Quantification of Visually Imperceptible Thin-Film Deposition Using Physics-Anchored Drift-Based State Conformance Metrics

Author: Stephen Francis
Affiliation: Phocoustic, Inc.
Date: March 1, 2026

Abstract

Thin films deposited on textured surfaces often produce distributed spectral redistribution without generating visually discernible boundaries or geometric discontinuities. Conventional vision-based inspection systems, which rely on edge contrast or predefined defect morphology, may struggle to detect such low-contrast, non-object-level perturbations.

This study presents a physics-anchored drift framework for quantifying visually imperceptible thin-film deposition relative to a defined reference state. A controlled experiment was conducted using a matte polymer substrate under fixed darkfield illumination with all adaptive camera functions disabled. A baseline (Golden) image and a detect image containing a thin alcohol film were captured under identical acquisition parameters.

Although the two regions were visually indistinguishable—even under brightness-enhanced inspection—quantitative drift metrics revealed measurable distributed deviation. Baseline metrics remained at zero (drift_mean = 0), while the thin-film condition exhibited elevated distributed activation (drift_mean = 58.79; padr_dist_score = 59.02) without object-level structural emergence (strt_Lcc = 0.0103).

These results demonstrate reference-anchored detection of distributed conformance loss without reliance on machine learning or probabilistic inference. The findings support drift-based State Conformance measurement as a viable framework for detecting sub-perceptual surface perturbations under controlled acquisition conditions.

Table of Contents

Quantification of Visually Imperceptible Thin-Film Deposition Using Physics-Anchored Drift-Based State Conformance Metrics.....1

Abstract....1
1. Introduction....2
1.1 The Thin-Film Detection Problem....2
1.2 Objective....3
2. Experimental Setup....3
2.5 ROI Consistency and Spatial Alignment....5
2.6 Drift Computation Integrity....5
3. Region of Interest (ROI) Definition....6
3.1 Full-Frame Context....7
3.2 ROI Cropping....10
3.3 Brightness-Enhanced Visualization....11
4. Drift-Based Quantification Framework....12
4.1 Physics-Anchored Drift Extraction....12
4.2 Spatial Reference Tiling (STRT)....13
4.3 Distributed Drift Quantification (PADR)....13
5. Quantitative Results....13
5.1 Golden Baseline....13
5.2 Thin-Film Detect Metrics....14
5.3 Interpretation of Results....14
6. State Conformance Interpretation....15
7. Measurement Integrity and Deterministic Conformance....15
8. Discussion....16
8.1 Distributed Field Instability....16
8.2 Comparison to Conventional Vision....16
8.3 Deterministic Operation....17
9. Industrial Implications....17
10. Limitations and Future Work....18
11. Conclusion....18
Appendix A — Definitions of Terms....19
Appendix B — Symbol and Metric Definitions....20
Appendix C — Imaging Parameter Lockdown Specification....21
Appendix D — ROI Selection Criteria....22
Appendix E — Interpretation Matrix for Distributed vs Structural Perturbation....22
Appendix F — Deterministic Conformance Principle....23
Appendix G — Limitations of Visual Inspection....23
Appendix H — Quantization and Signal Integrity....23

1. Introduction

1.1 The Thin-Film Detection Problem

Thin films present a unique challenge in surface inspection. Unlike cracks, scratches, or geometric defects, thin films:

- Do not necessarily form visible boundaries.
- Do not produce strong edge gradients.
- Often manifest as distributed micro-scattering changes.
- May remain below human perceptual thresholds.

Traditional computer vision systems rely heavily on:

- Edge detection.
- Intensity thresholding.
- Machine learning classification trained on labeled defect examples.

Such approaches may fail when perturbations are distributed rather than localized.

1.2 Objective

The objective of this study is to evaluate whether a deterministic, physics-anchored drift framework can:

1. Detect thin-film deposition that is visually indistinguishable from baseline.
2. Quantify deviation relative to a defined physical reference.
3. Distinguish distributed haze from object-level defect emergence.
4. Operate without reliance on machine learning training.

---

2. Experimental Setup

2.1 Imaging System and Optical Configuration

All image acquisition was performed using a fixed camera geometry under controlled darkfield illumination. The imaging configuration was optimized for quantitative drift measurement rather than aesthetic visualization.

The camera position, working distance, focus, and illumination geometry were mechanically fixed and not altered between baseline (Golden) and detect (thin-film) captures. The substrate remained stationary throughout the experiment to preserve pixel-level spatial correspondence.

Images were acquired in native sensor output format without post-capture normalization, histogram equalization, or adaptive contrast adjustment prior to drift computation. Brightness-enhanced figures included in this paper are presentation-only renderings and were not used in any quantitative analysis.

2.2 Manual Parameter Locking and Automation Disablement

To ensure deterministic measurement integrity, all adaptive camera functions within the ICentral acquisition software were explicitly disabled prior to image capture. The following automatic adjustments were disabled:

- Auto exposure
- Auto gain
- Auto white balance
- Auto contrast normalization
- Auto gamma correction
- Any adaptive sharpening or dynamic enhancement functions

All acquisition parameters were manually configured and held constant for both Golden and Detect captures.

Brightness and raw gain were intentionally set conservatively to preserve dynamic range and avoid saturation. This configuration may produce images that appear visually dark; however, it ensures that pixel intensity values remain within a linear operating regime and that measured variation reflects physical surface change rather than camera adaptation.

No parameter was altered between baseline and detect conditions. The presence of the thin film was the only experimental variable.

2.3 Signal Representation and Drift Domain

Drift computation was performed directly on raw pixel intensity values captured under fixed acquisition parameters.

