typically $\approx 1/L$).

Let $H \in \mathrm{GF}(2)^{(L-n)\times L}$ denote the standard Toeplitz parity-check matrix governed by a primitive generator polynomial of degree $n = 5$ over GF(2). The global syndrome of recurrence violations is evaluated as $\mathbf{s} = H\mathbf{b} \pmod 2$. To recover the underlying axis polynomial, the framework identifies the optimal configuration via:

$$a^* = \arg\min_{a \in A} \|\mathbf{e}_a\|_0, \tag{35}$$

selecting the axis polynomial a from the ensemble A exhibiting the minimal recovered errors under the constraint that the residual syndrome remains algebraically recoverable.

To avoid complex matrix decompositions over noisy streams, the pipeline formulates error recovery as a greedy footprint elimination problem over $\mathrm{GF}(2)^L$. The sharp contrast between error probabilities in erasure gaps ($P = 0.5$) and readable channel coordinates ($P \ll 0.5$) dictates a two-phase greedy strategy: soft pre-minimization over erasures followed by exact-matching error elimination over known positions.

Causal Horizon and Minimal Buffer Bound.

For a bit-flip residing at column index k, its response $\mathbf{g}_k = \mathbf{h}_k$ (k-th column of matrix H) lies entirely within a bounded causal horizon window of $n + 1$ check equations, spanning from row index $k - n$ up to row index k. As a raw channel error induces paired bit perturbations at adjacent positions $(k - 1, k)$, the combined response vector $\mathbf{g}_k = \mathbf{h}_{k-1} \oplus \mathbf{h}_k$ merges two single-bit signatures, spanning a paired causal horizon of $n + 2$ consecutive check equations over rows $i \in \{k - 1 - n, \ldots, k\}$. Every greedy step resolves an error via a single atomic vector flip $\mathbf{s} \to \mathbf{s} \oplus \mathbf{g}_k$.

By treating the paired bit-flip as a single compound physical perturbation event, its syndrome response simplifies algebraically to an atomic composite column signature. Although the expanded causal horizon spans $n + 2$ equations, this structural extension is strictly compensated for by the reduction in the combinatorial search space: while a single isolated error has L possible configurations within the window, a rigid adjacent pair possesses only $L - 1$ unique positional states. This dimensional shift guarantees that the collision entropy remains equivalent to the single-fault case. Consequently, uniquely resolving this compound signature requires that the number of check equations satisfies the fundamental bound $M \geq n + 1$. Expressed in terms of the window length $L = M + n$, this yields:

$$L_{\min} = 2n + 1 = 11 \text{ bits}. \tag{36}$$

This confirms that $L = 11$ bits ($M = 6$ check equations) represents the exact, tight algebraic bound necessary and sufficient to isolate a paired physical error. This derivation provides an exact physical and algebraic justification for the operational streaming buffer minimal length, showing that it is not merely an arbitrary heuristic choice. Rather, as established by the Hoeffding tail evaluation in Section 3.5 (summarized in Table 2), an observation window of $L = 11$ bits under $e = 2$ represents the tightest theoretical bound that simultaneously satisfies algebraic footprint constraints and guarantees a joint false-alarm probability below $\alpha = 1.0\%$.

Syndrome Masking.

Because value assignment within erasure gaps $\mathcal{U}$ is uninitialized ($P = 0.5$), these positions generate high-variance syndrome noise that can corrupt exact error detection on readable data. To shield the exact-matching pass, the decoder constructs a binary projection mask vector $\mathbf{p} \in \mathrm{GF}(2)^{L-n}$ in the syndrome space. Row i of H is masked out ($p_i = 0$) if it intersects any column corresponding to an erasure gap ($m_j = 0$):

$$p_i = \begin{cases} 0, & \text{if } \exists j \in \mathcal{U} \text{ such that } H_{i,j} \neq 0, \\ 1, & \text{otherwise.} \end{cases} \tag{37}$$

Rather than enforcing complete algebraic equivalence, the channel error elimination phase leverages the binary mask to perform a localized inner product validation over GF(2). The objective is to find a binary channel correction vector $\mathbf{e} \in \mathrm{GF}(2)^L$ that forces the discrepancy between the syndrome violation and the error footprint to lie entirely orthogonal to the unmasked subspace:

$$\langle \mathbf{p},\, \mathbf{s} \oplus H\mathbf{e} \rangle = 0 \implies \mathbf{p}^T(\mathbf{s} \oplus H\mathbf{e}) = 0. \tag{38}$$

Because this condition checks for a zero-norm result after projection, it ensures component-wise algebraic identity strictly on the readable data positions, protecting the exact-matching logic from noise cross-talk generated by random gap values.

Pre-Minimization over Erasure Subspace.

Before enforcing strict zero-residual constraints on valid bits, the pipeline executes an anti-causal sweep ($k = L-1, \ldots, 0$) operating strictly over erasure coordinates ($k \in \mathcal{U}$). Since erasure locations are known *a priori* but hold completely unconstrained values, this stage performs soft Hamming weight reduction on the local syndrome $\mathbf{s}$.

At each step $k \in \mathcal{U}$, an erasure bit flip is greedily committed if and only if its candidate column footprint $\mathbf{x}$ strictly compresses the syndrome weight:

$$\|\mathbf{x} \oplus \mathbf{s}\|_0 < \|\mathbf{s}\|_0, \quad k \in \mathcal{U}, \tag{39}$$

where $\mathbf{x}$ represents either the combined adjacent vector $\mathbf{g}_k = \mathbf{h}_{k-1} \oplus \mathbf{h}_k$ (tested with priority if $k-1 \in \mathcal{U}$) or the isolated column $\mathbf{h}_k$. Upon committing an inversion, the footprint is subtracted globally ($\mathbf{s} \leftarrow \mathbf{s} \oplus \mathbf{x}$). This preliminary pass suppresses noise cross-talk generated by random gap values.

Stage 1: Exact Syndrome Minimization.

Following erasure stabilization, the decoder sweeps anti-causally ($k = L-1, \ldots, 0$) strictly across valid indices ($k \in \mathcal{K}$) minimizing the objective given by Eq. (38). Under the assumption that $P(\text{error}) \ll 0.5$ on readable bits, this stage demands an exact zero-residual match within the unmasked projection subspace:

1. **Double-Error Hypothesis:** If $k-1 \in \mathcal{K}$, the paired column footprint $\mathbf{g}_k = \mathbf{h}_{k-1} \oplus \mathbf{h}_k$ is evaluated over the row slice $\{k-n-1, \ldots, k\}$. The hypothesis is accepted if the masked discrepancy between the footprint and the active syndrome vector is orthogonal to the projection vector:

$$\langle \mathbf{p},\, \mathbf{g}_k \oplus \mathbf{s} \rangle = 0. \tag{40}$$

   Upon verification, indices $\{k-1, k\}$ are logged into the candidate set E, and the syndrome is updated globally: $\mathbf{s} \leftarrow \mathbf{s} \oplus \mathbf{g}_k$.
2. **Single Isolated Bit-Flip Hypothesis:** Otherwise, the individual column $\mathbf{h}_k$ is cross-referenced over rows $\{k-n, \ldots, k\}$. The single flip is accepted if:

$$\langle \mathbf{p},\, \mathbf{h}_k \oplus \mathbf{s} \rangle = 0. \tag{41}$$

   Upon validation, index k is added into E, and the residual syndrome is flushed: $\mathbf{s} \leftarrow \mathbf{s} \oplus \mathbf{h}_k$.

If $\mathbf{s} = \mathbf{0}$ at the end of this pass, the stream is declared fully corrected, bypassing the syndrome reconstruction.

Stage 2: Syndrome Reconstruction.

If the greedy approach fails and the final syndrome remains non-zero, we reformulate the task as a Gauss-Jordan elimination problem. All accumulated discrete correction indices E and physical erasure gaps $\mathcal{U}$ are merged into an expanded set of candidate variables $\mathcal{U}' = \mathcal{U} \cup E = \{u_1, u_2, \ldots, u_N\}$, leaving the remaining positions strictly fixed as reliable data $\mathcal{K}' = \mathcal{K} \setminus E$.

To project the code space onto the candidate subspace, we introduce an axis selector matrix $S_{\mathcal{U}'} \in \mathrm{GF}(2)^{L \times N}$ that extracts only the matrix columns corresponding to indices in $\mathcal{U}'$. Right-multiplying the full parity-check matrix $H \in \mathrm{GF}(2)^{(L-n) \times L}$ by the selection matrix $S_{\mathcal{U}'}$ strips away all columns belonging to reliable positions $\mathcal{K}'$, yielding the reduced matrix $H' = HS_{\mathcal{U}'} \in \mathrm{GF}(2)^{(L-n) \times N}$. Reconstructing the sequence directly from the initial full syndrome vector $\mathbf{s}$ formulates the compact system over GF(2) for the sub-space error vector $\mathbf{e}' = S_{\mathcal{U}'}^{\top}\mathbf{e}$:

$$H'\mathbf{e}' = \mathbf{s}. \tag{42}$$

The augmented matrix $[H' \mid \mathbf{s}]$ is solved via Gauss-Jordan elimination over GF(2) to Row Echelon Form (REF):

1. **Exact Reconstruction:** If $\operatorname{rank}(H') = N$ and the system is algebraically consistent, the solver uniquely extracts $\mathbf{e}'$. The global correction vector is restored as $\mathbf{e} = S_{\mathcal{U}'}\mathbf{e}' \in \mathrm{GF}(2)^L$ and applied to produce the clean sequence $\mathbf{b}_{\text{final}} = \mathbf{b} \oplus \mathbf{e}$.
2. **Contradiction / Ambiguity Rejection:** If the reduction reveals an algebraic contradiction or $\operatorname{rank}(H') < N$ (an underdetermined system with free variables), the candidate hypothesis is rejected, forcing the decoder into a safe no-decision state.

3.7 Camera Intrinsic and Distortion Calibration on Hexagonal Lattice

Traditional camera calibration architectures typically enforce a joint optimization layout, simultaneously minimizing reprojection errors over both linear lens intrinsics (focal length f, optical center $\mathbf{c}$) and non-linear radial lens distortion parameters ($\boldsymbol{\kappa}$). However, such monolithic paradigms are governed by severe parametric cross-talk, which introduces parameter drift under weak perspective conditions [24]. To break this ill-conditioned coupling, we introduce an aperture-decoupled photogrammetric pipeline. This framework isolates distortion estimation from intrinsic matrix recovery by operating within an invariant, pixel-bounded coordinate domain, followed by a direct closed-form intrinsic solution on rectified grids.

This decoupled strategy builds directly upon the classical photogrammetric plumb-line paradigm pioneered by Duane Brown [25]. Indeed, the creator of the ubiquitous radial distortion model originally invented the plumb-line technique specifically to calibrate lens distortion independently of external orientation, principal point, and focal parameter spaces. By photographing perfectly tensioned weighted lines (plumb-lines) and optimizing radial coefficients solely based on a collinearity criterion, Brown isolated the lens distortion profile from high-dimensional parameter cross-talk. Mark Shortis subsequently formalized this operational sequence, demonstrating that utilizing a plumb-line pass to lock down radial parameters, followed by closed-form intrinsic tracking while keeping the lens profile strictly fixed, represents the most mathematically sound approach to defeating the strong parameter correlations inherent to monolithic calibrations [26].

