Trajectory-Only Structural Regime Detection in Deterministic Dynamical Systems

Aryan Gupta
guptaaryanr@gmail.com
ORCID: 0009-0001-8179-5145

Abstract

Deterministic dynamical systems are often analyzed with model-aware diagnostics, but those methods do not answer a narrower question: can a regime-relevant boundary be identified when detection is restricted to trajectories alone? This study evaluates that question with a minimal, non-ML framework for trajectory-only structural regime detection from multivariate/full-state trajectories. The framework builds a trajectory-only indicator atlas, aggregates selected trajectory-derived indicators into a structural score, and extracts a structural boundary along an ordered parameter sweep. Model-aware or benchmark-specific diagnostics are kept separate in an oracle validation layer used only for post hoc evaluation through a qualitative boundary and the lead distance between boundaries. On the FitzHugh-Nagumo benchmark, the structural boundary occurs at $\epsilon_{\text{struct}} = 0.1158$ and the oracle-defined qualitative boundary at $\epsilon_{\text{qual}} = 0.49719$, giving a lead distance of 0.38139. On an autonomous forced van der Pol benchmark, the structural boundary occurs at $A_{\text{struct}} = 1.025$ and the qualitative boundary at $A_{\text{qual}} = 1.25$, giving a lead distance of 0.225. Core robustness, supplemental scalar-delay robustness, specificity, ablation, integrity, and contract checks all pass in the frozen paper suite. The main result is simple: on both tested systems, the structural boundary occurs before the oracle-defined qualitative boundary.

1. Introduction

Deterministic physical dynamical systems can exhibit parameter-dependent changes in oscillatory organization, attractor structure, and long-time qualitative behavior [1,2]. Standard analyses of such changes are typically model-aware: they rely on linearization, continuation, bifurcation analysis, or other diagnostics that require direct access to governing equations or auxiliary dynamical structure [1,2]. Those tools are indispensable when available. They do not, however, answer a narrower question that arises naturally in data-centered settings: what regime-relevant structure can be identified when the detector is constrained to trajectories alone [3]?

This paper addresses that question for deterministic systems observed through multivariate/full-state trajectories and introduces a minimal, non-ML framework whose central object is a trajectory-only indicator atlas computed directly from trajectory data. The atlas is aggregated into a structural score and used to extract a structural boundary along an ordered parameter sweep. Model-aware or benchmark-specific diagnostics enter only afterward, through a separate oracle validation layer that defines qualitative boundaries for evaluation rather than detection. This separation matters because the detector is not allowed to borrow the validation signal it is later compared against. Without that separation, agreement between a trajectory-derived boundary and a qualitative transition would be hard to interpret.

The claim of the paper is narrow by design. It does not present a new Lyapunov method, a generic bifurcation detector, or a learned transition classifier. Instead, this paper uses the term structural regime operationally to denote a region of parameter space in which the trajectory-only atlas exhibits a coherent organization under a fixed analysis pipeline, and asks whether changes in that organization can be detected before an independently defined qualitative boundary. A structural boundary extracted from trajectories is therefore not presumed to coincide with a canonical bifurcation point unless that agreement is separately established.

On the two benchmarks studied here (FitzHugh-Nagumo and an autonomous forced van der Pol system), the trajectory-only framework identifies structural boundaries that precede the corresponding oracle-defined qualitative

boundaries. This study measures that separation by a lead distance in control-parameter space, positive when the structural boundary occurs earlier along the sweep than the qualitative boundary. Across both benchmarks, the lead distance is positive. Robustness, specificity, ablation, and reproducibility/integrity checks support the reported behavior, while scalar-observation and delay-embedding results are treated only as supplemental robustness rather than part of the paper's main claim [4,5].

The remainder of the paper is organized as follows. Section 2 formalizes the task and defines structural regime, structural boundary, qualitative boundary, and lead distance. Section 3 describes the trajectory-only core, the separated oracle validation layer, and the associated extraction and validation procedures. Section 4 presents results for the FitzHugh–Nagumo and autonomous forced van der Pol benchmarks together with nuisance, robustness, specificity, and ablation analyses. Sections 5 and 6 discuss limitations and conclude.