Let:

$I_{ref}(x, y)$ represent baseline intensity
$I_{detect}(x, y)$ represent detect intensity

Drift magnitude is defined as:

$$D(x, y) = |I_{detect}(x, y) - I_{ref}(x, y)|$$

All drift metrics (drift_mean, drift_p95, drift_max, padr_dist_score, strt_S, strt_Lcc) were computed exclusively from these raw intensity differences.

No gamma correction, tone mapping, or nonlinear scaling was applied prior to computation. Visualization scaling shown in figures was applied only for presentation clarity and was not used in analysis.

2.4 Sensor Linearity Assumption

Under the selected exposure and gain parameters, the imaging sensor is assumed to operate within a linear response regime. Conservative brightness and gain settings were chosen to avoid:

- Saturation
- Clipping
- Nonlinear amplification
- Auto-compensation artifacts

Because all adaptive features were disabled and illumination was fixed, pixel intensity variation is assumed to be linearly proportional to reflectance variation within the ROI.

Therefore, measured drift values reflect proportional physical changes in surface micro-scattering response rather than camera-induced normalization.

2.5 ROI Consistency and Spatial Alignment

An identical pixel-coordinate ROI was extracted from both Golden and Detect frames. The ROI was defined prior to quantitative analysis and applied symmetrically without modification.

This eliminates:

- Cropping asymmetry
- Spatial selection bias
- Algorithmic region scanning effects

All drift metrics were computed exclusively within this fixed ROI.

2.6 Drift Computation Integrity

Drift-based metrics were computed directly from raw image captures. No brightness enhancement, visualization scaling, or image adjustment was applied prior to metric computation.

Brightness-enhanced images included in this paper are presentation-only renderings and were not used in any quantitative analysis.

This ensures that reported metrics reflect physically captured optical response rather than post-processing artifacts.

---

3. Region of Interest (ROI) Definition

Accurate quantification of thin-film deposition requires controlled spatial comparison between a defined reference state and a detect state. In distributed perturbation scenarios—such as thin-film haze—full-frame analysis may dilute subtle deviations by incorporating unaffected regions. For this reason, a fixed Region of Interest (ROI) was defined and analyzed identically across both Golden (baseline) and Detect (thin-film) captures.

The ROI serves as the bounded spatial domain within which deterministic drift quantification is performed. It was manually defined prior to analysis and applied symmetrically to both frames without modification. The selected region satisfies the following constraints:

1. Identical pixel coordinates in both reference and detect images
2. Fixed optical and geometric acquisition conditions
3. Representative substrate texture
4. Absence of boundary artifacts or illumination gradients

This approach eliminates cropping asymmetry, frame misalignment, and algorithmic region-selection bias. Any measured deviation therefore arises exclusively from physical surface change rather than acquisition variability.

ROI control is particularly critical in thin-film detection because the perturbation is spatially distributed rather than localized. A thin alcohol film does not form a visible boundary or high-gradient structure; instead, it modifies micro-scattering characteristics across the affected area. Constraining analysis to a fixed ROI enables:

- Spatially localized drift aggregation
- Tile-based topology evaluation (STRT)
- Distributed activation quantification (PADR)
- Structural coherence assessment

By isolating a reproducible spatial domain, the experiment transitions from general image comparison to controlled spatial metrology. Quantitative metrics such as drift_mean, padr_dist_score, and strt_Lcc therefore reflect local conformance deviation rather than global scene variation.

