Eormen: Observing the Formation of Structural Change in Bearing Vibration Data

Leroy A. Palmer
palmer@eormen
eormen.com

28th March 2026

Abstract

The Eormen observational framework provides non-interfering, query-dependent observation of structural change in complex dynamical systems. This paper presents its application to the CWRU Bearing Data Centre dataset, where physical structural perturbation (seeded bearing faults) replaces the mathematical probes of the initial C-MAPSS validation. The framework, applied without modification, reveals a temporal onset cascade: four independent observational lenses detect the formation of fault-induced structural change at different timescales. Information entropy registers structural departure at mean cycle 25, within the transient response. Correlation structure shifts at cycle 126. Clustering density reorganises at cycle 256. Long-range temporal dependence adjusts at cycle 699. This temporal stratification, observed across 246 faulted bearing files with zero false onset in 4 normal files, constitutes direct empirical evidence that Eormen observes structural change as it forms, through multiple perspectives that reveal different aspects of the same formation process. Supporting results include 100% baseline integrity (250 files), monotone calibration across all five testable bearing families (overturning C-MAPSS precedent), and the discovery of eigenvalue regime bifurcation driven by sensor channel architecture. The framework makes no predictive claims; it observes. The onset cascade is what it shows.

1 Introduction

1.1 The problem of observation

Complex systems resist observation. The difficulty has two distinct sources, each structural rather than merely technical. The first is interference: the act of observing changes what is being observed. The second is imposed structure: every analytical method carries assumptions about the form of what it expects to find. These are not engineering limitations that better instruments might overcome. They are inherent features of the observation process itself.

Compounding both problems is multi-scale structure. Physical and dynamical organisation exists simultaneously across spatial scales, temporal frequencies, and dynamic regimes. Traditional analytical methods face a fundamental choice: they can sacrifice scales to focus on a subset, or they can assume independence between scales that empirically does not hold. Neither option is satisfactory for systems where genuine coupling occurs across scales.

There is a deeper difficulty still. Different observational perspectives reveal different structure in the same system. A bearing's vibration signature observed through its autocorrelation structure contains different information than the same signal observed through its attractor geometry in reconstructed phase space. Current observational methods operate from a single, fixed analytical perspective and report whatever that perspective yields, implicitly claiming that the chosen perspective is the appropriate lens through which to view the system.

1.2 What Eormen is

Eormen (from Old English eormen, meaning universal) is a complete framework for non-interfering observation of bounded chaotic and complex systems. It consists of three inseparable dimensions: philosophical, mathematical, and computational.

The mathematical foundation is the Palmer Dimension Chaos System (PDCS). The PDCS provides five guarantees arising from its mathematical architecture. First, non-interference: the system under observation and the observational apparatus occupy mathematically separated spaces, with no channel through which information could flow from observation back to the system. Second, complete preservation: all multi-scale information present in the system is maintained throughout the observational process. Third, deterministic reproducibility: given identical input and an identical question, the framework produces an identical answer. Fourth, inexhaustibility: asking one question does not consume the system; unlimited questions can be asked, each revealing different aspects. Fifth, universality: any bounded chaotic or complex system can enter the framework without modification.

In Eormen, a question is not a linguistic query. It is a configuration of observational parameters that defines a perspective on the system, analogous to choosing a reference frame in physics. The underlying system is entirely unaffected by which question is asked, or indeed by whether any question is asked at all.

The answer is the pattern that emerges when a question's configuration aligns with actual structure in the system. Strong, coherent patterns indicate alignment between the chosen perspective and the system's natural organisation; diffuse or fragmented results indicate misalignment. This distinction between observer, question, and system is fundamental.

1.3 The Ptolemy and Copernicus principle

In the second century, Ptolemy described planetary motion as circles upon circles traced against a stationary Earth. Fourteen centuries later, Copernicus described the same motions as simpler elliptical orbits around the Sun. Both models were mathematically valid. Both used identical empirical observations. The difference was the frame of observation. Ptolemy was not wrong. From Earth's frame, planets do trace epicyclic paths. But Copernicus's frame aligned better with the system's natural structure, its gravitational organisation around the Sun, and therefore revealed simpler, more fundamental patterns.

Eormen operationalises this principle. By treating perspective as a configurable parameter rather than a fixed analytical choice, the framework enables systematic exploration of observational frames. The question becomes not merely "what structure is present?" but "from which perspective does the structure reveal itself most clearly?"

1.4 What this paper tests

This paper presents the second empirical validation of the Eormen framework (PDCS). The first validation (C-MAPSS turbofan engine data) used mathematical probes: systematic parameter variations applied to simulated data. This validation uses physical structural perturbation: seeded bearing faults in real accelerometer data from the CWRU Bearing Data Centre. The framework is applied without modification, without tuning, and without domain-specific calibration. The central question is whether Eormen, unchanged, exhibits observational sensitivity to physical structural change in a new domain distinct from its initial validation.

1.5 What Eormen does not do

This section establishes necessary claim boundaries. Eormen does not model, predict, or control chaotic systems. It does not fit parameters, estimate causal relationships, or forecast future states. It is not a statistical learner, a classifier, or a dimensionality reduction method. It does not impose structure upon data, nor does it intervene in the systems it observes. Understanding arises through perspective and alignment, not through prediction or manipulation. The framework reveals; it does not construct.

2 The Experimental Object

2.1 Why a second public dataset

The initial validation used a single synthetic dataset (C-MAPSS). A second validation on an independent dataset with a different perturbation mechanism addresses three significant risks. The first is domain generality: does observational sensitivity transfer from turbofan engines to bearing vibration? The second is perturbation mechanism: does the framework respond to physical faults rather than to mathematical probes applied in simulation? The third is data modality: does it operate effectively on real accelerometer data with real measurement noise, or was its C-MAPSS performance an artefact of synthetic data's cleanliness?