Problem statement: given only ordered multivariate trajectories across a parameter sweep, can a trajectory-only detector extract a stable structural boundary that precedes an independently defined qualitative boundary?

Scope of claim: this study does not propose a generic bifurcation detector or a new stability theory. It evaluates a fixed trajectory-only atlas, with oracle diagnostics reserved for post hoc validation, on two deterministic benchmark families.

2. Problem Formulation

This paper considers a family of deterministic dynamical systems parameterized by a scalar control parameter λ:

$$\dot{x} = f(x;\lambda), \qquad x(t) \in \mathbb{R}^d, \qquad \lambda \in \Lambda \subset \mathbb{R}.$$

The analysis is carried out on an ordered finite sweep

$$\Lambda_N = \{\lambda_1 < \lambda_2 < \cdots < \lambda_N\},$$

and for each λ_i a trajectory is generated,

$$X_i = \{x(t_k;\lambda_i)\}_{k=1}^{T},$$

after the prescribed transient removal and sampling procedure. In the main setting, X_i is a multivariate/full-state trajectory. A supplemental extension considers scalar observations $y_i(t) = h(x(t;\lambda_i))$ represented through delay coordinates, but this extension is not part of the main claim [4,5].

The detector receives only X_i. From each trajectory it computes a collection of trajectory-derived indicators

$$Atlas(X_i) = (I_1(X_i), \ldots, I_m(X_i)) \in \mathbb{R}^m,$$

which is called the trajectory-only indicator atlas. These indicators are aggregated into a structural score

$$S_i = \Phi(Atlas(X_i)),$$

where the aggregation rule Φ is fixed before any comparison with oracle outputs. The mapping $X_i \mapsto S_i$ constitutes the trajectory-only core.

This paper uses structural terminology operationally. A structural regime is a contiguous portion of the ordered sweep over which the indicator atlas and structural score exhibit a coherent pattern under the fixed pipeline. A structural boundary is a parameter value

$$\lambda_{struct} = B_{struct}(\{(\lambda_i, S_i)\}^N_{i=1})$$

extracted solely from the structural score across the sweep. It is intended to mark a reproducible change in trajectory organization. By construction, λ_{struct} is not defined through model-aware stability analysis, tangent-space quantities, or oracle diagnostics.

For validation only, each benchmark is equipped with a separate oracle diagnostic that produces an auxiliary output sequence Q_i from which a qualitative boundary

$$\lambda_{qual} = B_{qual}(\{(\lambda_i, Q_i)\}^N_{i=1})$$

is extracted. The qualitative boundary serves as an external reference and is not used in constructing $Atlas(X_i)$, S_i, or λ_{struct}. This paper uses the term qualitative boundary in this benchmark-specific validation sense; whether a given oracle boundary coincides with a mathematically sharp bifurcation parameter depends on the oracle and is not assumed here. This separation between trajectory-only core and oracle validation is a defining requirement of the framework.

The primary comparison quantity is the lead distance

$$\Delta_{lead} = \lambda_{qual} - \lambda_{struct},$$

with the sweep oriented so that increasing λ moves towards the oracle-defined qualitative transition. Under this convention, $\Delta_{lead} > 0$ means that the structural boundary is detected earlier along the sweep than the qualitative boundary, $\Delta_{lead} = 0$ indicates coincidence at the resolution of the sweep and extraction rule, and $\Delta_{lead} < 0$ means that the structural boundary lags the qualitative boundary.

The central question of the paper can therefore be stated narrowly: given only the trajectory family $\{X_i\}^N_{i=1}$, does the trajectory-only core extract a stable structural boundary with positive lead distance relative to an independently defined qualitative boundary on the tested deterministic benchmarks? The remaining components of the study (robustness, specificity, ablation, and reproducibility/integrity checks) constrain the interpretation of that question rather than enlarge it.