Geometric Invariance of Line Straightness on Hexagonal Axes.

The validity of this decoupling approach follows directly from first principles: the collinearity of a segment is a central projective invariant, strictly preserved under linear pinhole projection but violated solely by non-linear optical distortion. Within our framework, continuous collinear tracks are constructed along the three primary topological axes $(0^\circ, 60^\circ, 120^\circ)$ of the hexagonal lattice, supplemented by the orthogonal 90° directions derived from lattice transversals. Whether a projected grid lane remains straight or exhibits radial curvature is an intrinsic property of the sensor's raw image plane, entirely independent of the camera's downstream focal length f.

Because focal parameters merely act as linear scaling transformations that stretch or compress the spatial canvas without introducing or mitigating spatial curvature, they have no physical or mathematical justification for entering the distortion-rectification operator. Forcing focal dependencies into the lens-unwarping formulation introduces an artificial coupling that severely degrades numerical stability.

Parametric Coupling in the Brown-Conrady Formulation.

The standard Brown-Conrady lens model maps an ideal, undistorted normalized coordinate vector $\mathbf{x}_n = (x_n, y_n)^\top$ to its physically warped location via a non-linear radial scaling factor $\mathcal{D}(\boldsymbol{\kappa}) = 1 + \kappa_1 r_n^2 + \kappa_2 r_n^4 + \dots$ [21, 24]. Crucially, this formulation inherently implies that the center of radial symmetry coincides with the principal optical axis. These normalized single-frame coordinates are tied explicitly to both the focal lengths (f_x, f_y) and the principal point $\mathbf{c}_0 = (c_x, c_y)^\top$:

$$x_n = \frac{x_{\text{pixel}} - c_x}{f_x}, \quad y_n = \frac{y_{\text{pixel}} - c_y}{f_y}, \quad r_n^2 = x_n^2 + y_n^2. \tag{43}$$

Expanding this transformation explicitly exposes the structural dependency connecting the focal scale space to the polynomial distortion coefficients:

$$\mathcal{D}(\boldsymbol{\kappa}) = 1 + \kappa_1 \left(\frac{(x_{\text{pixel}} - c_x)^2}{f_x^2} + \frac{(y_{\text{pixel}} - c_y)^2}{f_y^2} \right) + \dots. \tag{44}$$

Under single-frame observation where perspective foreshortening gradients are shallow, any numerical shift in the estimated principal point $\mathbf{c}_0$ or focal scale factor f can be almost completely absorbed by a reciprocal translation or scaling of $\boldsymbol{\kappa}$. This high-dimensional parameter cross-talk flattens the optimization energy landscape into an ill-conditioned, elongated valley, causing traditional joint solvers to fall into non-physical local minima or drift along degenerate parametric manifolds.

Aperture-Normalized Geometric Decoupling.

To resolve this structural coupling, our framework exploits the implicit Brown-Conrady property dictating that the center of radial distortion and the principal point occupy the exact same physical coordinates on the image plane. By capitalizing on this geometric intersection, the proposed pipeline decouples the parameter space. We explicitly isolate and solve for the principal point $\mathbf{c}_0 = [c_x, c_y]^\top$ purely as the center of radial distortion in a focal-free domain, subsequently leveraging this estimation to compute the intrinsics matrix $\mathbf{K}$ in closed form.

Mathematically, the camera projection model is reformulated into a strictly decoupled multi-step pipeline:

$$\mathbf{x}_u = \mathbf{K} \cdot \mathbf{P} \cdot \mathbf{X}, \quad \mathbf{x}_d - \mathbf{c}_0 = (\mathbf{x}_u - \mathbf{c}_0)\mathcal{D}(\boldsymbol{\kappa}; \mathbf{x}_u), \tag{45}$$

where $\mathbf{X}$ represents the 3D world coordinates, $\mathbf{P}$ is the extrinsic camera pose matrix, $\mathbf{x}_u$ denotes the ideal undistorted pixel coordinates, and $\mathbf{x}_d$ represents the distorted pixel coordinates captured on the sensor matrix.

Standard calibration routines optimize $\mathbf{K}$ and the lens distortion operator $\mathcal{D}$ simultaneously, triggering parameter interference. To break this dependency, our strategy calibrates $\mathcal{D}$ completely independently in the first pass, enabling the immediate recovery of the rectified metric field $\mathbf{x}_u$ and the subsequent closed-form extraction of $\mathbf{K}$ without cross-talk.

To achieve focal-free optimization, we replace the standard focal-normalized coordinates with a static, resolution-bounded metric: the *Aperture-Normalized Radius* (r_{ap}). We define a fixed bounding horizon radius R_{ap} as the maximum Euclidean distance from the principal point to the furthest corner vertex of the active sensor frame bounds. The independent distortion mapping operates strictly within this pixel-bounded domain, guaranteeing that the normalized radial field remains strictly bounded ($r_{\text{ap}} \leq 1.0$) across the active canvas:

$$\mathbf{x}_d - \mathbf{c}_0 = (\mathbf{x}_u - \mathbf{c}_0)\left(1 + \kappa_{\text{ap}} \frac{\|\mathbf{x}_u - \mathbf{c}_0\|^2}{R_{\text{ap}}^2}\right). \tag{46}$$

Because R_{ap} is determined entirely by the static sensor bounds, the aperture-normalized distortion coefficient κ_{ap} is mathematically immune to focal scale drift. This enables the solver to straighten the hexagonal lattice axes purely within raw sensor coordinates, producing a rectified spatial topology $\mathbf{x}_u$ optimized for the closed-form recovery of $\mathbf{K}$.

Finally, once the linear focal parameters are resolved, the standard Brown-Conrady representation is recovered by projecting the aperture coefficient back to the effective focal plane:

$$\kappa_1 = \kappa_{\text{ap}} \cdot \left(\frac{f_{\text{eff}}}{R_{\text{ap}}}\right)^2, \tag{47}$$

where the optimal scalar f_{eff} is defined via the harmonic mean of the squared focal parameters: $1/f_{\text{eff}}^2 = \frac{1}{2}(1/f_x^2 + 1/f_y^2)$. As proven in Lemma 1, this mapping corresponds strictly to the Chebyshev circular conic approximation of the Image of the Absolute Conic (IAC) matrix manifold under a Frobenius matrix norm.

Connection to the Thin Prism Model.

The physical manifestation of lens-to-sensor tilt is classically formalized in photogrammetry via the Thin Prism Model [25]. This paradigm dictates that a structural tilt transforms the isotropic projection cone into an anisotropic, non-coaxial optical wedge, introducing tangential bending and rotating the principal distortion ellipse by an arbitrary wedge angle. In conventional monolithic solvers, fitting standard Cartesian radial polynomials to this tilted profile triggers a severe parameter drift.

In contrast, our aperture-decoupled framework effectively bypasses this non-coaxial coupling. By performing the first-pass line rectification within an inherently rotation-invariant, isotropic aperture

domain, the pipeline isolates the geometric unwarping from the wedge orientation. As it is shown in Lemma 1, the cross-axial components induced by the optical wedge can be identically nullified on the diagonal manifold. This guarantees that the recovered effective focal length f_{eff} remains an invariant from assembly-induced tilt artifacts.

3.8 Decoupled Radial Lens Distortion Estimation via Spatial Plumb-Line Invariants

Traditional camera calibration pipelines typically solve for the pinhole projection matrix and lens distortion parameters concurrently using strict point-to-point correspondence optimization. However, such methods are prone to numerical parameter coupling traps, where errors in the estimated focal lengths (f_x, f_y) directly propagate into, and corrupt, the radial distortion coefficients (κ). To eliminate these parameter correlation vectors, we implement a decoupled calibration paradigm derived from the classical photogrammetric plumb-line approach [25], which states that the projection of a straight line in the physical world must remain perfectly straight on the normalized image plane if radial distortion is completely nullified.

By leveraging the topological continuity of the fully decoded B-HGP hexagonal lattice, parallel target tracks are extracted and organized into isolated directional linear manifolds. Utilizing the hexagonal topology matrix $\mathcal{M}$, these collinear sub-pixel node tracking paths are collected across the U, V, W, and vertical axis H to compose a globally coordinated geometric asset designated as the frame bundle:

$$\mathcal{B} = \mathcal{L}_U \cup \mathcal{L}_V \cup \mathcal{L}_W \cup \mathcal{L}_H. \tag{48}$$

By isolating the distortion extraction loop from the camera's projective projection layers, we solve for the radial profile inside a low-dimensional parameter space: $\mathbf{\Theta}_{\text{ap}} = [\kappa_{\text{ap}}, c_x, c_y]^\top$. Because the objective function is computed purely on the geometric straightness of the lanes rather than algebraic reprojection errors, the estimation loop requires zero initial approximations or knowledge regarding the effective focal variables.

Menger Curvature Optimization and Spatial Weighting

Our straightness cost function evaluates a local geometric penalty driven by the discrete Menger curvature $c(\mathbf{p}_A, \mathbf{p}_B, \mathbf{p}_C) = 1/R$ of a point triplet, where R is the circumscribed over $\triangle ABC$ circle radius. Defining the triangle edge vectors as $\mathbf{a} = \mathbf{p}_B - \mathbf{p}_A$, $\mathbf{b} = \mathbf{p}_C - \mathbf{p}_B$, and $\mathbf{c} = \mathbf{p}_C - \mathbf{p}_A$, this curvature is directly proportional to the area of the spanned triangle ABC divided by the product of its side lengths [27]:

$$c(\mathbf{p}_A, \mathbf{p}_B, \mathbf{p}_C) \propto \frac{\text{Area}(\triangle ABC)}{\|\mathbf{a}\| \cdot \|\mathbf{b}\| \cdot \|\mathbf{c}\|} \propto \frac{\|\mathbf{a} \times \mathbf{b}\|}{\|\mathbf{a}\| \cdot \|\mathbf{b}\| \cdot \|\mathbf{c}\|}. \tag{49}$$

Because constant scaling multipliers do not alter the position of the objective function's extrema, they are dropped for numerical efficiency. To evaluate local bending, we extract a symmetric node triplet $(\mathbf{p}_{i-s}, \mathbf{p}_i, \mathbf{p}_{i+s})$ centered at node $\mathbf{p}_i$ along a rectified lane segment $L \in \mathcal{B}$ with a discrete stride s. The discrete Menger curvature loss across the bundle is formulated as:

$$\mathcal{L}_{\text{Menger}}(\mathbf{\Theta}_{\text{ap}}) = \sum_{L \in \mathcal{B}} \sum_i w_i \cdot \frac{|(\mathbf{p}_i - \mathbf{p}_{i-s}) \times (\mathbf{p}_{i+s} - \mathbf{p}_i)|}{\|\mathbf{p}_i - \mathbf{p}_{i-s}\| \cdot \|\mathbf{p}_{i+s} - \mathbf{p}_i\| \cdot \|\mathbf{p}_{i+s} - \mathbf{p}_{i-s}\|}. \tag{50}$$

To achieve true sub-pixel calibration precision, two key spatial factors govern the formulation of Eq. (50):

1. **Symmetric Central Node Weighting (w_i):** Due to out-of-plane perspective slants and spatial scaling variations across the frame, projected target features exhibit non-uniform physical sizes on the sensor. Control points near the optical center appear at maximum blob size, whereas peripheral features undergo significant geometric compression. For larger blobs, boundary blur, pixel integration effects, and raster noise are distributed over a wider spatial area, which increases the variance and uncertainty of the estimated center. Conversely, for highly compressed target features, the physical center is geometrically constrained within a footprint of just a few pixels, which drastically reduces the maximum localization error. To leverage this physical reliability, a scalar weight $w_i \propto 1/r_i$ is attached explicitly to the central node $\mathbf{p}_i$ of each triplet based on its

detected blob radius r_i. This scheme naturally assigns higher confidence to the tightly bounded, high-precision peripheral positions where line bending is most pronounced, yielding a per-triplet loss contribution $\mathcal{L}_i \propto (1/r_i) \cdot (1/R_i)$.