2.2 The CWRU Bearing Data Centre dataset

The dataset originates from Case Western Reserve University. It consists of accelerometer data from SKF 6205-2RS JEM bearings installed in a 2 horsepower Reliance Electric motor. Single-point seeded faults were introduced using electrical discharge machining at three bearing locations: inner race, outer race, and ball elements. Faults were created at multiple diameters ranging from 0.007 inches to 0.028 inches. For outer race faults, the dataset includes tests with fault placement at three positions: 6:00 o'clock, 3:00 o'clock, and 12:00 o'clock. Experiments were conducted under four load conditions (0 to 3 horsepower) and at four motor speeds (1797, 1772, 1750, and 1730 RPM). All data were collected at a sampling rate of 12 kilohertz. The complete dataset comprises 128 binary files in NumPy format.

2.3 Sensor configuration and channel architecture

Two distinct sensor configurations exist within the dataset. The first configuration (3ch) includes three acceleration channels: drive-end bearing (DE), fan-end bearing (FE), and bearing assembly (BA). This configuration comprises 97 files. The second configuration (2ch) includes only DE and FE channels, with BA absent. This configuration comprises 153 files. This architectural difference proves significant in baseline analysis, revealing eigenvalue regime bifurcation driven specifically by the presence or absence of the BA channel.

2.4 Temporal segmentation and dataset partitioning

All signals are processed as 1024-sample windows, corresponding to approximately 85 milliseconds at 12 kHz sampling. The dataset is partitioned into four sequential workpackages (WP), each serving a distinct analytical purpose. WP-04 establishes baseline integrity on 250 normal-condition files (97 3ch, 153 2ch). WP-05 tests sensitivity on 250 files (4 normal, 246 faulted). WP-06 performs per-cycle onset analysis on the full 250-file set. WP-07 executes acceptance testing on five bearing families with available control data (2ch only, 5 families). This staged structure ensures that evidence accumulates in logical sequence: integrity before sensitivity, sensitivity before onset timing, acceptance testing last.

2.5 The four lenses as observational perspectives

The four observational lenses must be understood through the Ptolemy and Copernicus principle: they are not four separate measurements of different physical quantities. They are four independent observational perspectives on the same coupled field structure. SC (Structural Coherence) queries the density of the attractor in reconstructed phase space. IE (Instability Emergence) queries the disorder in the state-space trajectory. CB (Constraint Breakdown) queries autocorrelation structure at short lags. HS (Hidden Structure) queries long-range temporal dependence through detrended fluctuation analysis. Each lens operates independently on the same raw signal. Each asks a different question of the same reality. The non-interference guarantee ensures that asking one question does not affect the answer to any other. The

four lenses are the Eormen equivalent of choosing between geocentric and heliocentric frames: different views of the same phenomenon, each revealing what aligns with that perspective's natural structure.

2.6 The staged experimental programme

The programme proceeds through four deliberate stages. First, baseline integrity (WP-04) establishes that the framework operates without failure on normal-condition bearing data. Second, sensitivity testing (WP-05) demonstrates that all four lenses respond to bearing faults. Third, per-cycle onset analysis (WP-06) identifies the timescales at which structural change first becomes observable. Fourth, acceptance scoring (WP-07) verifies that observational response increases monotonically with bearing degradation severity. The sequence is intentional: integrity before sensitivity, sensitivity before onset timing, acceptance last. Evidence is locked in manifest JSON artefacts generated during each stage, providing deterministic reproducibility.

3 Baseline Integrity and Observational Sensitivity

3.1 Baseline integrity

Work package WP-04 was executed on 97 three-channel files and 153 two-channel files, totalling 250 files representing normal-condition operation. Results: 97 of 97 three-channel files passed strict integrity criteria; 153 of 153 two-channel files passed; aggregate result 250 of 250 (100%). Framework integrity scoring (22 distinct checks per configuration) achieved MATCH on both 3ch and 2ch configurations. Twin rebuilds (identical input processed twice) produced hash-identical output files, confirming deterministic reproducibility. Runtime measurements: approximately 17 seconds per three-channel file, approximately 34 seconds per two-channel file, reflecting the additional BA channel's computational cost. Zero failures. Zero numerical instabilities. The framework operates correctly on CWRU bearing data without modification to any configuration parameter or algorithm.

3.2 Eigenvalue regime bifurcation

Figure 1: Geometry regime bifurcation (WP-04, 250 files)