3.1 Full-Frame Context

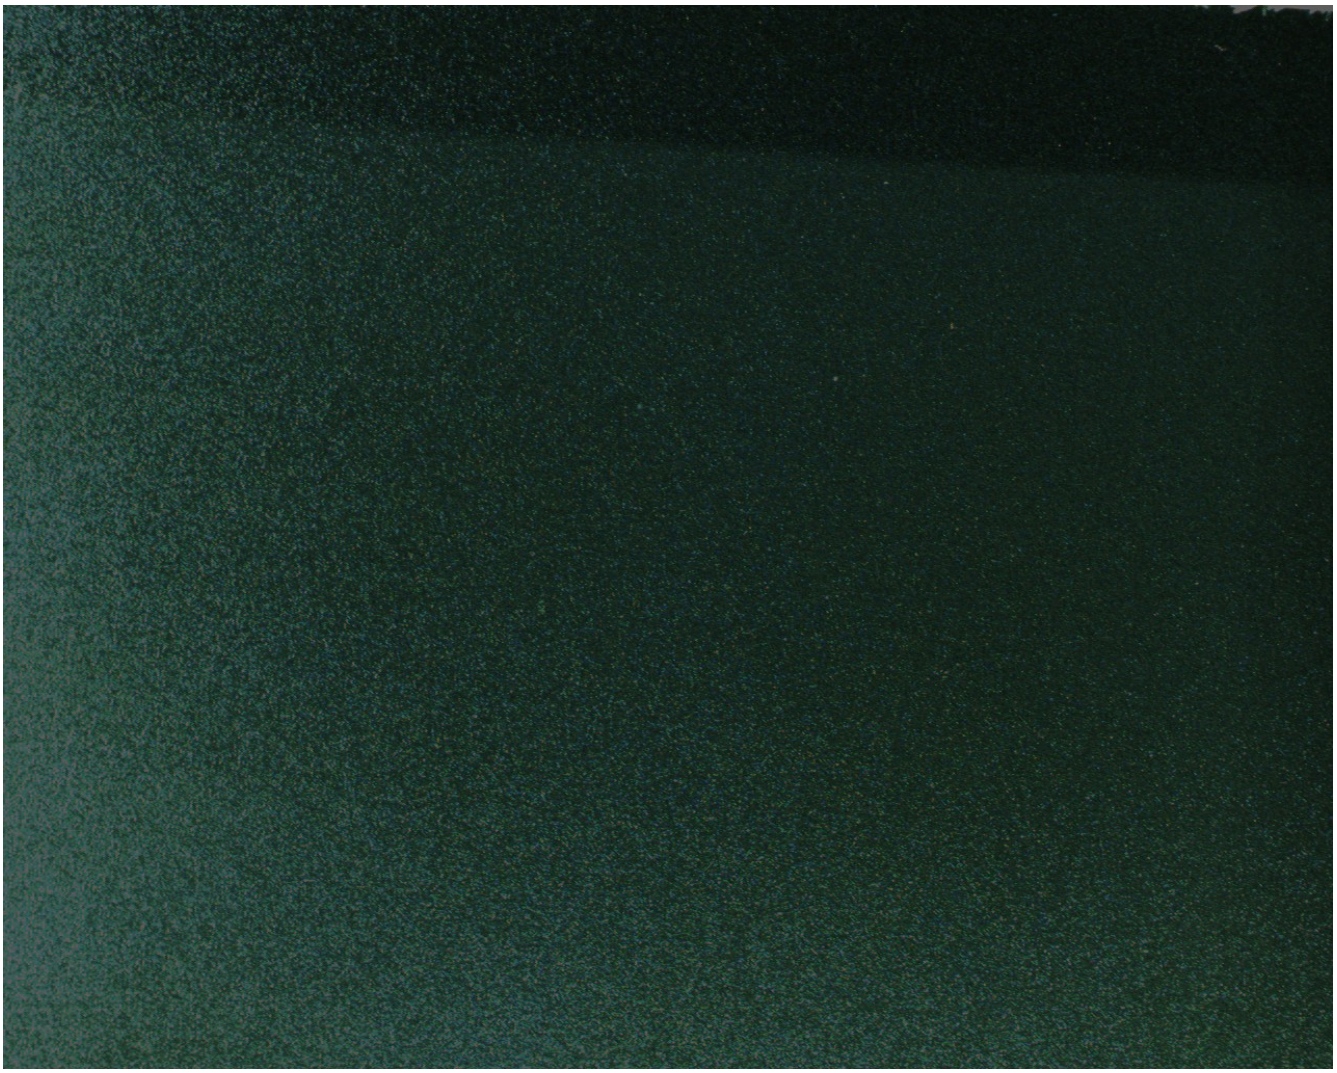

Figure 1 — Full Frame Golden (Baseline)

Full-frame Golden (baseline) capture under controlled darkfield illumination. This image defines the expected physical scattering state of the textured substrate prior to thin-film deposition. Drift metrics for this frame were zeroed, confirming baseline conformance.