2. **Chord-Length Sensitivity Control (s):** The stride parameter s defines the geometric baseline chord length $\|\mathbf{p}_{i+s} - \mathbf{p}_{i-s}\|$. As lines become progressively straight during optimization, the cross-product numerator in Eq. (50) vanishes, making the loss vulnerable to raster quantization noise. Expanding the stride s increases the triangle baseline, lifting the geometric residual above the noise floor and maintaining gradient sensitivity near convergence.

Cascaded Nelder–Mead Simplex Optimization.

To minimize the non-linear Menger curvature loss $\mathcal{L}_{\text{Menger}}(\mathbf{\Theta}_{\text{ap}})$, the pipeline exploits a modified derivative-free Nelder–Mead search algorithm [28]. This choice is motivated by the absence of analytical Jacobians for discrete curvature mappings and the presence of sub-pixel quantization noise.

To prevent optimization bias caused by the disparate physical units of radial distortion coefficients and pixel coordinates, all parameters are dynamically mapped onto a normalized, dimensionless workspace $\mathbf{x} = [\kappa_{\text{norm}}, x_{\text{norm}}, y_{\text{norm}}]^\top$. The optimization operates through a coarse-to-fine cascaded search scheme, dynamically coordinating two operational scales: the stride s along lines and the parameter search domain scale δ. The optimization progresses iteratively across discrete search stages $k \in \{1, 2, \ldots, K\}$ as follows:

1. **Initialization and Tetrahedron Setup:** At stage k, the search domain is centered at the current best parameter estimate $\mathbf{\Theta}_{\text{ap}}^{(k)}$. A non-degenerate, perfectly isotropic regular tetrahedron (4 vertices in $\mathbb{R}^3$) is constructed in the normalized parameter space with a characteristic edge length $\delta^{(k)}$. Edge isotropy ensures uniform volume expansion and contraction during Nelder–Mead steps, preventing dimensional collapse along any single parameter axis.
2. **Coarse Stage (Compact Stride s and Wide Search δ):** During initial convergence iterations, strong global curvature guarantees an adequate optimization signal. The stride parameter is kept moderate ($s = s_{\text{coarse}}$) to capture local non-linear warping without over-smoothing. Simultaneously, the parameter search scale $\delta^{(1)}$ is initialized broadly to enclose expected physical parameter tolerances.
3. **Fine Stage (Stride Expansion s and Narrow Search δ):** As intermediate convergence thresholds are met ($\texttt{fatol} < \tau^{(k)}$), the solver transitions to stage $k+1$. To prevent numerical gradient stagnation as target lines flatten, the stride parameter is systematically increased ($s_{\text{fine}} > s_{\text{coarse}}$). This expands the baseline chord length of the evaluated triplets, amplifying sub-pixel residual bends above the sensor quantization floor. Concurrently, the parameter space re-centers, and the search tetrahedron contracts ($\delta^{(k+1)} = \eta \cdot \delta^{(k)}$, where $\eta < 1$).

By expanding the geometric evaluation baseline s while recursively tightening the search domain δ, the framework preserves numerical gradient responsiveness at sub-pixel levels. This dual-scale adaptation bypasses floating-point precision constraints, enabling strict convergence criteria to isolate the exact optical center $\mathbf{c}_0$ and aperture distortion parameter κ_{ap}.

3.9 Closed-Form Intrinsic Recovery via DLT

Once distortion optimization completes, the estimated parameter vector $[\kappa_{\text{ap}}, c_x, c_y]^\top$ is used to rectify all lines in the frame bundle $\mathcal{B}$. Because the grid target tracks are straightened, the remaining linear intrinsic parameters can be extracted using an analytical closed-form solver without iterative reprojection loops.

This recovery relies on two geometric invariants. First, because the calibration target layout provides parallel lines along the lattice structures, we compute the intersections of these parallel bundles to isolate the raw vanishing points. All corresponding spatial model direction vectors $\mathbf{V}_i^0$ lie strictly on the infinitely distant line of the target plane:

$$\mathbf{V}_i^0 = (u_i^0, v_i^0, 0)^\top, \tag{51}$$

where $i \in \{U, V, W, H\}$ indexes the directional lattice channels.

Second, to isolate focal scaling degrees of freedom from translational shifts, the intrinsic camera matrix $\mathbf{K}$ is factored into an optical center translation operator $\mathbf{T}_c$ and an isotropic focal scaling

diagonal $\mathbf{K}_f$:

$$\mathbf{K} = \mathbf{T}_c\mathbf{K}_f = \begin{bmatrix} 1 & 0 & c_x \\ 0 & 1 & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} f_x & 0 & 0 \\ 0 & f_y & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{52}$$

For raw vanishing points extracted in homogeneous coordinates $\mathbf{V}_i = [V_{x,i}, V_{y,i}, V_{w,i}]^\top$, we apply $\mathbf{T}_c^{-1}$ to decouple them from principal point offsets and shift the origin directly to the optical center. For each channel, the centered homogeneous vanishing point vectors $\mathbf{V}_i^c$ are formulated as:

$$\mathbf{V}_i^c = \mathbf{T}_c^{-1}\mathbf{V}_i = \begin{bmatrix} 1 & 0 & -c_x \\ 0 & 1 & -c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} V_{x,i} \\ V_{y,i} \\ V_{w,i} \end{bmatrix} = \begin{bmatrix} V_{x,i} - c_x V_{w,i} \\ V_{y,i} - c_y V_{w,i} \\ V_{w,i} \end{bmatrix} = \begin{bmatrix} u_i \\ v_i \\ w_i \end{bmatrix}. \tag{53}$$

For a given camera pose, the structural mapping between centered image vanishing points $\mathbf{V}_i^c = [u_i, v_i, w_i]^\top$ and 3D model directions $\mathbf{V}_i^0$ is governed by projective collinearity $\mathbf{V}_i^c \times (\mathbf{H}\mathbf{V}_i^0) = \mathbf{0}$. Expanding this cross-product directly yields three linear equations per directional channel. Dropping the third redundant row, the first two linearly independent rows construct the local Direct Linear Transformation (DLT) system matrix block $\mathbf{A}_{\text{DLT}}\mathbf{h} = \mathbf{0}$:

$$\begin{bmatrix} w_i u_i^0 & w_i v_i^0 & 0 & 0 & -u_i u_i^0 & -u_i v_i^0 \\ 0 & 0 & w_i u_i^0 & w_i v_i^0 & -v_i u_i^0 & -v_i v_i^0 \end{bmatrix} \begin{bmatrix} h_{11} \\ h_{12} \\ h_{21} \\ h_{22} \\ h_{31} \\ h_{32} \end{bmatrix} = \mathbf{0}. \tag{54}$$

Evaluating the Singular Value Decomposition (SVD) of the global matrix $\mathbf{A}_{\text{DLT}}$ yields the unique frame direction homography parameter vector $\mathbf{h} = [h_{11}, h_{12}, h_{21}, h_{22}, h_{31}, h_{32}]^\top$ corresponding to the minimum singular value.

Selective View Pooling Strategy

Following Zhang [24], the columns of the extracted homography $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2]$ must satisfy geometric constraints imposed by the Image of the Absolute Conic (IAC), $\boldsymbol{\omega} = \mathbf{K}^{-\top}\mathbf{K}^{-1} = \text{diag}(g_x, g_y, 1)$, where $g_x = 1/f_x^2$ and $g_y = 1/f_y^2$. Enforcing orthogonality ($\mathbf{h}_1^\top\boldsymbol{\omega}\mathbf{h}_2 = 0$) and equal scale ($\mathbf{h}_1^\top\boldsymbol{\omega}\mathbf{h}_1 - \mathbf{h}_2^\top\boldsymbol{\omega}\mathbf{h}_2 = 0$) generates a local 2×2 linear system $\mathbf{M}\mathbf{g} = \mathbf{N}$:

$$\begin{bmatrix} h_{11}h_{12} & h_{21}h_{22} \\ h_{11}^2 - h_{12}^2 & h_{21}^2 - h_{22}^2 \end{bmatrix} \begin{bmatrix} g_x \\ g_y \end{bmatrix} = \begin{bmatrix} -h_{31}h_{32} \\ h_{32}^2 - h_{31}^2 \end{bmatrix}. \tag{55}$$

To ensure numerical stability against scale variations across views, each local equation block is normalized by its inverse Frobenius norm $s = 1/\|\mathbf{M}\|_F$.

Stage 1: Multi-View Anisotropic Focal Estimation ($f_x \neq f_y$).

The two-degree-of-freedom solver (g_x, g_y) relies strictly on observable perspective gradients. Views exhibiting flat affine behavior ($h_{31}, h_{32} \to 0$) carry zero information regarding directional focal anisotropy; including them introduces pure raster noise into the system, which destabilizes the solver.