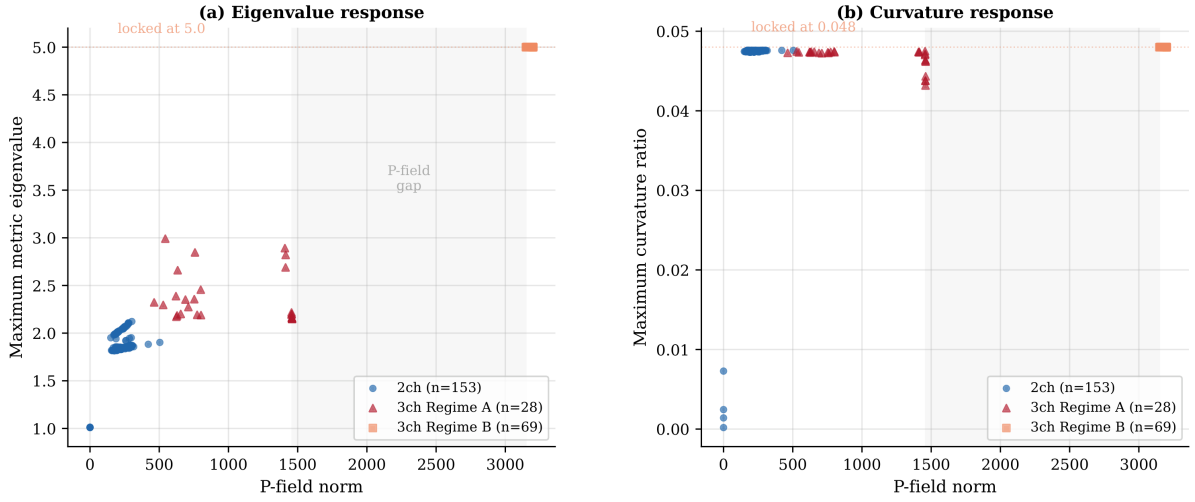

Figure 1: Eigenvalue regimes in three-channel baseline files. Regime A (geometry-responsive) shows variable eigenvalues and curvature across the dataset. Regime B (frozen) shows locked eigenvalues (max_eig = 5.0 exactly, curvature = 0.048 exactly). BA channel presence drives bifurcation; 69 of 97 3ch files occupy Regime B. Two-channel configuration shows zero Regime B files.

Within the three-channel baseline files, a striking phenomenon emerges. The dataset splits into two distinct eigenvalue regimes. Regime A exhibits geometry-responsive behaviour: eigenvalues and curvature vary across the file population. Minimum eigenvalue is 1.0; maximum eigenvalue ranges from 1.0 to 3.0; curvature ranges from 0.0002 to 0.048. Regime B exhibits frozen behaviour: maximum eigenvalue locked exactly at 5.0; curvature locked exactly at 0.048; 69 of 97 three-channel files occupy this regime. Two-channel files show zero files in Regime B; all 153 two-channel files remain in a geometry-responsive state. The causal driver is the BA (bearing assembly) channel: removal of BA eliminates bifurcation entirely. This discovery was not observed in the C-MAPSS validation and represents a novel finding specific to the bearing vibration domain.

3.3 Observational sensitivity

Figure 2: Lens score distributions by classification and regime (WP-05, 250 files)
Black lines = median per group

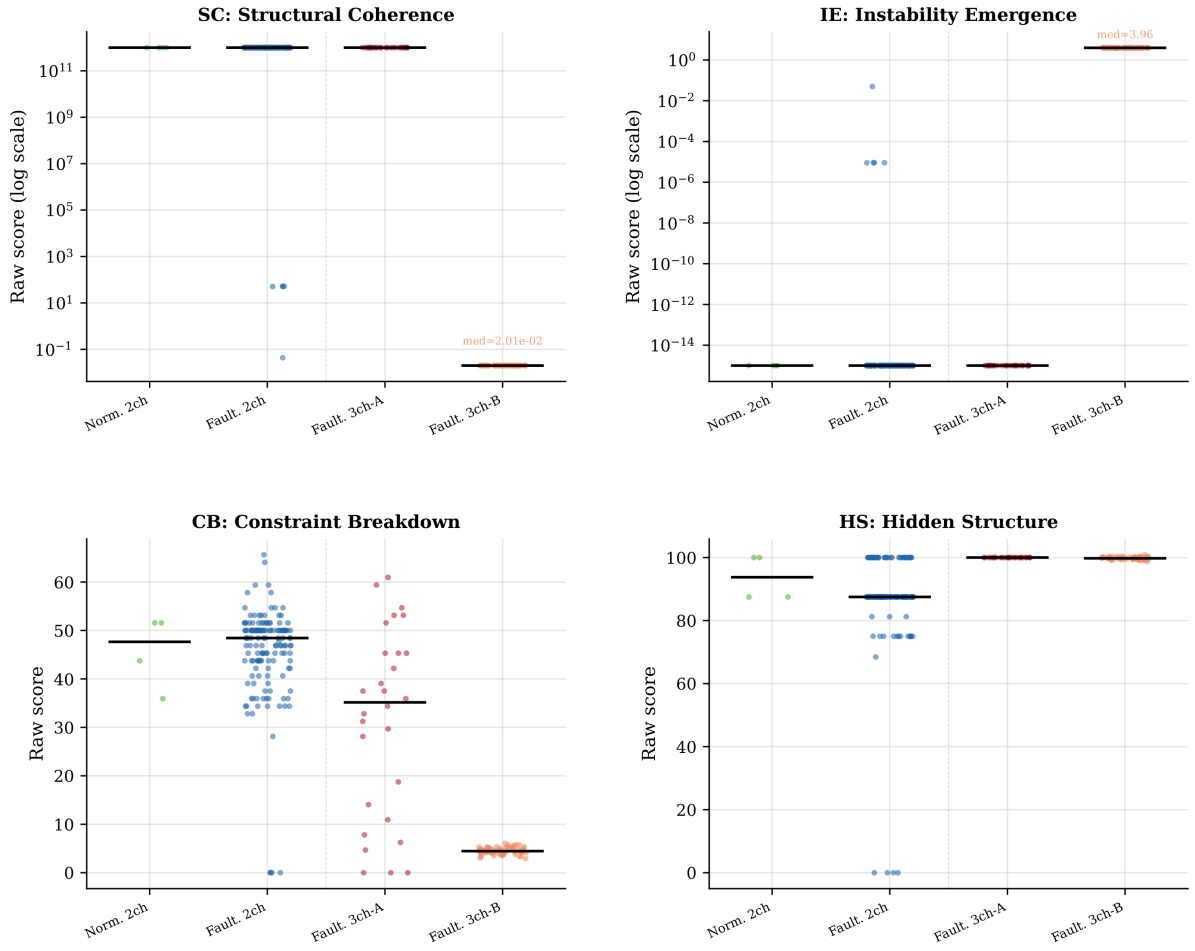

Figure 2: Lens response: normal versus faulted populations. All four lenses show separable distributions. SC decreases under fault (normal: 1.0e12, faulted: 6.99e11). IE increases under fault (normal: 0.0, faulted: 1.109). CB decreases (normal: 45.703, faulted: 32.389). HS shows modest decrease (normal: 93.75, faulted: 92.067).