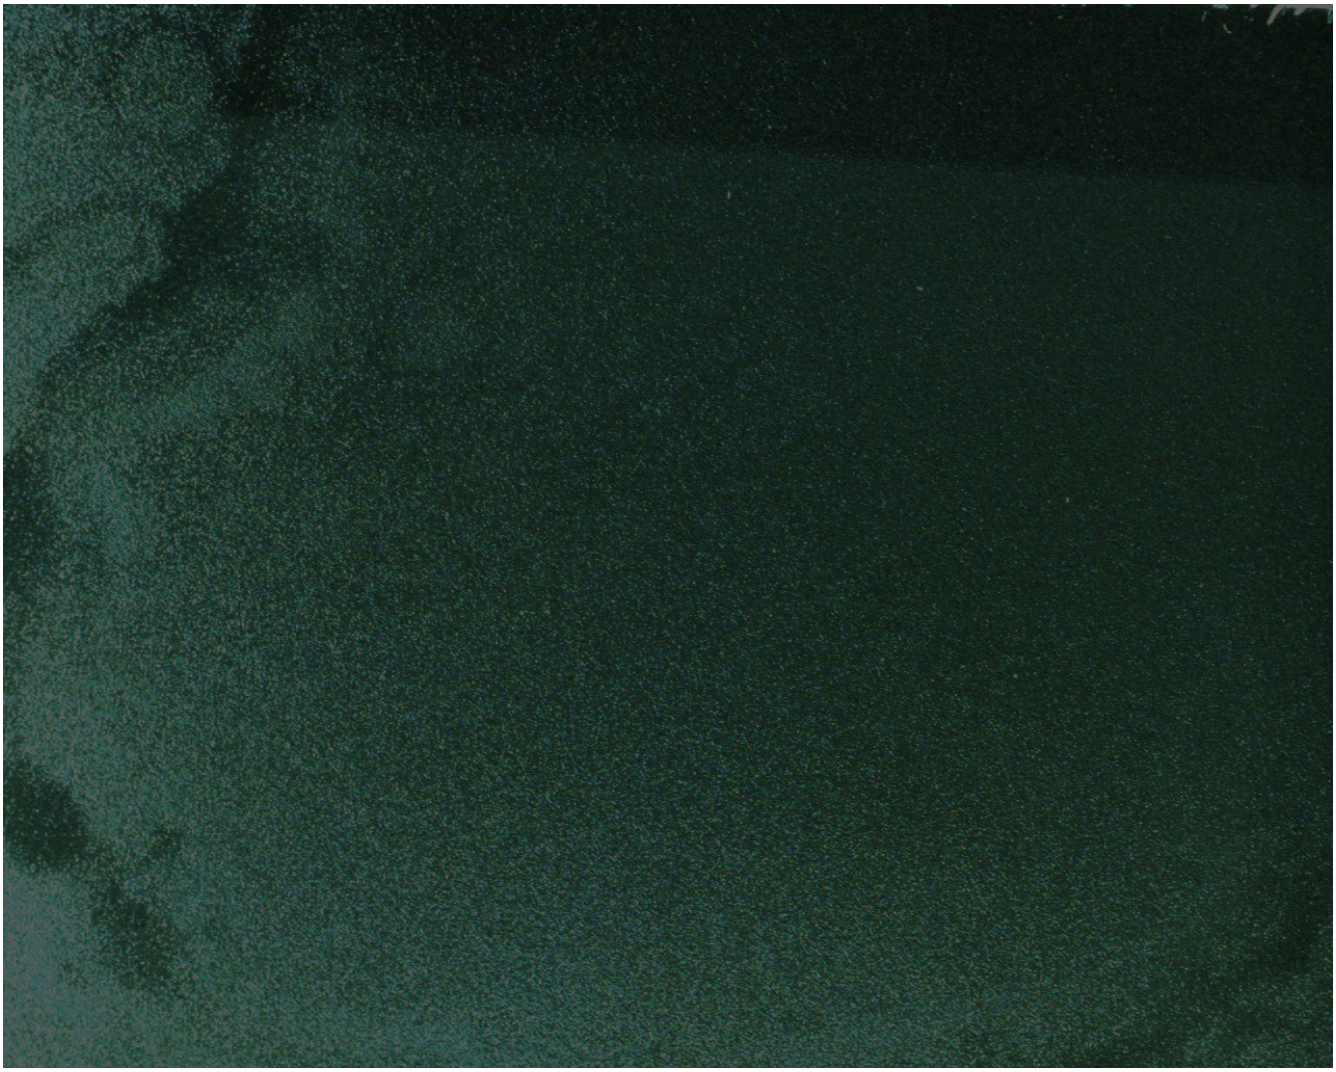

Figure 2 — Full Frame Detect (Thin Film)

Full-frame detect image captured after application of a thin alcohol film. Imaging geometry, illumination, and exposure were identical to the Golden capture. No macroscopic boundary, haze gradient, or structural discontinuity is visually discernible. Alcohol film alters refractive index boundary conditions, modifying angular scattering response under darkfield.