A frame $k \in \{1, \ldots, K\}$ is admitted into the asymmetric subset $\mathcal{K}_{\text{asym}}$ (comprising K' valid frames) if its local 2×2 system block matrix $\mathbf{M}^{(k)} \in \mathbb{R}^{2\times 2}$ is well-conditioned ($|\det(\mathbf{M}^{(k)})| > \tau_{\text{det}}$) and yields a valid aspect ratio ($0.4 < f_x/f_y < 2.5$). Admitted frames append their scaled blocks $\mathbf{M}^{(k)}$ and vectors $\mathbf{N}^{(k)} \in \mathbb{R}^{2\times 1}$ to an asymmetric system:

$$\mathbf{M}_{\text{asym}}\mathbf{g} = \mathbf{N}_{\text{asym}}, \quad \text{where } \mathbf{M}_{\text{asym}} = \begin{bmatrix} s^{(1)}\mathbf{M}^{(1)} \\ \vdots \\ s^{(K')}\mathbf{M}^{(K')} \end{bmatrix}, \quad \mathbf{N}_{\text{asym}} = \begin{bmatrix} s^{(1)}\mathbf{N}^{(1)} \\ \vdots \\ s^{(K')}\mathbf{N}^{(K')} \end{bmatrix}. \tag{56}$$

Solving this system via Ordinary Least Squares (OLS), $\mathbf{M}_{\text{asym}}^{\top}\mathbf{M}_{\text{asym}}\mathbf{g} = \mathbf{M}_{\text{asym}}^{\top}\mathbf{N}_{\text{asym}}$, decouples g_x and g_y into explicit scalar components over the filtered subset $k \in \mathcal{K}_{\text{asym}}$:

$$g_x = \frac{C \cdot D_x - B \cdot D_y}{\Delta}, \qquad g_y = \frac{A \cdot D_y - B \cdot D_x}{\Delta}. \tag{57}$$

where the normal system matrix $\mathbf{M}_{\text{asym}}^{\top}\mathbf{M}_{\text{asym}} = \begin{bmatrix} A & B \\ B & C \end{bmatrix}$ and vector $\mathbf{M}_{\text{asym}}^{\top}\mathbf{N}_{\text{asym}} = \begin{bmatrix} D_x \\ D_y \end{bmatrix}$ are explicitly accumulated as:

$$\begin{bmatrix} A & B \\ B & C \end{bmatrix} = \sum_{k \in \mathcal{K}_{\text{asym}}} (s^{(k)})^2 \begin{bmatrix} (M_{11}^{(k)})^2 + (M_{21}^{(k)})^2 & M_{11}^{(k)}M_{12}^{(k)} + M_{21}^{(k)}M_{22}^{(k)} \\ M_{11}^{(k)}M_{12}^{(k)} + M_{21}^{(k)}M_{22}^{(k)} & (M_{12}^{(k)})^2 + (M_{22}^{(k)})^2 \end{bmatrix}, \tag{58}$$

$$\begin{bmatrix} D_x \\ D_y \end{bmatrix} = \sum_{k \in \mathcal{K}_{\text{asym}}} (s^{(k)})^2 \begin{bmatrix} M_{11}^{(k)}N_1^{(k)} + M_{21}^{(k)}N_2^{(k)} \\ M_{12}^{(k)}N_1^{(k)} + M_{22}^{(k)}N_2^{(k)} \end{bmatrix}. \tag{59}$$

Here, A and C scale the primary directional responses along the X and Y axes, B captures cross-axis correlation, and $\Delta = AC - B^2$ is the system determinant. If $\Delta > \epsilon$ and both estimates satisfy physical sensor boundaries ($1/f_{\max}^2 < g_x, g_y < 1/f_{\min}^2$), the directional focal lengths are assigned as $f_x = 1/\sqrt{g_x}$ and $f_y = 1/\sqrt{g_y}$.

Stage 2: Isotropic Fallback ($f_x = f_y = f$).

If Stage 1 fails due to rank deficiency ($\Delta \to 0$) or ill-conditioned view distributions, the framework contracts the search space to a symmetric model ($g_x = g_y = g = 1/f^2$). Unlike the asymmetric solver, the symmetric formulation does not require strict perspective gradient filtering; hence, it pools all available frames $k \in \{1, \dots, K\}$ to maximize statistical robustness.

Contracting each block via $\mathbf{M}_{\text{sym}}^{(k)} = \mathbf{M}^{(k)}\mathbf{1}$ (where $\mathbf{1} = [1, 1]^{\top}$) over all K frames projects the system into a robust scalar estimator:

$$g = \frac{\sum_{k=1}^{K} (s^{(k)})^2 \left((M_{11}^{(k)} + M_{12}^{(k)})N_1^{(k)} + (M_{21}^{(k)} + M_{22}^{(k)})N_2^{(k)} \right)}{\sum_{k=1}^{K} (s^{(k)})^2 \left((M_{11}^{(k)} + M_{12}^{(k)})^2 + (M_{21}^{(k)} + M_{22}^{(k)})^2 \right)}. \tag{60}$$

If $1/f_{\max}^2 < g < 1/f_{\min}^2$, the uniform focal length is accepted as $f = 1/\sqrt{g}$.

Rejection of Affine Fallbacks and Distortion Update.

In a purely affine regime ($h_{31}, h_{32} \to 0$), the physical focal length f and pattern depth become fundamentally uncoupled, causing f to scale dynamically with the relative depth of each frame. Pooling such views within a multi-frame estimator is mathematically ill-posed.

To preserve strict calibration integrity, our pipeline explicitly rejects heuristic affine fallbacks for multi-view intrinsic extraction. If both Stage 1 and Stage 2 fail, the solver returns a structural degeneracy alert, enforcing the acquisition of non-parallel target poses with sufficient perspective tilt.

Finally, the standard Brown–Conrady radial distortion coefficient is recovered by projecting the optimized aperture-normalized coefficient back to the active focal plane with Eq. (47).

4 Processing Pipeline and Implementation Notes

4.1 Single-Frame Calibration Workflow

The pipeline executes sequentially through a series of deterministic processing phases:

1. **Feature Extraction:** Detects candidate nodes and classifies their topological profiles (circle versus triangle features) to isolate raw code-space primitives (see Section 3.2).
2. **Lattice Reconstruction:** Assembles unorganized sub-pixel barycenters into stable hexagonal coordinate networks via a Delaunay nucleation and front-propagation queue, producing isolated, internally coherent crystalline topological graphs (see Section 4.2).

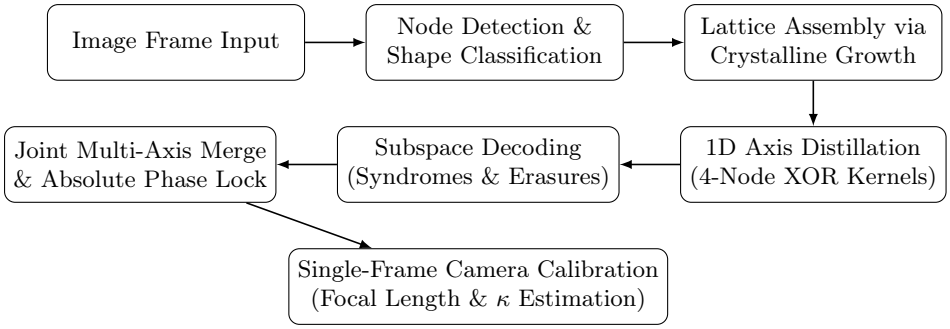


Fig. 2: B-HGP processing pipeline architecture, decoupling structural spatial grid assembly from downstream algebraic phase decoding.

3. **Axis Distillation:** Evaluates each independent crystal segment by applying local 4-node XOR sliding kernels, isolating the target dU and dW' channels into distinct 1D stream vectors representing noisy fragments of an M-sequence (see Section 3.4).
4. **Subspace Decoding:** Processes each extracted 1D stream through the adaptive M-sequence analyzer pipeline, repairing bit-flips and filling missing dropouts via Gauss-Jordan erasure resolution (see Section 3.6).
5. **Coordinate Phase Locking:** Merges the resolved phase positions of the independent orthogonal axes to establish absolute, global (u, v, w) coordinate vectors and maps them directly to the reference template rows and columns (see Section 3.3).
6. **Geometric Optimization:** Leverages the fully mapped point correspondences to perform high-precision single-frame camera calibration, optimizing intrinsic focal length parameters (see Section 3.9) and the primary radial distortion coefficient κ (see Section 3.8).

4.2 Topological Lattice Assembly via Delaunay Front-Propagation

While the decoding process is straightforward, assembling the lattice from an unorganized point cloud is a non-trivial phase of the algorithm. This section is entirely devoted to resolving this problem.

Following the initial feature extraction, the system delivers a two-dimensional point cloud of sub-pixel coordinates heavily contaminated by sensory erasures and false background detections. To reconstruct a globally ordered coordinate grid without incurring the exponential complexity of combinatorial graph matching, the system implements a spatial grid assembly framework inspired by the physical crystallization of hexagonal lattices. The framework operates sequentially by isolating geometric structural assembly, resolving the topological network up to an arbitrary global discrete translation and rotation.

Delaunay Seed Allocation.

The unorganized point cloud is first structured into a global proximity network via 2D Delaunay triangulation. To prevent long-range connections on the convex hull, an upper boundary threshold $d_{\max}$ is established using the median edge length augmented by its median absolute deviation (MAD).

Because perspective distortion and optical aberrations introduce non-uniform spatial scaling, candidates cannot be reliably evaluated using a global geometric metric. Instead, the algorithm employs a localized, scale-invariant observer based on the isotropic inscribed radius $r_{\text{in}} = 2\mathcal{A}/\mathcal{P}$ of each triangular cell, where $\mathcal{A}$ and $\mathcal{P}$ denote its area and perimeter, respectively. By accumulating spatial running averages of r_{in} across partitioned canvas regions, the observer constructs a baseline map of representative cell scales robust to non-uniform distortion.

Triangular faces exhibiting minimal scale deviation relative to their local baseline are retained and oriented counter-clockwise (CCW). To guarantee dense coverage across the entire grid while preventing spatial clustering of seeds, Spatial Non-Maximum Suppression (NMS) is applied directly between candidate triangles. Using a suppression radius tuned to the local grid scale, NMS greedily selects the most representative triangle within each region, yielding a spatially uniform, uncorrupted backbone of seeds for wave propagation.

Vertex-Reflection Queue and Frontier Propagation.

Starting from the established seed faces, the labeled sub-graphs expand outward using a front-propagation queue. For any active boundary triangle with vertices $(\mathbf{p}_a, \mathbf{p}_b, \mathbf{p}_c)$, the engine predicts the physical position of the reflected vertex $\mathbf{p}'_a$ via vector lattice addition $\mathbf{p}'_a = \mathbf{p}_b + \mathbf{p}_c - \mathbf{p}_a$ (Eq. 8). To map this analytical projection back to physical sensor data, a spatial k-d tree is queried for the two nearest candidate points ($k = 2$) within a bounded tolerance radius $r \leq \alpha \cdot s$, where s denotes the baseline lattice step, and the empirical tolerance is configured to $\alpha = 0.10$, remaining strictly below the theoretical upper bound ($\alpha < 0.35$) derived in Section 3.1.

Candidate ambiguity is evaluated directly using Lowe's ratio test with $\tau_{\text{Lowe}} = 0.9$. Rather than guarding against adjacent lattice nodes—which are inherently isolated by the compressed Voronoi uncertainty zone (Section 3.1)—this ratio test primarily acts as a safeguard against detector artifacts and phantom responses, where optical noise yields duplicate detections near a single true target, triggering an immediate rejection of the ambiguous candidates.