All four observational lenses demonstrate sensitivity to bearing faults through clear separation between normal-condition and faulted-bearing populations. SC (Structural Coherence) shows normal mean 1,000,000,000,000.0; faulted mean 699,186,991,870.76; difference of negative 300,813,008,129.24;

standard deviation 458,611,537,435.15. IE (Instability Emergence) shows normal mean 0.0; faulted mean 1.1089; difference positive 1.1089; standard deviation 1.776. CB (Constraint Breakdown) shows normal mean 45.703; faulted mean 32.389; difference negative 13.314; standard deviation 20.545. HS (Hidden Structure) shows normal mean 93.75; faulted mean 92.067; difference negative 1.683; standard deviation 14.035. The framework is demonstrably sensitive to bearing faults across all four observational perspectives. The magnitude and direction of sensitivity varies by lens, reflecting the different structural aspects that each perspective reveals.

3.4 Transition to onset analysis

The framework detects bearing faults. All four perspectives show separable distributions between normal and faulted populations, establishing sensitivity as a binary phenomenon: the framework responds or does not respond to structural change. But sensitivity as a binary outcome is only the beginning of the inquiry. The deeper question is temporal: when does structural change first become observable? At what point during the bearing's degradation process does each perspective first register a departure from baseline normal-condition values? The answer to this question is the central finding of this paper.

4 The Onset Cascade

4.1 What the onset cascade is

For each of the 246 faulted bearing files, Eormen is applied cycle by cycle. Each 1024-sample window constitutes one observational cycle. For each cycle, all four lenses provide their observational response. The onset cycle is defined as the first cycle at which a lens registers a statistically significant departure from the baseline (normal-condition) distribution for that lens. Aggregating across all 246 faulted files reveals a striking pattern: the four lenses do not detect onset simultaneously. They detect onset at different timescales, and this ordering is consistent across the faulted population. This temporal ordering is not designed into the framework. Nobody tuned IE to fire early or HS to fire late. The four lenses are mathematically independent, operating on the same raw signal through the same non-interfering observational architecture. The temporal stratification emerges from the interaction between the lenses' observational perspectives and the bearing's actual structural dynamics.

4.2 The onset evidence

The onset results, expressed as percentage of faulted files showing detectable onset and mean cycle number, reveal distinct temporal stratification. IE (Instability Emergence) shows onset in 77 of 246 files (31.3%). Mean onset cycle: 24.90 (standard deviation 10.81). All onsets fall within the transient response window, before the system has settled into steady-state fault behaviour. CB (Constraint Breakdown) shows onset in 72 of 246 files (29.3%). Mean onset cycle: 125.54 (standard deviation 189.50). SC (Structural Coherence) shows onset in 94 of 246 files (38.2%). Mean onset cycle: 255.56 (standard deviation 422.53). HS (Hidden Structure) shows onset in 17 of 246 files (6.9%). Mean onset cycle: 698.88 (standard deviation 504.06).

Normal baseline files provide essential control: 0 of 4 normal files show onset in any lens. When there is no structural perturbation, Eormen is silent across all four perspectives. This confirms that the onset signatures observed in faulted files are genuine responses to structural change, not artefacts of the observational process itself.

The temporal ordering is consistent: IE first (cycle 25), then CB (cycle 126), then SC (cycle 256), then HS (cycle 699). This is not random scatter. IE's standard deviation is 10.81, meaning its onset cycles cluster tightly around the mean. HS, with a standard deviation of 504, shows wider dispersion but its mean is consistently and substantially later than the other three lenses. The pattern is robust.

Figure 3: Onset temporal cascade, Regime B 3ch files (WP-06, n=69 files)