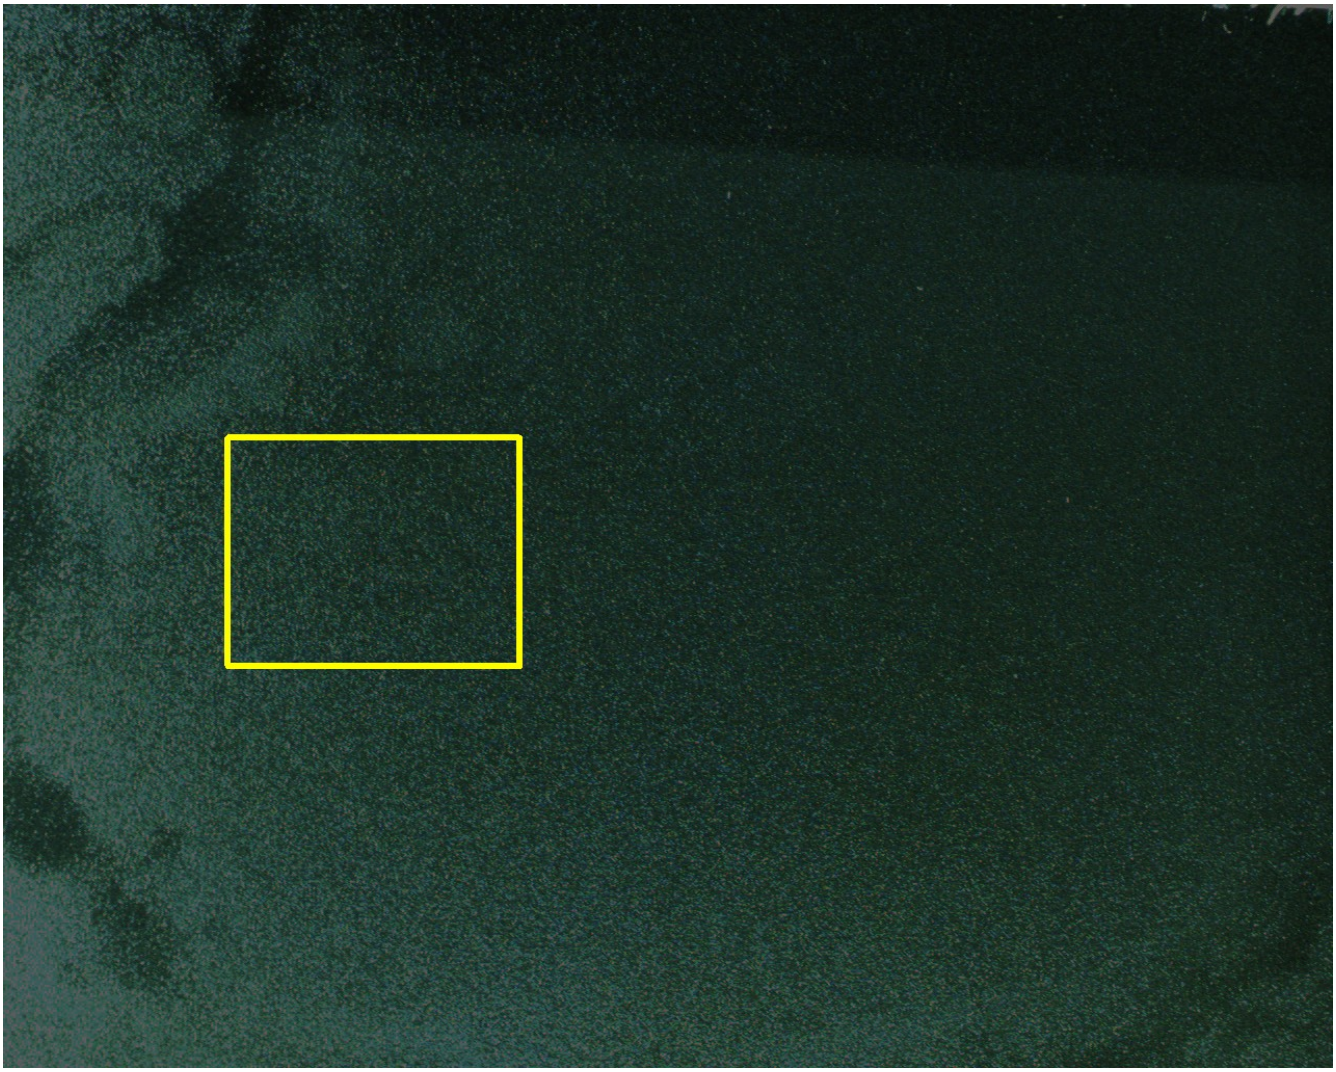

Figure 3 — Detect Frame with ROI Overlay

Detect frame with fixed Region of Interest (ROI) overlay (approximate dimensions). The ROI was defined prior to quantitative analysis and applied identically to both Golden and Detect frames to eliminate spatial selection bias and ensure deterministic comparison.

3.2 ROI Cropping

An identical ROI was extracted from both Golden and Detect frames.

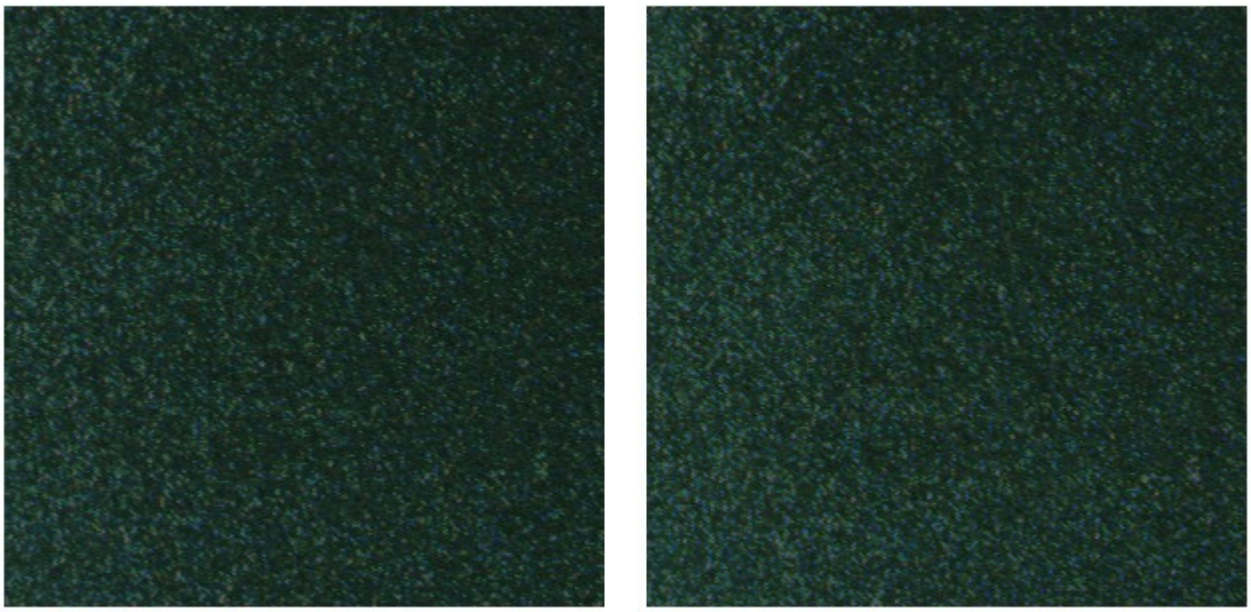

Figure 4 — Cropped Golden vs Detect (Raw)

3.3 Brightness-Enhanced Visualization

For conceptual comparison only, brightness enhancement was applied in Inkscape. Drift analysis was performed exclusively on raw images.