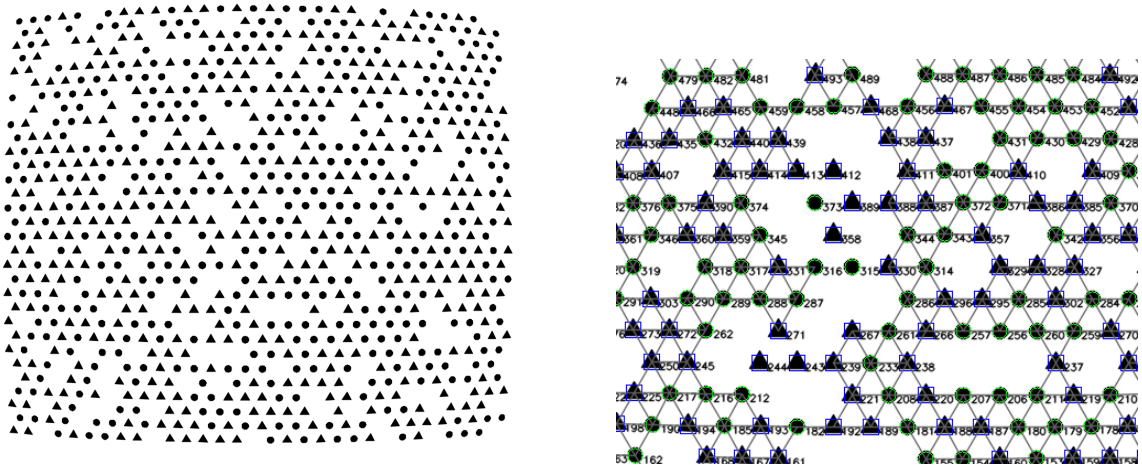
Any candidate node $\mathbf{p}_{\text{new}}$ undergoes topological verification to confirm geometric non-degeneracy via a signed area pseudo-scalar threshold ($(\mathbf{p}_b - \mathbf{p}_{\text{new}}) \times (\mathbf{p}_c - \mathbf{p}_{\text{new}}) < -S_{\min}$) alongside a backward edge-deformation check bounded by the same tolerance parameter. Successfully verified nodes receive global topological coordinates and enter the active propagation queue.

Disjoint-Set Expansion.

To manage concurrent seeds growth across distinct regions of the image frame, each expanding domain is tracked as an isolated crystal island inside a Disjoint-Set Union (DSU). When an edge reflection yields a valid node match index, the algorithm evaluates a tri-modal topological routing rule to maintain grid properties:

- **Cycle Closure:** If the matched node already belongs to the current crystal root component, the decoder executes a metric loop verification, ensuring that cyclic paths close onto the tracking lattice without phase deviations.
- **Island Fusion:** If the expanding front collides with a neighboring independent crystal, an island boundary clash occurs. The algorithm invokes a fusion routine to project collinear directional rays along the coordinates of the underlying axes. The sub-lattices are fused if and only if their spatial tracks align seamlessly across the junction, completely eliminating axis misalignment drift.
- **Single Point Absorption:** If the front intersects an isolated singleton node, the local coordinate grid is extended by absorbing the new vertex, and the disjoint components are dynamically merged within the forest.

This propagation loop executes iteratively across sequential wave steps until the Disjoint-Set lattice has fully stabilized. Once stabilized, it is structured into a dense, row-ordered tracking matrix. Finally, the clean data buffer is forwarded directly to the downstream syndrome decoder to resolve the absolute polynomial phase and orientation.



(a) Initial point cloud.

(b) Crystalline front propagation result.

Fig. 3: Topological lattice growth, illustrating the structural expansion of the triangular mesh under a simulated 15% random node dropout constraint.

5 Experimental Results

The codebase includes synthetic benchmarks using controlled flips and erasures and a Blender-based scene generator to systematically analyze tracking performance and structural robustness under severe perspective warping and image noise. A representative simulation sample displaying these adverse conditions is illustrated in Fig. 4.

Two datasets are considered: the first is a rotational sequence with 6 pure roll rotations spaced by 60° alongside a single tilted view to capture perspective effects; the second contains a baseline zero-rotation view and 8 randomly oriented images. The rotational dataset exhibits severe radial distortion ($\kappa_1 = -0.20$) with isotropic focal parameters. In contrast, the tilted dataset features anisotropic focal parameters with a moderate average distortion of $\kappa_1 = -0.15$. The former effectively captures distortion due to its uniform coverage of the image canvas, whereas the latter is designed to evaluate anisotropic focal distances under steep perspective gradients.

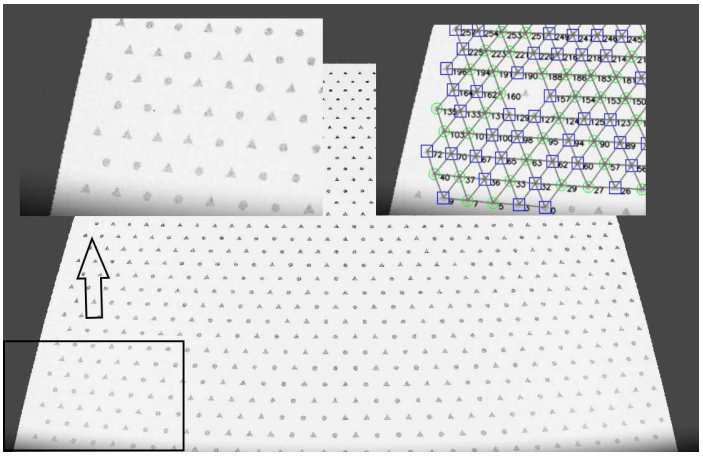


Fig. 4: Blender simulation of a camera tilt under high noise and radial distortion (`oblique_tilt_high`). The inset displays the boundary node degradation and the resulting topological lattice restoration.

5.1 Pattern Registration Performance

The physical dimensions of the pattern features are tightly bounded, yielding an average feature diameter of approximately 10 pixels on the image canvas. Throughout the experiments, the centroid approximation is applied to determine key-point projection centers. Across all evaluation sequences, the low-level spot localization achieves a highly stable pixel-displacement tracking accuracy, with Root Mean Square Error (RMSE) bounded between 0.6 and 0.9 pixels. Consequently, we do not account for localization bias caused by the centroid approximation (see Section 2.1), as the current sub-pixel centroid precision proves completely sufficient for robust parameter estimation.

Table 3: Comprehensive Topological Matching Performance: Comparison Metrics Between Simple Rotational and Random Tilted Sequences

Test Case	GT Nodes	Total Det.	Misclass.	False Det.	TP	Skip (Isolated)	Recall (%)
Simple Rotational Sequence							
`clean_baseline`	961	961	2	0	961	0	100.000
`oblique_tilt_high`	952	889	63	0	888	1	93.277
`roll_120`	924	890	8	8	880	2	95.238
`roll_180`	961	976	14	17	958	1	99.688
`roll_240`	922	636	60	10	561	65	60.846
`roll_300`	924	887	11	9	877	1	94.913
`roll_60`	922	900	12	18	882	0	95.662
Random Tilted Sequence							
`clean_baseline`	961	961	4	0	961	0	100.000
`compound_rotation_0`	916	920	39	24	896	0	97.817
`compound_rotation_1`	928	774	18	5	767	2	82.651
`compound_rotation_2`	440	425	5	29	396	0	90.000
`compound_rotation_3`	957	641	61	2	505	134	52.769
`compound_rotation_4`	450	418	12	8	407	3	90.444
`compound_rotation_5`	633	618	13	38	580	0	91.627
`compound_rotation_6`	944	898	75	7	885	6	93.750
`compound_rotation_7`	953	578	45	20	369	189	38.720

In the presence of heavy illumination noise and aggressive geometric warping, the Galois parity-check prevents lattice tracking mismatches. By filtering out non-convex noise aggregates during feature extraction, the architecture maintains a precision score of 100% after algebraic correction across all tests, yielding zero false positives, misalignments, or phantom nodes.

As quantitatively mapped in Table 3, all metric ratios are evaluated relative to the total number of visible pattern nodes to guarantee that decoding accuracy is benchmarked strictly against valid, in-frame projected keypoints, isolating local reconstruction errors from view-dependent spatial truncation.

Under extreme non-linear perspective slants, non-uniform lighting, and spatial compression, the low-level key-point detection degrades. This triggers a peak node classification failure rate of 7.94% of the entire pattern layout in sequence `compound_rotation_6`, where 75 node shape labels are corrupted. Simultaneously, extreme perspective tilts force massive segments of the pattern out of the active camera viewport, inducing a maximum pattern omission rate of 61.28% in sequence `compound_rotation_7`, where only 369 true positives are retrieved.

Despite these concurrent defects, the structural lattice reconstruction layer successfully bridges the broken nodes and resolves shape classification conflicts. By ensuring that every accepted node undergoes zero index mismatches, the framework guarantees perfect pattern registration across all test configurations, enabling unbiased downstream camera matrix recovery.

This resilience highlights a fundamental property of the framework: node shape classification (circles vs. triangles) is strictly required only during the initial pattern localization and coordinate phase-locking stages. Once the binary decoding resolves the global topological indices, the physical shape label becomes secondary.

Because sub-pixel centroid calculations rely on blob intensity moments, the geometric center of a node remains accurate regardless of its shape misclassification. Consequently, the system salvages

these misclassified points as valid anchors for the downstream calibration goal. Furthermore, the verified triangular primitives can be exploited as deterministic sources for sub-pixel anchors located at the triangle corners, drastically scaling up the geometric constraint density under extreme wide-angle distortion for the final non-linear optimization pass.

5.2 Radial Distortion and Intrinsics Estimation Performance

The comparative experimental results presented in Table 4 validate that the proposed cascaded calibration framework completely suppresses parameter cross-talk and structural degeneracies across distinct capture trajectories.

Calibration under pure rotational (roll) configurations traditionally introduces a well-known homographic rank degeneracy. Because the calibration plate remains strictly orthogonal to the optical axis, the projective equations lose perspective translation variance, inducing a strong algebraic correlation between focal scaling parameters (f_x, f_y) and scene depth. In standard joint optimization loops, this ill-conditioned landscape leads to parameter drift or optimization collapse. In contrast, our cascaded architecture completely isolates the lens distortion field κ_1 beforehand using localized non-linear geometric straightness metrics. This spatial decoupling enables the system to converge with an isotropic focal error of just 0.26%, while capturing the lens distortion field within a 4.88% error bound ($\kappa_1 = -0.2098$ solved vs. -0.2000 ground truth).

Conversely, the tilted dataset demonstrates superior performance in capturing lens distortion and successfully recovers anisotropic focal parameters owing to deeper perspective gradients. Unlike classical joint optimization frameworks that suffer from parameter cross-talk and inherently require close initial guesses to avoid local minima, our virtual-aperture regularized approach extracts absolute lens parameters directly from raw line primitives with high precision.

Table 4: Evaluation of Calibration Invariants Across Test Sequences

Intrinsic Parameter	**Simple Rotational Sequence**			**Random Tilt Sequence**		
	GT	**Solved**	**Error, %**	**GT**	**Solved**	**Error, %**
Focal Length f_x (px)	1150.00	1152.97	**0.26**	1250.00	1246.77	**0.26**
Focal Length f_y (px)	1150.00	1152.97	**0.26**	1150.00	1144.15	**0.51**
Principal Point c_x (px)	960.00	961.23	**0.06**	965.00	960.19	**0.25**
Principal Point c_y (px)	540.00	538.45	**0.14**	543.00	540.23	**0.26**
Distortion Coefficient κ_1	-0.2000	-0.2098	**4.88**	-0.1500	-0.1493	**0.45**

Table 5 presents single-frame calibration performance as a function of keypoint recall. The data clearly demonstrates that broad keypoint coverage across the image boundary and perspectivity gradient is essential for reliable intrinsic recovery.