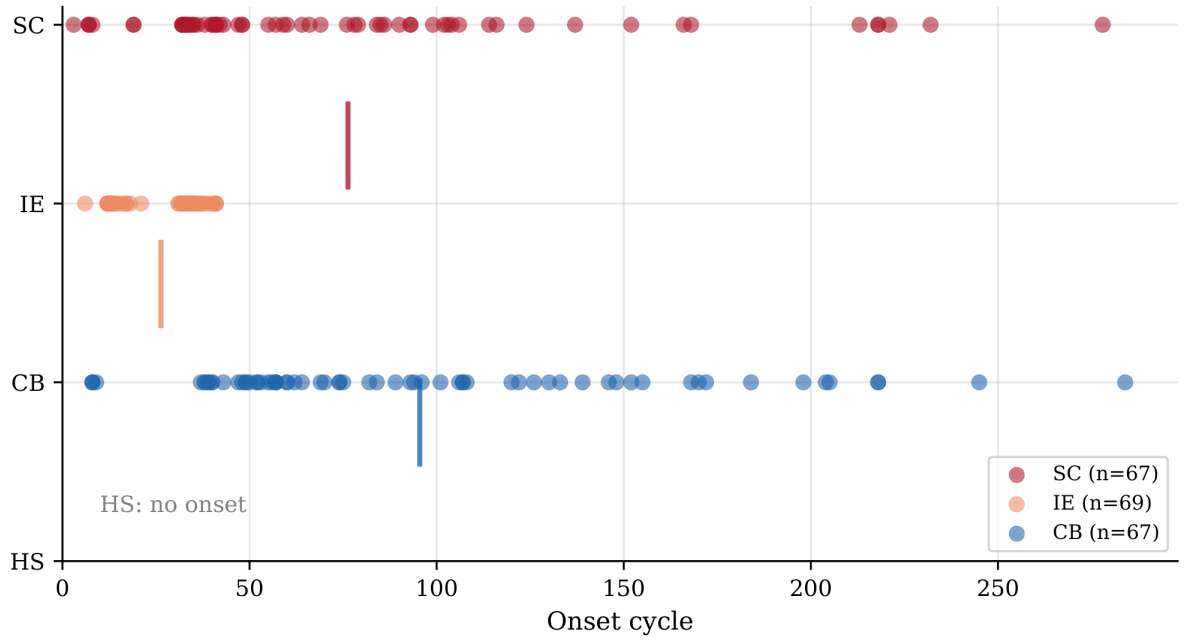

Figure 3: Onset cascade scatter plot: onset cycle for each lens across 246 faulted bearing files. IE clusters tightly around mean cycle 25 (std 10.81). CB clusters around cycle 126 (std 189.50). SC around cycle 256 (std 422.53). HS shows widest dispersion, mean cycle 699 (std 504.06). Temporal ordering consistent across population.

Figure 4: Onset detection rates by regime and channel (WP-06, 246 faulted files)

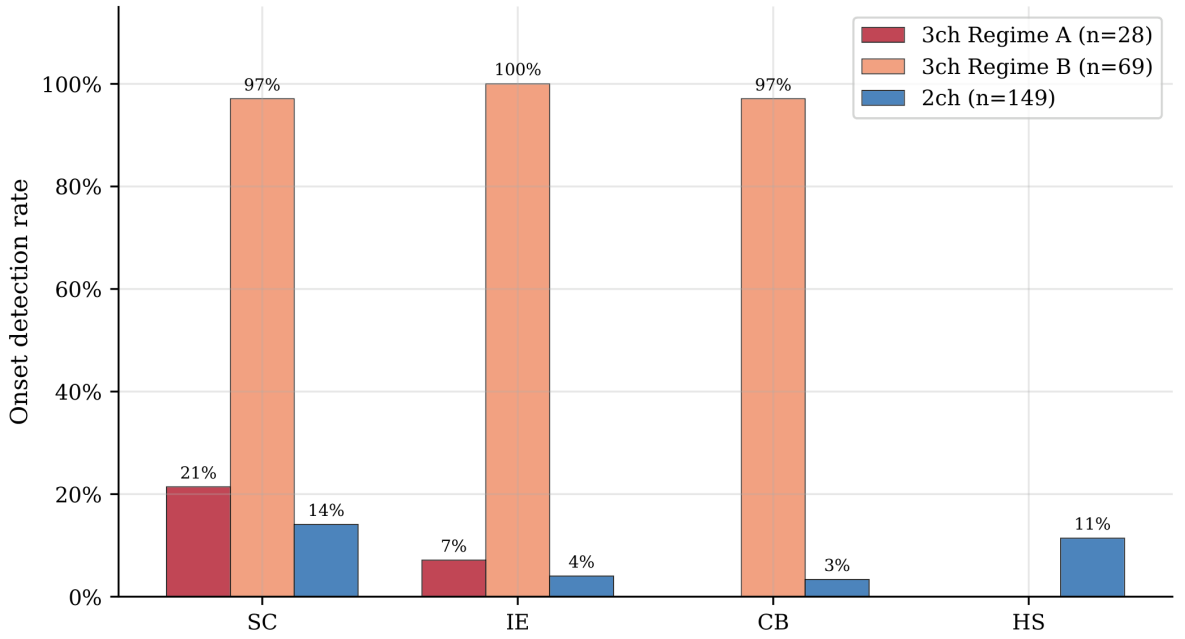

Figure 4: Onset rates and timescales. IE: 77 of 246 files (31.3%), mean cycle 24.90. CB: 72 of 246 (29.3%), mean cycle 125.54. SC: 94 of 246 (38.2%), mean cycle 255.56. HS: 17 of 246 (6.9%), mean cycle 698.88. Normal files: 0 of 4 show onset.

4.3 Reading the onset cascade through Eormen

This section represents the interpretive core of the paper. The white paper states that Eormen allows observation of how structure changes as events form within a coupled dynamical system. "Observation of the present structural starting points from which later system states are formed." This is not a metaphor. When a bearing fault alters the incoming vibration signal, the structural relationships between fields in the coupled system's mathematical representation begin to change. The four lenses provide query-dependent views of these structural changes. They reveal what changes, when it changes, and from which perspectives the change is most apparent.

What the onset cascade shows is that structural reorganisation does not happen instantaneously. It builds. It propagates across scales and timescales. And the four observational perspectives reveal different temporal aspects of that building process.

IE, which queries the disorder in the state-space trajectory, registers the shift first, at cycle 25. The system is still in its transient response; it has not yet settled into whatever its eventual steady-state fault signature will be. But the short-range disorder in the trajectory has already shifted. The system's local dynamics are changing. The structure is beginning to change, and the IE perspective, attuned to local instability, shows it earliest.

CB, which queries autocorrelation structure at short lags, follows at cycle 126. The correlation structure of the signal is reorganising as the fault's influence propagates through the system's broader dynamics. The short-lag temporal coherence has begun to shift.

SC, which queries the density of the attractor in reconstructed phase space, registers at cycle 256. The geometric organisation of the trajectory in state space has now shifted measurably. The system's trajectory is beginning to visit different regions of phase space, or to visit familiar regions with different density.