3. Methods

3.1 Systems and Data Generation

All computations in this section refer to the frozen paper-suite implementation archived with this study. Each benchmark run consisted of two separate sweeps over the same ordered parameter grid: a structural sweep, which computed only trajectory-only indicators, and an oracle sweep, which computed only benchmark-specific validation diagnostics. The

separation between trajectory-only core and oracle validation was executable rather than rhetorical: score construction was programmatically blocked unless every boundary-driving metric was tagged as trajectory-only.

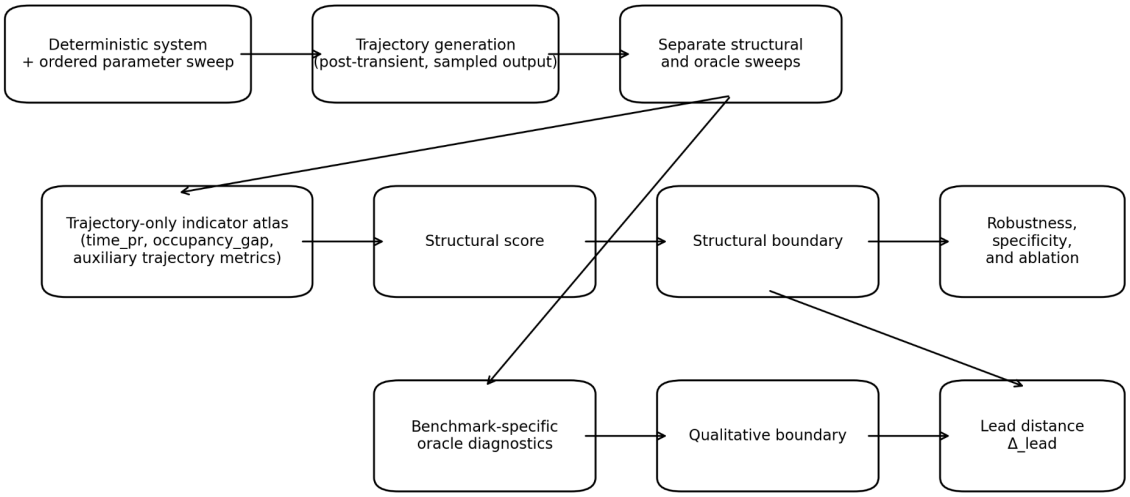

FIGURE 1: Workflow and method overview

The first benchmark is the FitzHugh-Nagumo system [6,7],

$$\dot{v} = \frac{1}{\epsilon}\left(v - \frac{v^3}{3} - w + I\right), \qquad \dot{w} = v + a - bw,$$

with fixed parameters

$$a = 0.7, \qquad b = 0.8, \qquad I = 0.5$$

for the main sweep. The structural control parameter is ϵ, interpreted here as the fast-slow separation parameter. The frozen main sweep used an explicit monotone 17-point grid in ϵ from 0.02 to 0.50, while the nuisance sweep fixed $\epsilon = 0.08$ and varied I over a 9-point grid from 0.495 to 0.505. Numerical integration used adaptive ODE solvers with uniformly sampled output after transient removal [8,9]. For the main FitzHugh-Nagumo sweep, the archived configuration used Radau with output spacing $\Delta t = 0.03$, transient length 180, post-transient window length 260, relative tolerance 10^{-7}, and absolute tolerance 10^{-9}. For the nuisance sweep, the same solver and output spacing were used with transient length 120 and post-transient window length 220. After transient removal, time was reset to zero. Initial conditions were generated from the frozen system default with small seed-controlled perturbations, and robustness replicates varied only the seed.