Under single-frame layouts with heavy dropouts (e.g., `compound_rotation_7`), topological graph fragmentation limits pattern registration to its dataset minimum of 369 nodes. The lack of peripheral constraints induces focal-distortion parameter compression. To systematically account for quality degradation, the recall metric R_i is defined as the ratio of the number of correctly decoded pattern nodes TP_i to the total expected nodes N_i within the image bounds of frame i.

The mathematical coupling between estimation error and keypoint recall is explicitly confirmed by the data trends mapped in Table 5. For instance, frames with a high recall, such as `compound_rotation_4` (0.904), consistently preserve sub-percent focal errors (0.8% and 1.1%), whereas a severe drop down to 0.387 in `compound_rotation_7` triggers immediate optimization drift.

Assuming a constant per-keypoint error variance σ_0^2, the estimation variance of the calibration parameters derived from frame i is inversely proportional to the decoded node density $\mathrm{TP}_i = N_i \cdot R_i$:

$$\sigma_i^2 = \frac{\sigma_0^2}{N_i \cdot R_i}. \tag{61}$$

Crucially, a high recall alone does not guarantee intrinsic observability. As demonstrated by `compound_rotation_5`, despite achieving a high recall of 0.916, the estimation suffers from severe parameter divergence (38.5% error in f_x and 50.5% in f_y). This degradation is caused by poor geometric conditioning, which leads to numerical instability and rank deficiency in the focal solver.

Table 5: Random Tilt Sequence: Parameter Confidence Intervals and Multi-Frame Joint Consensus

Test Case	Nodes		f_x (px)		f_y (px)		Distortion κ_1	
	TP	Recall	Solved	Error,%	Solved	Error,%	Solved	Error,%
Ground Truth	**961**	**1.000**	**1250.0**	**Ref.**	**1150.0**	**Ref.**	**-0.1500**	**Ref.**
`compound_rotation_0`	896	0.978	1295.2	3.6	1215.9	5.7	-0.1778	18.5
`compound_rotation_1`	767	0.827	1203.9	3.7	1062.2	7.6	-0.1316	12.3
`compound_rotation_2`	396	0.900	1175.9	5.9	1083.5	5.8	-0.1227	18.2
`compound_rotation_3`	505	0.528	1285.9	2.9	1173.4	2.0	-0.1566	4.4
`compound_rotation_4`	407	0.904	1239.6	0.8	1137.9	1.1	-0.1420	5.3
`compound_rotation_5`	580	0.916	1730.7	38.5	1730.7	50.5	-0.3147	109.8
`compound_rotation_6`	885	0.938	1174.1	6.1	1078.6	6.2	-0.1610	7.3
`compound_rotation_7`	369	0.387	1386.6	10.9	1376.8	19.7	-0.2568	71.2
Weighted Mean **Student** $\Pr = 0.95$, $\nu = 5.52$	—	—	**1237.3** ± **66.3**	**1.0** ± **5.3**	**1141.6** ± **90.3**	**0.7** ± **7.9**	**-0.1565** ± **0.0355**	**4.3** ± **23.7**
Joint Estimation	—	—	**1246.8**	**0.3**	**1144.2**	**0.5**	**-0.1493**	**0.5**

To evaluate baseline single-view expectations, `compound_rotation_5` is strictly classified as an outlier and excluded from the valid sample pool ($N = 7$ remaining views). By weighting the valid single-view estimates proportionally to their keypoint recall ($w_i \propto 1/\sigma_i^2 \propto R_i$), the system implicitly models the posterior likelihood of each camera pose configuration, shielding the interval estimation from depth-translation artifacts.

To address the information loss caused by non-uniform frame reliability, we apply weighted statistics with Kish's effective sample size N_{eff}:

$$N_{\text{eff}} = \frac{\left(\sum_{i=1}^{N} w_i\right)^2}{\sum_{i=1}^{N} w_i^2} = \frac{1}{\sum_{i=1}^{N} \tilde{w}_i^2}, \tag{62}$$

where $\tilde{w}_i = w_i / \sum_{k=1}^{N} w_k$ represents the normalized frame weight such that $\sum \tilde{w}_i = 1$.

The weighted consensus mean for the parameter vector $\boldsymbol{\theta}$ across independent single-frame estimates $\hat{\boldsymbol{\theta}}_i$ is computed as:

$$\hat{\boldsymbol{\theta}} = \sum_{i=1}^{N} \tilde{w}_i \hat{\boldsymbol{\theta}}_i. \tag{63}$$

The unbiased weighted sample variance of the j-th component of the vector $\hat{\boldsymbol{\theta}}$ is defined as:

$$s_{\text{eff},j}^2 = \frac{N_{\text{eff}}}{N_{\text{eff}} - 1} \sum_{i=1}^{N} \tilde{w}_i \left(\hat{\boldsymbol{\theta}}_{i,j} - \hat{\boldsymbol{\theta}}_j\right)^2. \tag{64}$$

Consequently, the $(1-\alpha)$ Student's t confidence interval for the j-th component with $\nu = N_{\text{eff}} - 1$ effective degrees of freedom is constructed as:

$$\text{CI}_{1-\alpha,\, j} = \hat{\boldsymbol{\theta}}_j \pm t_{1-\alpha/2,\, N_{\text{eff}}-1} \cdot \frac{s_{\text{eff},j}}{\sqrt{N_{\text{eff}}}}. \tag{65}$$

Applying this quality-weighted probability profile establishes precise bounds for parameter estimation precision, which is essential for blind calibration without ground truth. Notably, simple parameter post-averaging across independent outputs remains insufficient, as it fails to cancel unobserved single-view structural biases, at least for small view sets.

To overcome this limitation, our proposed joint multi-frame framework bypasses independent post-averaging. By passing the filtered sequence directly into a joint multi-frame bundle consensus layer and aggregating single-frame vanishing points across all valid frames simultaneously, individual view errors cancel out dynamically within the global optimization loop. As a result, this joint geometric minimization completely outperforms statistical post-processing, forcing the overall estimation accuracy of the focal parameters to converge firmly within a tight 0.5% error margin.

6 Conclusion

The Hexagonal Galois Pattern (HGP) framework provides a self-localized target architecture and calibration pipeline designed for high-precision camera parameter recovery under severe optical and geometric degradation. By combining the optimal two-dimensional spatial packing density of hexagonal lattices with the algebraic noise immunity of Galois field codes expanded over two spatial dimensions, the proposed architecture effectively decouples structural grid assembly from both algebraic location decoding and intrinsic parameter extraction.

Experimental validation demonstrates that while single-frame estimation can become ill-conditioned under weak perspective gradients and sparse feature coverage—leading to numerical parameter drift—the joint multi-frame consensus layer successfully stabilizes the solution space. The framework achieves tight focal length error margins ($\approx 0.5\%$), confirming its suitability for high-accuracy metrological applications and camera calibration.

The explicit integration of a full joint optimization loop (e.g., standard Levenberg-Marquardt bundle adjustment) to refine extrinsic poses alongside Perspective-n-Point (PnP) solvers is deliberately omitted from this evaluation. While a non-linear joint solver can theoretically refine parameters near a local minimum, it inherently requires a precise initial intrinsic guess and strict extrinsic initialization via homography factorization [21] to prevent optimization divergence. Because this work focuses strictly on the algebraic and topological restoration of intrinsics, incorporating a subsequent non-linear PnP refinement pass remains outside the current scope and stands as a natural downstream engineering extension.

The core structural and algorithmic contributions of the framework are summarized as follows:

- **Minimal Cropping Area:** The algebraic decoding window operates with a lower structural bound requiring a pattern fragment of at least 12×2 elements, preserving full localization capabilities even when degraded by misclassification errors or feature dropouts.
- **Robust Lattice Construction:** The grid reconstruction is modeled as a physical hexagonal lattice crystallization process propagating from initial seeds across a 6-neighbor adjacency, enabling stable lattice expansion across complex boundaries and partial occlusions.
- **Balanced Pseudorandom Feature Distribution:** Complementing the maximal packing density of the hexagonal grid, an almost ideal uniform binary two-dimensional distribution is mathematically verified. This pseudorandom factorization establishes an algebraic foundation to break constraint coherence and eliminate systematic parameter bias during optimization.
- **Plumb-Line Bundles:** The three-axis lattice generates dense collinear point traces along 30° angular increments, yielding continuous plumb-line bundles that optimize line-based lens distortion solvers and vanishing point estimation.
- **Decoupled Multi-Stage Calibration:** The pipeline enforces a strict separation of calibration stages, isolating non-linear radial lens distortion via straightness invariants prior to analytical closed-form intrinsic parameter extraction. This decoupled architecture demonstrates structural applicability for thin-prism model extensions and explicit lens-to-sensor tilt compensation.

Limitations and Future Work

Despite its resilience, specific operational boundaries define the limits of the HGP pipeline:

- **Parity Saturation Under Error Bursts:** Localized error correction remains strictly bounded by the underlying generator polynomial; dense error bursts exceeding the linear subspace capacity degrade the Gauss-Jordan erasure resolution performance.
- **Complexity Overhead:** Maintaining a Disjoint-Set Union (DSU) of crystal seeds and performing Galois field arithmetic introduces a higher computational complexity compared to generic, unstructured fiducial markers.
- **Centroid Sensitivity at Low SNR:** The current implementation does not account for asymmetric centroid perspective bias, leaving sub-pixel accuracy susceptible to boundary uncertainty under severe motion blur or low signal-to-noise ratios.
- **Grid Reconstruction Vulnerability:** High densities of phantom nodes and clustered erasures under severe noise can distort the hexagonal lattice wave propagation front, requiring further systematic investigation.

Future work will focus on addressing these theoretical constraints. Integrating deep-learning-based sub-pixel node detectors represents a promising direction to improve initial primitive

localization in low-SNR environments, alongside generalizing the algebraic subspace decoder toward higher-order Galois structures to expand the code space volume and burst-error correction boundaries. Furthermore, we intend to explore the mathematical influence of introducing triangular vertex features on geometric calibration precision, and systematically extend this decoupled calibration framework to resolve full, explicit lens-to-sensor tilt parameters in non-coaxial optical systems.

Acknowledgments. The authors express sincere gratitude to the faculty and research staff of the State University of Information and Communication Technologies (Kyiv, Ukraine) for their insightful discussions and supportive academic environment.

Funding. This work was completed without external funding.

Conflict of interest/Competing interests. The authors declare no competing or financial conflicts of interest.

Ethics approval and consent to participate. Not applicable.

Consent for publication. Not applicable.

Use of generative AI. The authors acknowledge the supportive role of a generative AI assistant (Large Language Model) utilized during the drafting phase of this manuscript. The model assisted in refining textual style, formatting LaTeX equations, and implementing selected concept validation scripts in Python. All conceptual breakthroughs, physical pattern topologies, algebraic decoding strategies, and experimental setups were conceived, executed, and validated solely by the authors, who assume full accountability for the integrity and accuracy of the published work.