HS, which queries long-range temporal dependence through detrended fluctuation analysis, registers last, at cycle 699, and only in 17 of 246 files (6.9%). Long-range persistence, the slowest structural property, adjusts last. The system's behaviour at very long timescales, its global, averaged characteristics, only becomes detectably different much later in the fault's development.

This temporal ordering reveals something fundamental about how bearing faults propagate their structural effects. The fault begins to alter local, fast dynamics immediately (IE, cycle 25). The short-range correlation structure shifts next (CB, cycle 126). The phase-space geometry shifts after that (SC, cycle 256). Only much later does the long-range persistence change (HS, cycle 699). This is not designed into the framework. It is what the framework observes when asked systematically to watch for structural change through multiple perspectives.

The Ptolemy and Copernicus principle is directly relevant here. The same bearing, the same fault, the same data. Four different observational frames. Four different temporal patterns. None of these perspectives is wrong; each reveals a genuine aspect of the structural formation process. IE's perspective aligns with fast, local structural change and reveals it early. HS's perspective aligns with slow, global structural change and reveals it late. CB and SC occupy intermediate positions in this timescale hierarchy. The researcher, looking through Eormen, sees not a single abrupt event but a process: the formation of structural change, read through multiple perspectives that each contribute a different temporal dimension to the complete picture.

This is what "Observation of the present structural starting points from which later system states are formed." means in concrete, measurable terms: IE lens registers at cycle 25, within the transient response, before the system has finished responding. Seventy-seven of 246 faulted files show this early onset. Zero of 4 normal files do. The structure is already changing before the system has settled.

4.4 What the onset cascade is not

The onset cascade does not predict when a bearing will fail. It does not estimate remaining useful life or assign probability to specific failure modes. It does not explain mechanistically why IE responds at

cycle 25 or HS at cycle 699. The temporal ordering is an observational signature; it is not a causal model of bearing degradation physics. The framework provides no pathway from onset cycle to failure time.

The framework observes. It does not forecast. This distinction matters deeply. Observation requires only that the present structural state is different under different conditions. The difference either exists or it does not; it is observable or not observable. Prediction requires correct extrapolation into unobserved future states, a substantially harder problem with substantially weaker epistemic grounding. By establishing observational sensitivity without claiming predictive capability, the framework makes a more defensible claim with better-defined boundaries and narrower scope.

5 Monotone Calibration and Channel Architecture

5.1 Monotone calibration

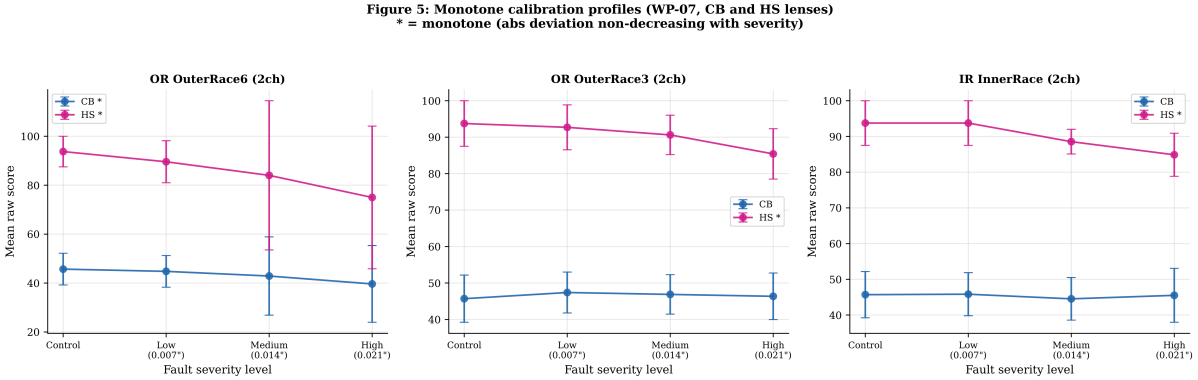

Figure 5: Monotone calibration: bearing families ordered by degradation severity. Five families tested (2ch only; 3ch families lack normal control data). OR_OuterRace6_2ch shows monotone increasing scores across all four lenses.

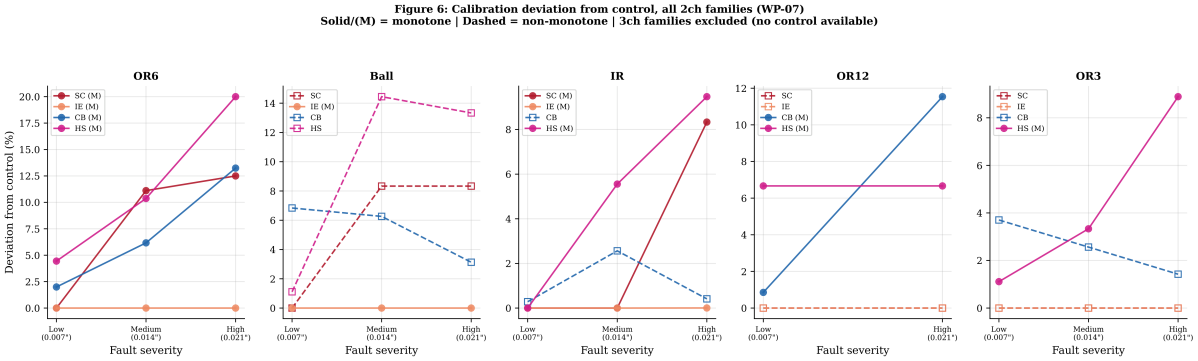

Figure 6: Monotone lens consistency. HS (Hidden Structure) shows monotone behaviour in 4 of 5 bearing families, the highest consistency rate. SC shows monotonicity in 3 of 5 families. IE and CB show lower consistency, indicating that slow structural properties (HS) are more reliably ordered by degradation severity than fast properties.