The second benchmark is an autonomous lift of the periodically forced van der Pol oscillator [10,11],

$$\dot{x}_1 = x_2, \qquad \dot{x}_2 = \mu(1 - x_1^2)x_2 - x_1 + A\sin\phi, \qquad \dot{\phi} = \omega.$$

The main sweep fixed

$$\mu = 8.0, \qquad \omega = 0.9$$

and varied the forcing amplitude A over an explicit 18-point monotone grid from 0.2 to 1.6. The nuisance sweep fixed $A = 0.8$ and $\mu = 8.0$ and varied ω over a 9-point grid from 0.86 to 0.94. The archived simulation settings used RK45 with output spacing $\Delta t = 0.05$, maximum internal step 0.05, transient length 80, post-transient window length 160, relative tolerance 10^{-6}, and absolute tolerance 10^{-8}. As in the FitzHugh-Nagumo runs, time was reset after transient removal, and seed-controlled default initial conditions were used.

The paper-freeze configurations used explicit parameter lists rather than generated uniform grids. This choice kept every reported structural boundary tied to a real sampled parameter value and every reported qualitative boundary tied to a fixed interpolation rule on a frozen grid. The exact sweep values, fixed parameters, solver settings, and oracle thresholds are summarized in Table 1 and archived in the configuration files.

Benchmark	FitzHugh-Nagumo	Autonomous forced van der Pol
Equations	$\dot{v} = (\frac{1}{\epsilon})(v - \frac{v^3}{3} - w + I)$, $\dot{w} = v + a - bw$	$\dot{x_1} = x_2$, $\dot{\phi} = \omega$, $\dot{x_2} = \mu(1 - x_1^2)x_2 - x_1 + A\sin\phi$
Main sweep	epsilon: 0.02 to 0.5 (17 points)	A: 0.2 to 1.6 (18 points)
Nuisance sweep	I: 0.495 to 0.505 (9 points) at epsilon=0.08	omega: 0.86 to 0.94 (9 points) at A=0.8, mu=8.0
Fixed parameters	a=0.7, b=0.8, I=0.5 (main)	mu=8.0, omega=0.9 (main)
Detector observation	(v,w) full state	(x1,x2) full state; phi excluded
Structural solver	Radau; dt=0.03; transient=180; post-transient=260	RK45; dt=0.05; max_step=0.05; transient=80; post-transient=160
Nuisance solver	Radau; dt=0.03; transient=120; post-transient=220	RK45; dt=0.05; max_step=0.05; transient=80; post-transient=160
Oracle	Tail amplitude of v; first crossing of 0.2	Stroboscopic cluster count; first crossing of 4

TABLE 1: Systems, equations, parameters, oracles, and sweep design

3.2 Trajectory-Only Indicator Atlas

For each parameter value λ_i, the structural sweep produced an observed trajectory

$$z_0, z_1, \ldots, z_M \in \mathbb{R}^{d_{obs}}$$

from the post-transient uniformly sampled output. In the main setting, $z_k = (v(t_k), w(t_k))$ for FitzHugh-Nagumo and $z_k = (x_1(t_k), x_2(t_k))$ for the autonomous forced van der Pol benchmark; the lifted phase variable ϕ was excluded from the detector and retained for oracle construction only. Supplemental scalar-observation cases first formed a coordinate observation y_k and then a delay embedding

$$Y_k = (y_k, y_{k+\tau}, \ldots, y_{k+(m-1)\tau}),$$

with fixed embedding dimension $m = 3$ and case-dependent delay τ [4,5].

The trajectory-only atlas contains two families of indicators. The first family is based on local secants of the observed trajectory. For stride $s \geq 1$, define

$$\Delta z_k = z_{k+s} - z_k, \qquad \Delta t_k = t_{k+s} - t_k, \qquad \Delta s_k = ||\Delta z_k||^2,$$

and, for nonzero segments,

$$u_k = \frac{\Delta z_k}{||\Delta z_k||^2}.$$

Given nonnegative weights w_k, define the weighted directional covariance

$$C_w = \frac{\sum_k w_k u_k u_k}{\sum_k w_k},$$

with eigenvalues $\lambda_1, \ldots, \lambda_{d_{obs}}$. The associated participation ratio is

$$PR(C_w) = \frac{(\sum_j \lambda_j)^2}{\sum_j \lambda_j^2}.$$

This study used two weightings:

$$time\_pr = PR(C_{\Delta t}), \qquad arc\_pr = PR(C_{\Delta s}),$$

and their difference

$$occupancy\_gap = arc\_pr - time\_pr.$$

The same secant data also defined two auxiliary speed-heterogeneity summaries,

$$speed\_cv = \frac{std(v_k)}{mean(v_k)}, \qquad speed\_log\_spread = log(\frac{q_{0.9}(v_k)}{q_{0.1}(v_k)}),$$

where $v_k = \Delta s_k / \Delta t_k$ and q_p denotes the p-quantile.