Data availability. Data sharing is not applicable to this article as no static datasets were generated or analyzed during the current study. All evaluation data are generated dynamically using the open-source Blender scripts available in the GitHub repository https://github.com/gkis-conda/robust_calibration_pattern.

Materials availability. All figures, diagrams, and graphical elements presented in this manuscript were generated directly by the authors and are available from the corresponding author upon reasonable request.

Code availability. The complete Python framework, pattern generation routines, image processing pipelines, and calibration scripts used to process data and generate tables/figures are publicly available in the GitHub repository at https://github.com/gkis-conda/robust_calibration_pattern under the AGPL-3.0 license.

Author contribution. All authors contributed to the study design, software testing, and result processing. Main conceptualization, methodology, and software implementation were performed by Gennadiy Kis.

Appendix A Initial states for reversible dual sequences

A.1 Notation and conventions

All linear algebra is over the field GF(2). We use the LSB-first convention. The n-bit register state at time t is

$$s_t = (s_{t,0}, s_{t,1}, \ldots, s_{t,n-1})^\top \in \mathrm{GF}(2)^n,$$

So, that $s(t) \in \mathrm{GF}(2)^n$ – the state vector "before" the update at time t, $s(t+1) = A\,s(t)$, where A - linear state propagation matrix. Let find it explicitly. A right logical shift `state » 1` corresponds to the linear map $S : \mathrm{GF}(2)^n \to \mathrm{GF}(2)^n$

$$S = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \qquad Ss = (s_1, s_2, \ldots, s_{n-1}, 0)^\top,$$

so S is the $n \times n$ matrix with ones on the superdiagonal: $S_{i,i+1} = 1$ for $0 \le i < n-1$ and all other entries zero.

Let the feedback polynomial be

$$f(x) = a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0, \qquad a_k \in \mathrm{GF}(2),$$

and collect coefficients in the vector $a = (a_0, a_1, \ldots, a_n)^\top$.

Denote by $e_0 = (1, 0, \ldots, 0)^\top$ the LSB selector, where $b_t = s_{t,0} = e_0^\top s_t$ is the sequence output bit produced by the code construct `state & 1`.

For the dual (reciprocal) polynomial the feedback column appears in reversed order; accordingly we may write the dual step matrix as

$$A = S + p\, e_0^\top, \qquad A' = S + p'\, e_0^\top = S + J_n\, p\, e_0^\top,$$

where $e_0 = (1, 0, \ldots, 0)^\top$, $p = (a_1, a_2, \ldots, a_{n-1}, 1)^\top$, $p' = (a_{n-1}, a_{n-2}, \ldots, a_1, 1)^\top$, and J_n is the reversal operator affecting the starting block of length $n-1$ and doesn't change the last element.

Explicit block form of A and A':

$$A = \begin{pmatrix} a_1 & 1 & 0 & \cdots & 0 \\ a_2 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & 0 \\ a_{n-1} & 0 & 0 & \cdots & 1 \\ 1 & 0 & 0 & \cdots & 0 \end{pmatrix}, \qquad A' = \begin{pmatrix} a_{n-1} & 1 & 0 & \cdots & 0 \\ a_{n-2} & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & 0 \\ a_1 & 0 & 0 & \cdots & 1 \\ 1 & 0 & 0 & \cdots & 0 \end{pmatrix}. \tag{A1}$$

One sees directly that the only difference between A and A' is their first column: p vs. p'.

Assume both registers are initialized from an identical state $s(0) = s$. To satisfy the output matching constraint $e_0^\top A^t s = e_0^\top (A')^{-t} s$ during the early evolution stages, the state vector must possess an inherent structural symmetry.

Let the register state be partitioned as $s = (0, u)^\top$, where the active LSB is explicitly zero and the remaining sub-vector $u \in \mathrm{GF}(2)^{n-1}$ is a strict binary palindrome. This specific palindromic configuration serves as a simple sufficient condition for trajectory symmetry due to two direct properties:

- **Feedback Suppression:** Because $s_0 = 0$, the non-zero feedback coefficients in the first column of both the forward matrix A and the dual matrix A' are completely suppressed on the initial step. Consequently, both operators collapse into an identical, pure geometric right-shift: $As = A's$.
- **Symmetric Bit Streaming:** As subsequent clock cycles shift the elements of the palindrome u into the LSB slot, the spatial symmetry of u guarantees that the forward evolution and the time-reversed evolution produces the exact same bit sequence.

Therefore, this palindromic pattern creates a localized time-reversible horizon for exactly $t \le n-1$ steps, forcing the initial output sequences to coincide before the recurrence relation takes over.

A.2 Derive Initial Galois LFSR state

To evaluate backward-in-time trajectories, the exact matrix inverse $(A')^{-1}$ must reverse the right-shift into a left-shift while correctly propagating the feedback pipeline. By executing row-by-column verification over GF(2), the canonical inverse matrix $(A')^{-1}$ is derived as:

$$(A')^{-1} = \begin{pmatrix} 0 & 0 & \ldots & 0 & 1 \\ 1 & 0 & \ldots & 0 & a_{n-1} \\ 0 & 1 & \ldots & 0 & a_{n-2} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & a_1 \end{pmatrix} \tag{A2}$$

Elimination of Matrix Powers via Stepwise Invariant Induction The strict condition for phase-synchronous time reversibility demands that the output of the forward register moving t steps forward must identically match the output of the dual register moving t steps backward from the shared state s. Instead of explicitly evaluating the high-order matrix powers A^t and $(A')^{-t}$, we collapse the trajectory problem into a sequence of polynomial-independent mathematical induction steps.

The First Temporal Boundary ($t = 1$) Evaluating state propagation for a single time step yields:

$$e_0^T A s = (a_1 \cdot s_0) \oplus s_1 \tag{A3}$$

$$e_0^T (A')^{-1} s = s_{n-1} \tag{A4}$$

Equating these outputs forces the algebraic relation $(a_1 \cdot s_0) \oplus s_1 = s_{n-1}$. To eliminate the dependency on the polynomial coefficient a_1 and maintain universal validity across any dual primitive pair ($A \neq A'$), the state is strictly bound by:

$$s_0 = 0 \implies s_1 = s_{n-1}$$

The Inductive Stepwise Invariant ($t \geq 2$) When the initial boundary conditions are propagated forward, the computational need to evaluate matrix powers vanishes entirely. At each subsequent discrete time step t, the internal linear feedback operations containing the polynomial coefficients a_i cancel each other due to the spatial symmetries established in the preceding steps (e.g., $a_1 s_1 \oplus a_1 s_{n-1} = 0$ at $t = 2$).

Consequently, the multi-step operator tracking identity reduces inductively to a single, invariant geometric boundary rule mapping the outer register layers toward the center:

$$\boxed{s_t = s_{n-t}} \quad \text{for } t = 1, 2, \ldots, \frac{n-1}{2} \tag{A5}$$

A.3 Corollary

The time-reversal condition forces the active register state to systematically construct an internal, symmetric $(n-1)$-bit binary palindrome with a forced trailing zero. This restricts the total available degrees of freedom to exactly $\frac{n-1}{2}$, confirming the exact distribution of universal, polynomial-independent phase anchor solutions:

$$N = 2^{\frac{n-1}{2}} - 1 \tag{A6}$$

The total number of non-zero solutions generated by this structural constraint corresponds exactly to the formula (A6):

- **For degree $n = 3$ ($N = 1$ solution):**

$$v_0 = \{6\} \implies [1,1,0]_2 \xrightarrow{\gg 1} [1,1]_2$$

- **For degree $n = 5$ ($N = 3$ solutions):**

$$12 \implies [0,1,1,0,0]_2 \xrightarrow{\gg 1} [0,1,1,0]_2$$
$$18 \implies [1,0,0,1,0]_2 \xrightarrow{\gg 1} [1,0,0,1]_2$$
$$30 \implies [1,1,1,1,0]_2 \xrightarrow{\gg 1} [1,1,1,1]_2$$

- **For degree $n = 7$ ($N = 7$ solutions):**

$$v_0 = \{24, 36, 60, 66, 90, 102, 126\}$$

- **For degree $n = 9$ ($N = 15$ solutions):**

$$v_0 = \{48, 72, 120, 132, 180, 204, 252, 258, 306, 330, 378, 390, 438, 462, 510\}$$

This rule for state construction is applicable to any primitive polynomials for Galois LFSR to have back and forward sequences to be aligned by phases.

Appendix B The Effective Focal Length via Chebyshev Circular Conic Approximation

Expressing the circular aperture coordinates through the canonical normalized camera space transforms the isotropic distortion field into a non-coaxial anisotropic conic (an ellipse) mapped by the full, non-centered Image of the Absolute Conic (IAC) matrix ω. To map this tensor onto standard implementations with minimal structural distortion, we establish the following autonomous optimization framework:

Lemma 1 (Circular Conic Approximation)**.** *Given a non-centered anisotropic intrinsics profile driven by f_x and f_y, the optimal effective focal length f_{eff} is defined by the harmonic mean of the squared focal fields:*

$$\frac{1}{f_{eff}^2} = \frac{1}{2} \cdot \left(\frac{1}{f_x^2} + \frac{1}{f_y^2} \right) \tag{B7}$$

Remarkably, this metric representation corresponds exactly to the classical Chebyshev circular approximation of an ellipse evaluated within an integral L_2 matrix manifold.

Proof of Lemma 1 (Circular Conic Approximation).

The radial lens unwarping factor $\mathcal{D}$ represents an intrinsic, physical multiplier acting on a given spatial track point and its numerical value must remain invariant under any choice of homogeneous coordinate representation or scaling domain. Expressing this scaling invariance across both the Aperture-Normalized description and the canonical Brown-Conrady focal space dictates that the respective localized quadratic expansion fields should be mapped with minimal geometric discrepancy across the active sensor plane:

$$\kappa_{\text{ap}} \cdot r_{\text{ap}}^2 \approx \kappa_1 \cdot r_n^2 \implies \kappa_1 \approx \kappa_{\text{ap}} \left(\frac{r_{\text{ap}}}{r_n} \right)^2 . \tag{B8}$$

Because the true focal canvas exhibits directional anisotropy ($f_x \neq f_y$), the scale relation r_{ap}/r_n varies continuously with the polar angle across the image plane, making perfect analytical identity impossible. To preserve strict calibration integrity without introducing directional scale distortion, the pipeline must optimize the parameters globally to minimize the integrated discrepancy between these two quadratic radial representations.