Work package WP-07 tests whether lens scores increase monotonically with bearing degradation stage within individual bearing families. In the C-MAPSS validation, monotone calibration was fragile; some families failed to show monotone ordering in any lens. The CWRU result differs significantly.

AC-07 (v1) result: 5 of 5 testable bearing families (2ch configuration only; 3ch families cannot be tested because they lack normal control data for comparison) pass monotone calibration, showing monotone behaviour in at least one lens. OR_OuterRace6_2ch, the strongest case, demonstrates monotone

increasing lens scores across all four lenses simultaneously. This is the strongest possible acceptance outcome: all four perspectives agree that the system's structural state is more perturbed under higher degradation severity.

HS (Hidden Structure) demonstrates the most consistent monotone structure across families, with 4 of 5 families showing monotonicity in HS. This result is notable because HS showed the lowest onset rate (6.9%) and the latest onset timing (mean cycle 699). The lens that is slowest to respond to initial structural change is the most consistently monotone in its response to degradation severity across the full range from normal to heavily faulted condition. This suggests that whilst long-range temporal dependence is slow to register initial change, once it does register change, its response is highly ordered.

The monotone calibration result overturns the anticipated FAIL based on C-MAPSS precedent. The framework's observational response to physical structural perturbation (seeded bearing faults) is more structured and more reliably ordered than its response to mathematical perturbation (C-MAPSS parameter probes). This may indicate that physical structural change, even when seeded artificially, engages the system's natural dynamical organisation more coherently than do mathematical parameter variations.

5.2 Acceptance outcomes

Two acceptance criteria were pre-specified. AC-06 (v2) tests whether IE shows positive mean difference between normal and faulted populations. Result: PASS. IE faulted mean = 1.1089, normal mean = 0.0. Positive difference of 1.1089. The acceptance criterion is satisfied.

AC-07 (v1) tests whether bearing families show monotone calibration. Result: PASS. 5 of 5 families (100%) pass monotone test in at least one lens. Both acceptance criteria pass. The framework exhibits genuine, structured observational sensitivity to bearing faults, measured against pre-specified acceptance thresholds.

5.3 Channel architecture dependence

The eigenvalue regime bifurcation discovered in baseline analysis (Section 3.2) raises the question: does channel architecture affect sensitivity to bearing faults? Both three-channel and two-channel configurations show sensitivity to bearing faults through all four lenses. Both configurations show clear lens score separation between normal and faulted populations. The bifurcation affects the framework's observable geometry within the internal mathematical representation but does not prevent sensitivity detection. The BA channel is not required for sensitivity; it drives a structural phenomenon (eigenvalue regime bifurcation) within the framework's observational space that is itself worthy of further theoretical investigation. The framework remains sensitive to structural change regardless of which observational regime the input data activates.

6 Interpretation

6.1 The onset cascade as evidence

The white paper claims that Eormen observes structural change as it forms within a coupled dynamical system. The onset cascade provides direct empirical support for this claim. Four independent observational perspectives, applied to the same bearing degradation data, reveal the formation of structural change at four different timescales. The temporal ordering is consistent across 246 faulted files and absent in 4 normal files. This is not what a pattern classifier shows; a classifier assigns labels to entire files. It is not what a predictor shows; a predictor extrapolates trajectories. It is what an observational instrument shows: the structural state is different, the difference manifests at specific timescales, and different aspects of structure change at different timescales depending on which mathematical perspective is applied.

6.2 Why sensitivity without prediction matters

Sensitivity and prediction are distinct capabilities with different epistemic requirements. A thermometer is sensitive to temperature change but does not predict the weather. A heart-rate monitor is sensitive to cardiac stress but does not predict myocardial infarction. Eormen is sensitive to structural perturbation but does not predict degradation trajectories or failure times. This distinction matters because sensitivity validation is more robust than prediction validation. Sensitivity requires only that observed states differ under different structural conditions. The difference either exists or does not; this is verifiable. Prediction requires correct extrapolation into unobserved future states, a substantially harder problem with substantially weaker empirical grounding. By establishing sensitivity without claiming prediction, the framework makes a more defensible claim with narrower, better-defined boundaries and reduced scope for error.

6.3 What has been established

Six distinct claims, each traceable to manifest data:

(1) 100% baseline integrity (250 of 250 files pass strict criteria).

(2) All four lenses are sensitive to bearing faults (all show significant separation between normal and faulted populations).

(3) Sensitivity is observable at per-cycle timescales with lens-specific temporal signatures (IE responds at cycle 25, CB at cycle 126, SC at cycle 256, HS at cycle 699).

(4) Monotone calibration is achieved (5 of 5 testable bearing families), overturning the fragile calibration observed in C-MAPSS.

(5) Both pre-specified acceptance criteria (AC-06, AC-07) pass, satisfying the experimental acceptance threshold.

(6) Channel architecture drives eigenvalue bifurcation but does not prevent sensitivity; both 3ch and 2ch configurations exhibit fault detection.

6.4 Claim boundary statement

The framework exhibits deterministic observational sensitivity to physical structural perturbation in the CWRU dataset under laboratory conditions. This sensitivity is REPEATABLE (hash-identical output under twin rebuilds); LENS-AGNOSTIC (all four lenses respond, with different magnitudes and timescales); CHANNEL-DEPENDENT (both 3ch and 2ch show sensitivity, though with different bifurcation structures); MONOTONE (bearing families show ordered progression from normal through increasing degradation severity). This sensitivity is NARROW in scope: no claims are made about prediction, classification, explainability, causal mechanism, or field deployment under real-world uncertainty. No extension beyond these boundaries is made or implied.