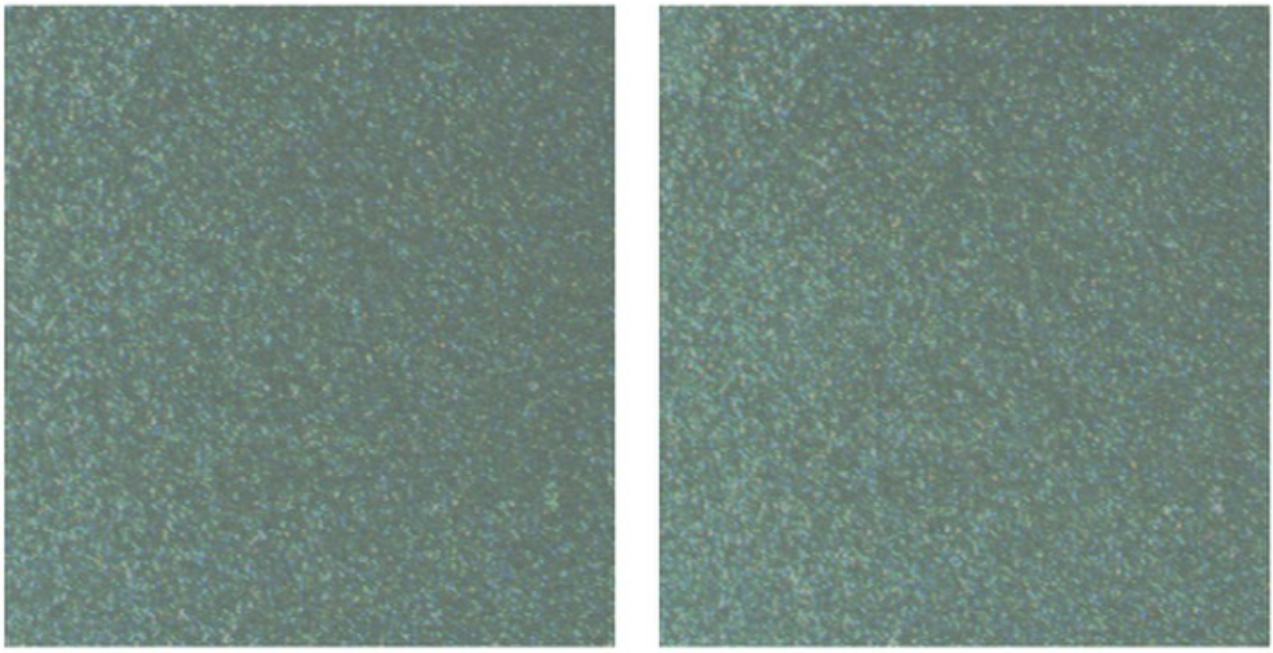

Figure 5 — Brightness-Enhanced Comparison (Presentation Only)

Brightness-enhanced comparison of cropped ROI for conceptual visualization only. Enhancement was performed exclusively for presentation clarity. All drift computations were performed on raw, unaltered images. Even under enhanced inspection, no visible thin-film boundary or haze gradient is apparent.

Observation:

Even after brightness enhancement, the Golden and Detect images remain visually indistinguishable.

There is:

- No visible haze gradient.
- No ring boundary.
- No structural discontinuity.

Conclusion: Under the defined acquisition and enhancement conditions, no visually discernible boundary or gradient is apparent.

---

4. Drift-Based Quantification Framework

4.1 Physics-Anchored Drift Extraction

The drift framework operates by establishing a fixed physical reference state and measuring deterministic deviation from that state under parameter-locked acquisition conditions.

The process consists of:

1. Defining a baseline reference state (Golden capture).
2. Measuring local deviation relative to that state.
3. Aggregating deviation into spatially structured metrics.

Let:

I_ref(x, y) represent baseline intensity at pixel (x, y)
I_detect(x, y) represent detect intensity

Drift magnitude is defined as:

$D(x, y) = |I_{detect}(x, y) - I_{ref}(x, y)|$

Because all acquisition parameters (exposure, gain, illumination geometry, and adaptive functions) are fixed, the difference term D(x, y) reflects physical change in surface scattering response rather than camera-induced normalization or contrast adaptation.

This reference-anchored formulation differs fundamentally from conventional contrast-based inspection. The system does not attempt to enhance edges, segment objects, or classify learned defect patterns. Instead, it evaluates deviation directly relative to a defined physical baseline. The absence of gamma correction, histogram equalization, or nonlinear scaling ensures that drift magnitude remains proportional to measured reflectance change within the sensor’s linear operating regime.

Under these constraints, drift extraction functions as a deterministic measurement operator rather than a probabilistic inference engine. The output is a spatial field of deviation values that can be aggregated into interpretable metrology-style metrics (drift_mean, drift_p95, padr_dist_score, strt_S, strt_Lcc).

No machine learning models, training sets, or learned parameters are used in this determination. Deviation is measured directly against a stable, parameter-locked reference state.

4.2 Spatial Reference Tiling (STRT)

STRT partitions the ROI into structured tiles and evaluates local deviation.

Key topology metrics:

- strt_S: spatial activation fraction
- strt_Lcc: largest connected component ratio

4.3 Distributed Drift Quantification (PADR)

The padr_dist_score metric quantifies distributed deviation across the ROI.

Thin films are expected to produce:

- Elevated distributed drift
- Low connected topology
- Minimal object coherence

---

5. Quantitative Results

5.1 Golden Baseline

As previously shown, all Golden (baseline) drift metrics were zero:

- **Mean Surface Drift (drift_mean)** = 0
- **Distributed Drift Score (padr_dist_score)** = 0
- **Activated Spatial Fraction (strt_S)** = 0
- **Structural Coherence Ratio (strt_Lcc)** = 0

These results indicate that no measurable deviation from the defined reference state was present within the ROI.

The absence of distributed activation, spatial clustering, or structural emergence confirms that the baseline capture represents a stable conformance condition under the defined imaging geometry and illumination parameters.

5.2 Thin-Film Detect Metrics

Internal Metric	Descriptive Name	Value	Interpretation
drift_mean	Mean Drift Magnitude	58.79	Average deviation from baseline across ROI
drift_p95	95th Percentile Drift	127	Upper-bound distributed deviation level
drift_max	Maximum Drift	255	Peak local deviation intensity
padr_dist_score	Distributed Drift Score	59.02	Strength of spatially distributed deviation
strt_S	Spatial Activation Fraction	0.014	Fraction of ROI tiles exceeding drift threshold
strt_Lcc	Largest Connected Component Ratio	0.0103	Degree of object-like structural emergence

The thin-film condition produces a substantial increase in distributed drift (mean = 58.79; distributed score = 59.02) while maintaining extremely low connected topology (Lcc = 0.0103), confirming diffuse spectral redistribution without object-level structural emergence.

5.3 Interpretation of Results

Drift Magnitude

Mean drift increased from 0 to 58.79, indicating measurable deviation from baseline.