Let $\mathbf{x} = [x_{\text{pixel}}, y_{\text{pixel}}, 1]^\top \in \mathbb{P}^2$ define the raw, non-centered homogeneous pixel coordinate vector captured on the sensor matrix. To isolate the quadratic radial metrics from the projective dimension bias, we project the IAC $\omega = \mathbf{K}^{-\top}\mathbf{K}^{-1}$ onto the 2D focal sub-manifold utilizing the ideal truncation operator $\mathbf{J} = \text{diag}(1, 1, 0)$ via the symmetric sandwich product: $\omega_2 = \mathbf{J}\omega\mathbf{J}$. Under this formulation, the resolution-bounded radial fields are expressed directly and exactly via the truncated metric tensors without unit offsets:

$$r_n^2 = \mathbf{x}^\top \mathbf{\Omega}_2 \mathbf{x}, \qquad r_{\text{ap}}^2 = \mathbf{x}^\top \mathbf{\Omega}_2^{\text{ap}} \mathbf{x}, \tag{B9}$$

where $\mathbf{\Omega}_2 = \mathbf{T}_c^{-\top} \text{diag}(1/f_x^2, 1/f_y^2, 0)\mathbf{T}_c^{-1}$ and $\mathbf{\Omega}_2^{\text{ap}} = \mathbf{T}_c^{-\top} \text{diag}(1/R_{\text{ap}}^2, 1/R_{\text{ap}}^2, 0)\mathbf{T}_c^{-1}$ define the decoupled two-dimensional embedding profiles of the IAC matrix [21] and the target aperture matrix respectively.

By leveraging the scale-invariant conversion rule $\kappa_1 = \kappa_{\text{ap}} \cdot (f_{\text{eff}}/R_{\text{ap}})^2$, the optimization problem over the bounded coordinate canvas translates to finding an optimal isotropic circular configuration that minimizes the discrepancy between these two quadratic representations:

$$\begin{aligned} \arg\min_{f_{\text{eff}}} \left(\max_{\mathbf{x}} \left| \kappa_{\text{ap}} \cdot r_{\text{ap}}^2 - \kappa_1 \cdot r_n^2 \right| \right) &= \arg\min_{f_{\text{eff}}} \left(\max_{\mathbf{x}} \left| \kappa_{\text{ap}} \left(\mathbf{x}^\top \mathbf{\Omega}_2^{\text{ap}} \mathbf{x} \right) - \kappa_1 \left(\mathbf{x}^\top \mathbf{\Omega}_2 \mathbf{x} \right) \right| \right) \\ &= \arg\min_{f_{\text{eff}}} \left(\max_{\mathbf{x}} \left| \mathbf{x}^\top \left(\kappa_{\text{ap}} \mathbf{\Omega}_2^{\text{ap}} - \kappa_{\text{ap}} \cdot \left(\frac{f_{\text{eff}}}{R_{\text{ap}}} \right)^2 \mathbf{\Omega}_2 \right) \mathbf{x} \right| \right) . \end{aligned} \tag{B10}$$

Because the common translation boundary matrix $\mathbf{T}_c$ acts as an isometric coordinate shift, the functional minimization over the tensor components maps directly onto the diagonal focal components. Factorizing out the invariant scale κ_{ap} and evaluating the system under the Frobenius matrix norm, the rescaled objective collapses directly onto the upper-left 2×2 focal manifold:

$$\arg\min_{f_{\text{eff}}} \left\| \left(\frac{f_{\text{eff}}}{R_{\text{ap}}}\right)^2 \text{diag}\left(\frac{1}{f_x^2}, \frac{1}{f_y^2}\right) - \text{diag}\left(\frac{1}{R_{\text{ap}}^2}, \frac{1}{R_{\text{ap}}^2}\right) \right\|_F^2 = \arg\min_{f_{\text{eff}}} \left\| \text{diag}\left(\frac{1}{f_x^2}, \frac{1}{f_y^2}\right) - \frac{1}{f_{\text{eff}}^2}\mathbf{I}_2 \right\|_F^2. \tag{B11}$$

Crucially, because the minimization space is governed entirely by the focal parameters, we can project this core relationship back onto the full projective 3×3 representation. By appending the homogeneous coordinate constants to the diagonal blocks without shifting the optimization landscape, Eq. (B11) confirms that the physical unwarping optimization corresponds exactly to finding the closest circular conic on the full IAC matrix manifold:

$$E\left(\frac{1}{f_{\text{eff}}^2}\right) = \left\| \text{diag}\left(\frac{1}{f_x^2}, \frac{1}{f_y^2}, 1\right) - \text{diag}\left(\frac{1}{f_{\text{eff}}^2}, \frac{1}{f_{\text{eff}}^2}, 1\right) \right\|_F^2 = \left(\frac{1}{f_x^2} - \frac{1}{f_{\text{eff}}^2}\right)^2 + \left(\frac{1}{f_y^2} - \frac{1}{f_{\text{eff}}^2}\right)^2 + (1-1)^2. \tag{B12}$$

Taking the first partial derivative with respect to the scalar variable $(1/f_{\text{eff}}^2)$, we isolate the algebraic extrema condition:

$$\frac{\partial E}{\partial(1/f_{\text{eff}}^2)} = -2\left(\frac{1}{f_x^2} - \frac{1}{f_{\text{eff}}^2}\right) - 2\left(\frac{1}{f_y^2} - \frac{1}{f_{\text{eff}}^2}\right) = 0. \tag{B13}$$

Evaluating the linear roots yields the definitive energy-preserving relation:

$$\frac{1}{f_{\text{eff}}^2} = \frac{1}{2} \cdot \left(\frac{1}{f_x^2} + \frac{1}{f_y^2}\right). \tag{B14}$$

This concludes the proof. By minimizing the Frobenius norm on the IAC manifold in Eq. (B11), the circular conic generated by f_{eff} is not merely a geometric focal approximation, but directly serves as the L_2-optimal circular conic approximation for the physical lens distortion field $\mathcal{D}(\mathbf{x})$.

References

[1] Malhotra, D., El-Sheimy, N., Gao, Y.: Extrinsic calibration of visual and inertial sensors for the autonomous vehicle. In: Proceedings of the International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2023), Denver, Colorado, USA, pp. 2131–2139 (2023). https://doi.org/10.33012/2023.19323

[2] Zhang, J., Wang, J., Wu, Q., Wang, S., Jiang, J.: Camera-lidar calibration methods in autonomous driving: A review. IEEE Transactions on Intelligent Transportation Systems **24**(12), 13110–13129 (2023) https://doi.org/10.1109/TITS.2023.3309833

[3] Qian, L., Azimi, E., Kazanzides, P., Nassir, N.: A review of calibration methods for optical see-through augmented reality displays. IEEE Transactions on Visualization and Computer Graphics **27**(4), 2484–2502 (2021) https://doi.org/10.1109/TVCG.2021.3054178

[4] Bradski, G.: The OpenCV Library. Dr. Dobb's Journal of Software Tools **25**(11), 120–125 (2000)

[5] Xu, J., Peetakul, J., Li, M., Zhou, J.: Geometric Wide-Angle Camera Calibration: A Review and Evaluation. arXiv preprint arXiv:2306.09014 (2023). Revised March 2024

[6] Olson, E.: Apriltag: A robust and flexible visual fiducial system. In: 2011 IEEE International Conference on Robotics and Automation, pp. 3400–3407 (2011). https://doi.org/10.1109/ICRA.2011.5979561

[7] Garrido-Jurado, S., Muñoz-Salinas, R., Madrid-Cuevas, F.J., Marín-Jiménez, M.J.: Automatic generation and detection of highly reliable fiducial markers under occlusion. Pattern Recognition **47**(6), 2280–2292 (2014) https://doi.org/10.1016/j.patcog.2014.01.005

[8] Lightbody, P., Krajník, T., Hanheide, M.: A versatile high-performance visual fiducial marker detection system with scalable identity encoding. In: Proceedings of the 32nd ACM Symposium

on Applied Computing. SAC '17, pp. 276–282. Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/3019612.3019709

[9] Romero-Ramirez, F.J., Muñoz-Salinas, R., Medina-Carnicer, R.: Speeded up detection of ArUco fiducial markers. Image and Vision Computing **76**, 38–47 (2018) https://doi.org/10.1016/j.imavis.2018.05.004

[10] Mallon, J., Whelan, P.F.: Precise calibration of circular features on a planar target. Machine Vision and Applications **16**(3), 132–141 (2005) https://doi.org/10.1007/s00138-004-0164-3

[11] Heikkilä, J.: Geometric Camera Calibration Using Circular Control Points. IEEE Transactions on Pattern Analysis and Machine Intelligence **22**(10), 1066–1077 (2000) https://doi.org/10.1109/34.879788

[12] Datta, A., Kim, R., Kanade, T.: Accurate camera calibration using iterative refinement of control points. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, pp. 120–119 (2009). https://doi.org/10.1109/CVPRW.2009.5204179 . IEEE

[13] Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. IEEE Transactions on Pattern Analysis and Machine Intelligence **27**(10), 1615–1630 (2005)

[14] Hughes, C., Denny, P., Jones, E., Glavin, M.: Wide-angle camera calibration using laser-projected circular patterns. Applied Optics **49**(10), 1876–1883 (2010) https://doi.org/10.1364/AO.49.001876

[15] Schwalbe, E.: A new approach for the calibration of digital fisheye cameras in forest canopy photography. International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences **36**(5), 417–422 (2005)

[16] Scaramuzza, D., Martinelli, A., Siegwart, R.: A flexible technique for accurate omnidirectional camera calibration and structure from motion. In: Proceedings of the IEEE International Conference on Computer Vision Systems (ICVS), pp. 45–45 (2006). https://doi.org/10.1109/ICVS.2006.3 . IEEE

[17] Kannala, J., Brandt, S.S.: A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses. IEEE Transactions on Pattern Analysis and Machine Intelligence **28**(8), 1335–1340 (2006) https://doi.org/10.1109/TPAMI.2006.153

[18] Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**(2), 91–110 (2004) https://doi.org/10.1023/B:VISI.0000029664.99615.94

[19] Middleton, L., Sivaswamy, J.: Hexagonal Image Processing: A Practical Approach. Springer, London, UK (2005). https://doi.org/10.1007/1-84628.046.0

[20] Stoyan, D., Kendall, W.S., Mecke, J.: Stochastic Geometry and Its Applications, 2nd edn. John Wiley & Sons, Chichester, UK (1995)

[21] Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision, 2nd edn. Cambridge University Press, Cambridge, UK (2003)

[22] Klette, R., Rosenfeld, A.: Digital Geometry: Geometric Foundations for Digital Picture Analysis. Elsevier Morgan Kaufmann, San Francisco, CA, USA (2004)

[23] Proakis, J.G., Salehi, M.: Digital Communications, 5th edn. McGraw-Hill, ??? (2008)

[24] Zhang, Z.: A flexible new technique for camera calibration. IEEE Transactions on Pattern Analysis and Machine Intelligence **22**(11), 1330–1334 (2000) https://doi.org/10.1109/34.888718

[25] Brown, D.C.: Close-range camera calibration. Photogrammetric Engineering **37**(8), 855–866 (1971)

[26] Shortis, M.R., Robson, S., Beyer, H.A.: Isolated determination of calibration parameters for close-range photogrammetric cameras. The Photogrammetric Record **15**(85), 37–59 (1995) https://doi.org/10.1111/j.1477-9730.1995.tb00350.x

[27] Menger, K.: Untersuchungen über allgemeine metrik. Mathematische Annalen **103**(1), 466–501 (1930) https://doi.org/10.1007/BF01455704

[28] Nelder, J.A., Mead, R.: A simplex method for function minimization. The Computer Journal **7**(4), 308–313 (1965) https://doi.org/10.1093/comjnl/7.4.308