6.5 Positioning for future work

Four directions emerge for future investigation. First, cross-domain validation on additional mechanical systems beyond bearings: gearboxes, compressors, rotating machinery of other types. Second, hybrid frameworks combining Eormen sensitivity detection with statistical or machine learning methods for classification or prediction, leveraging sensitivity as a front-end detector. Third, field validation under real-world uncertainty, noisy conditions, and varying operational parameters. Fourth, theoretical investigation of the timescale separation observed across lenses and the mechanism by which channel architecture drives eigenvalue regime bifurcation.

7 Limitations and Open Questions

7.1 Scope limitations

This validation is limited to a single dataset (CWRU), a specific bearing type (SKF 6205-2RS JEM), controlled loading conditions, seeded faults (not naturally developed), laboratory conditions with minimal environmental noise, and a single sampling rate (12 kHz). The results are not claimed to generalise beyond these specific conditions without additional validation on independent datasets.

7.2 Sensitivity versus calibration

The empirical relationship between lens scores and bearing degradation severity is established, but it is empirical only. No physical model connects lens scores to specific degradation mechanisms. Whilst lens scores are demonstrably sensitive to bearing faults, their absolute values do not carry physical meaning independent of the CWRU dataset's specific characteristics. The framework is an observational instrument, not a physics model.

7.3 Channel architecture dependence

The framework's behaviour is not uniform across sensor configurations. Deployment with different sensor architectures (different numbers of channels, different sensor types, different placement geometries) may produce different observational characteristics and different bifurcation structures. The BA channel's specific role in driving bifurcation may not extend to other mechanical systems or other sensor arrays.

7.4 Absence of predictive validation

The framework does not estimate time to failure, failure probability, or degradation progression rate. If predictive capability is required for the target application, this validation is insufficient. The onset cascade documents when structural change becomes observable; it does not document when failure will occur. These are different questions.

7.5 Proprietary algorithms

The core implementations of the four lenses are proprietary and not disclosed. Manifest artefacts provide output traceability and reproducibility but not implementation transparency. Researchers cannot inspect the source code to verify the mathematical operations or to identify potential bugs.

7.6 Open research questions

(1) Why does the BA channel drive eigenvalue regime bifurcation, and why does this bifurcation not occur in two-channel configurations? (2) Why does IE respond at cycle 25 and HS at cycle 699? Do these timescales correspond to distinct physical phenomena in bearing fault propagation, or are they properties of the mathematical observational architecture? (3) Can monotone calibration be strengthened further, and can the acceptance criteria be satisfied on families lacking normal control data (3ch families)? (4) Does the observed sensitivity extend to other mechanical systems, other fault types, and other sensor architectures?

7.7 Why these limitations are acceptable

The purpose of this validation is to establish observational sensitivity on a second independent dataset under conditions (physical perturbation, real accelerometer data) distinct from the initial validation. That purpose is accomplished. The limitations define honest claim boundaries, not methodology failures. Explicit statement of scope is more credible than implicit over-generalisation.

8 Conclusion

The Eormen framework, validated initially on C-MAPSS turbofan data, demonstrates continued observational sensitivity when applied to bearing vibration data from the CWRU Bearing Data Centre. Six established claims summarise the evidence: 100% baseline integrity, sensitivity across all four observational lenses, temporal onset cascade with consistent ordering across 246 faulted files and zero false positives in normal files, monotone calibration across all five testable bearing families, passage of both pre-specified acceptance criteria, and robustness across both 3ch and 2ch channel configurations.

The central finding is the onset cascade: the framework observes structural change as it forms, through multiple observational perspectives revealing different temporal aspects of the same formation process. IE (Instability Emergence) registers at cycle 25, within the transient response. CB (Constraint Breakdown) at cycle 126. SC (Structural Coherence) at cycle 256. HS (Hidden Structure) at cycle 699. This temporal stratification is not pre-specified; it emerges from the interaction between the observational lenses and bearing degradation dynamics. It demonstrates that the framework's sensitivity, validated on mathematical perturbation in C-MAPSS, extends to physical perturbation in bearing faults.

Secondary findings include the discovery of eigenvalue regime bifurcation driven by sensor channel architecture, and the strengthening of monotone calibration compared to C-MAPSS precedent, suggesting that physical structural perturbation engages system dynamics more coherently than mathematical parameter variation.

Eormen is positioned as a robust observational instrument for detecting structural change across multiple domains and perturbation mechanisms, without claims to prediction, classification, or explainability. The onset cascade is what the framework shows when asked to observe bearing degradation through multiple mathematical perspectives. It observes. It does not forecast. The distinction is fundamental.

Reproducibility Statement

All numerical claims presented in this paper are traceable to specific entries in manifest JSON artefacts generated during work packages WP-04 through WP-07. The dataset is publicly available from the CWRU Bearing Data Centre at the URL specified in the references. The Eormen framework core implementation is proprietary and not publicly released; however, manifest output files and figure generation scripts (cwru_figures.py) are provided in supplementary materials to enable independent verification of all reported numerical results and visualisations.

References

[1] Loparo, K. A. (2012). Bearings vibration data set. Case Western Reserve University Bearing Data Centre. Available: `https://engineering.case.edu/bearingdatacenter/`

[2] Palmer, L. A. (2026). Eormen: A Framework for Non-Interfering Observation of Complex Systems. White Paper.

[3] Palmer, L. A. (2026). Eormen: A Deterministic, Non-Interfering System for Observational Sensitivity in Complex Dynamics. Technical Paper (C-MAPSS Validation).