The 95th percentile drift of 127 confirms that elevated deviation is not confined to isolated pixels.

Relative to baseline, mean drift increased by 58.79 units within the ROI, while structural coherence remained near zero.

The separation between baseline and detect conditions exceeds the quantization floor by more than two orders of magnitude.

Distributed Activation

padr_dist_score ≈ 59.02 confirms distributed deviation across the ROI.

This is consistent with thin-film spectral redistribution rather than localized defect formation.

Topological Signature

Low strt_Lcc (0.0103) indicates:

- No dominant connected structure.
- No object-level boundary formation.

- No crack-like topology.

The thin film produces distributed field instability rather than geometric anomaly.

While this study demonstrates deterministic separation between baseline and thin-film conditions, future work will include repeated trials to quantify repeatability and variance bounds.

---

6. State Conformance Interpretation

This system does not classify anomalies.

Instead, it:

1. Establishes a defined reference state.
2. Measures deterministic deviation from that state.
3. Quantifies conformance loss.

Golden state:
Conformance confirmed.

Thin film:
Conformance degraded via distributed drift increase.

No object-level emergence threshold is crossed.

---

7. Measurement Integrity and Deterministic Conformance

The validity of deterministic State Conformance measurement depends critically on acquisition stability and parameter control. In this experiment, all automatic camera adjustments were disabled, and all imaging parameters were manually fixed across baseline and detect captures. This eliminates adaptive normalization effects that could otherwise mask or exaggerate surface variation.

The use of conservative brightness and gain settings preserves dynamic range and maintains linear signal response. By preventing camera-induced compensation for reflectance change, the system ensures that measured drift reflects physical surface state change alone.

This strict acquisition discipline aligns with metrology principles and reinforces the deterministic nature of the State Conformance framework. Conformance degradation is not inferred probabilistically; it is measured directly against a stable, parameter-locked reference state.

8. Discussion

8.1 Distributed Field Instability

Thin films modify surface micro-scattering characteristics. The resulting signature:

- Elevated distributed drift
- Low structural coherence
- No localized boundary

This aligns with theoretical expectations of distributed spectral redistribution.

8.2 Comparison to Conventional Vision

Conventional vision-based inspection systems are typically optimized for detecting geometric discontinuities, high-contrast edges, or predefined defect classes. These systems commonly rely on:

- Gradient-based edge detection
- Intensity threshold segmentation
- Feature extraction followed by trained classification

Such approaches assume that a perturbation manifests as a visible boundary, localized structure, or recognizable defect morphology.

In the present experiment, none of these conditions are satisfied. The thin-film deposition produces:

- No strong gradient discontinuity
- No visually discernible boundary
- No object-level structural emergence
- No predefined defect class

The perturbation is spatially distributed and sub-perceptual, manifesting as micro-scale redistribution of scattering characteristics rather than geometric alteration.

As a result, a conventional edge-driven or classification-based system would lack a deterministic separation criterion within this ROI. There is no explicit object to segment, no boundary to trace, and no labeled defect signature to match.

In contrast, the physics-anchored drift framework operates relative to a defined reference state rather than searching for object-level anomalies. Deviation is measured as a distributed change in surface

response, enabling detection of sub-visible state transitions without reliance on edge contrast or trained defect exemplars.

This distinction reflects a fundamental difference in methodology: conventional systems attempt to recognize anomalies, whereas the present framework quantifies conformance loss relative to a fixed physical baseline.

8.3 Deterministic Operation

This system:

- Requires no training data.
- Produces interpretable physical metrics.
- Operates against a defined baseline.
- Aligns with metrology-style conformance verification.

---

9. Industrial Implications

Applications include:

- CMP haze detection.
- Thin-film coating validation.
- Surface contamination detection.
- Cleaning verification.
- Early distributed instability monitoring.

The ability to detect sub-visible distributed perturbations enables earlier intervention in process control environments.

---

10. Limitations and Future Work

Current limitations:

- Single film thickness.

- Single capture per condition.
- No temporal evaporation sequence.

Future work should include:

- Repeated baseline and detect captures.
- Film thickness gradient study.
- Time-decay analysis.
- Comparative structural defect testing.

Repeatability testing is required to establish statistical variance bounds and detection thresholds for industrial deployment.

---

11. Conclusion

This experiment demonstrates:

- Thin-film deposition can be visually indistinguishable from baseline.
- Deterministic drift metrics detect measurable distributed deviation.
- Mean drift increases from 0 to 58.79.
- Distributed activation score increases from 0 to 59.02.
- No object-level topology emerges (strt_Lcc ≈ 0.0103).

Physics-anchored drift quantification demonstrates measurable separation under controlled conditions of distributed thin-film conformance loss without machine learning, probabilistic classification, or visible defect formation.

The results support deterministic State Conformance measurement as a robust framework for distributed surface perturbation detection.

Appendix A — Definitions of Terms

Drift Field (D(x,y))

The pixel-level deviation magnitude measured relative to a defined Golden reference state.

Mean Surface Drift (drift_mean)

The average deviation magnitude across the defined Region of Interest (ROI).

95th Percentile Drift (drift_p95)

The drift value below which 95% of ROI pixel deviations fall. Provides a high-percentile distributed deviation indicator.

Distributed Drift Score (padr_dist_score)

A structured aggregation metric that quantifies spatially distributed deviation across tiled regions of the ROI.

STRT (Spatial Reference Tiling)

A tiling-based spatial evaluation framework that partitions the ROI into structured cells for activation and topology analysis.

Spatial Activation Fraction (strt_S)

The fraction of ROI tiles whose drift magnitude exceeds a defined activation threshold.

Indicates the spatial extent of deviation.

Largest Connected Component Ratio (strt_Lcc)

The ratio of the largest contiguous activated tile cluster relative to total activated tiles.

- Low values → distributed perturbation
- High values → object-like structural emergence

In this study, strt_Lcc = 0.0103 indicates absence of object-level topology.

State Conformance

A deterministic evaluation of whether a measured surface state matches a defined reference state within tolerance bounds.

Appendix B — Symbol and Metric Definitions

This appendix defines the primary quantitative metrics used throughout this study.

A.1 Drift Field

Let:

D(x, y)

represent the local deviation magnitude at pixel coordinate (x, y) relative to the Golden reference state.

Drift is computed as a deterministic function of intensity deviation under fixed acquisition parameters.

A.2 Mean Surface Drift (drift_mean)

drift_mean = (1/N) Σ D(x, y)

where N is the number of pixels within the ROI.

Represents average deviation magnitude across the ROI.

A.3 95th Percentile Drift (drift_p95)

drift_p95 represents the 95th percentile value of D(x, y) within the ROI.

Provides an upper-bound distributed deviation indicator that is less sensitive to single-pixel outliers than drift_max.

A.4 Maximum Drift (drift_max)

Maximum observed D(x, y) value within the ROI.

Represents peak local deviation intensity.

A.5 Distributed Drift Score (padr_dist_score)

A spatially aggregated metric quantifying distributed deviation across structured tiles within the ROI.

Higher values indicate spatially widespread deviation rather than isolated pixel noise.

A.6 Spatial Activation Fraction (strt_S)

Fraction of ROI tiles exceeding a defined drift activation threshold.

Indicates spatial extent of deviation.

A.7 Largest Connected Component Ratio (strt_Lcc)

Ratio of the largest connected activation region relative to total activated area.

Low values indicate distributed activation.
High values indicate object-like structural emergence.

Appendix C — Imaging Parameter Lockdown Specification

To ensure measurement integrity, the following acquisition constraints were enforced:

- Auto exposure: disabled
- Auto gain: disabled
- Auto white balance: disabled
- Auto contrast normalization: disabled
- Auto gamma correction: disabled
- Adaptive enhancement: disabled

All parameters were manually configured and held constant between Golden and Detect captures.

Brightness and raw gain were set conservatively to preserve dynamic range and maintain linear sensor response.

No post-processing normalization was applied prior to drift computation.

This ensures deterministic state comparison.

Appendix D — ROI Selection Criteria

The Region of Interest (ROI) was defined according to the following constraints:

1. Identical pixel coordinates in Golden and Detect frames

2. No boundary adjacency
3. Representative substrate texture
4. Absence of pre-existing structural anomalies
5. Spatial isolation from illumination gradient edges

The ROI was defined prior to quantitative analysis to eliminate algorithmic selection bias.

Appendix E — Interpretation Matrix for Distributed vs Structural Perturbation

Metric Pattern	Interpretation
High drift_mean + High padr_dist_score + Low strt_Lcc	Distributed thin-film redistribution
High drift_mean + High strt_Lcc	Object-level structural defect
Low drift_mean + Low activation	Stable conformance
High drift_max only	Localized spike / noise candidate

In this study:

- drift_mean = 58.79
- padr_dist_score = 59.02
- strt_Lcc = 0.0103

This corresponds to a distributed thin-film signature.

Appendix F — Deterministic Conformance Principle

State Conformance is defined as:

A system state S(t) conforms to reference state S_0 if D(x,y,t) ≈ 0 within defined tolerance bounds across the ROI.

Conformance degradation is quantified deterministically as:

ΔC = f(drift_mean, padr_dist_score, strt_S, strt_Lcc)

No probabilistic classification or learned model parameters are used in this determination.

Deviation is measured relative to a fixed physical baseline.

Appendix G — Limitations of Visual Inspection

Human visual perception is limited by:

- Contrast sensitivity thresholds
- Spatial frequency bias
- Adaptation to uniform illumination
- Lack of quantitative aggregation

Thin films often modify micro-scattering properties without forming perceptually salient edges.

This experiment demonstrates measurable distributed deviation despite visual indistinguishability.

Appendix H — Quantization and Signal Integrity

Appendix H — Quantization and Signal Integrity

All drift computations were performed on raw image data captured under fixed manual acquisition parameters. No adaptive normalization, histogram rescaling, or nonlinear tone mapping was applied prior to analysis.

Bit Depth
Images were processed in 8-bit intensity space (0–255), consistent across both Golden and Detect captures. Quantization resolution was therefore fixed and identical between conditions. Because acquisition parameters were locked and no rescaling was applied, the quantization floor remained constant for all comparisons.

Channel Handling
Drift computation was performed on grayscale intensity values derived from the native sensor output. No per-channel RGB drift separation was used in this study. Color channels were not independently analyzed, and no color-space transformations were applied prior to drift extraction.

Drift Domain Definition
Drift magnitude D(x, y) was computed directly on grayscale intensity differences between Golden and Detect frames. This ensures that reported deviation reflects luminance-based scattering variation rather than color-channel amplification or selective channel weighting.

Because the acquisition pipeline was parameter-locked and linearized, measured deviation arises exclusively from physical surface state change rather than quantization variability, channel rebalancing, or acquisition normalization artifacts.

Linearity validation (operational). Acquisition parameters were fixed and conservative (no auto exposure/gain/gamma). Under these conditions, drift was computed directly from raw intensity differences, assuming operation within the sensor's approximately linear response regime. A full linearity characterization (step-response / flat-field sweep) is reserved for future work.