The second family is a short-horizon delay-space divergence diagnostic computed from a scalar observation using a Rosenstein-style construction [12]. Nearest-neighbor pairs were selected in delay space subject to a Theiler exclusion window, and the mean log-separation curve

$$D(l) = \frac{1}{|P_l|} \sum_{(i,j)\in P_l} log||Y_{i+l} - Y_{j+l}||^2$$

was regressed against a lag time over a fixed short horizon to obtain an auxiliary finite-time divergence rate. In the frozen implementation, finite-data failures of this auxiliary regression were retained as missing auxiliary outputs rather than used to halt an otherwise valid sweep.

The saved trajectory-only atlas therefore contains the scalar curves $\{time\_pr,\ arc\_pr,\ occupancy\_gap,\ speed\_cv,\ speed\_log\_spread,\ delay\_div\_rate,\ delay\_div\_quality\}$, together with the delay-divergence regression correlation. This study reports the atlas and the score separately: the atlas is the full trajectory-only descriptor set, whereas the reported structural boundary is driven only by the contract-fixed subset $\{time\_pr,\ occupancy\_gap\}$. In the frozen study contract, the boundary-driving subset was fixed to

$$\{time\_pr,\ occupancy\_gap\},$$

with equal weights and fixed orientation:

$$time\_pr \mapsto increasing, \qquad occupancy\_gap \mapsto decreasing.$$

For each boundary-driving metric $M_j(\lambda_i)$, define the oriented curve $O_j(M_j(\lambda_i))$ and the robust standardized curve

$$Z_j(\lambda_i) = \frac{O_j(M_j(\lambda_i)) - median_i O_j(M_j(\lambda_i))}{\sigma_j^{rob}},$$

where

$$\sigma_j^{rob} = 1.4826\ MAD_i(O_j(M_j(\lambda_i)))$$

with standard-deviation fallback when the median absolute deviation degenerates. Each Z_j was smoothed by a centered moving average of width 3, and the structural score was then defined by the equal-weight average

$$S(\lambda_i) = \frac{1}{2}(\tilde{Z}_{time\_pr}(\lambda_i) + \tilde{Z}_{occupancy\_gap}(\lambda_i)).$$

This deliberately simple score construction serves two purposes. First, it keeps the detector interpretable and non-ML. Second, it makes the trajectory-only claim executable. Every saved metric carries source metadata, and the implementation validates before score construction that every boundary-driving input has source class `trajectory_only`. Oracle metrics, equation-based stability quantities, and any other non-trajectory inputs are therefore excluded from the structural score by construction rather than by convention alone.

3.3 Oracle Validation Layer

Oracle validation was computed in a separate sweep on the same parameter grid and saved to a separate artifact bundle. In this study, oracle validation means benchmark-specific diagnostics that are reserved for post hoc evaluation of the trajectory-only core. Some of these diagnostics are directly model-aware, whereas others are benchmark-specific summary rules applied to the full simulated state. In all cases, they are excluded from the structural score and from structural boundary extraction.

For FitzHugh-Nagumo, the primary oracle was a tail-amplitude summary of the v-component. Let T_{tail} denote the final quarter of the post-transient trajectory. The oracle amplitude was defined as

$$A_{tail}(\epsilon) = \max_{t \in T_{tail}} v(t;\epsilon) - \min_{t \in T_{tail}} v(t;\epsilon).$$

The qualitative oracle label was

$$oscillation \quad if\ A_{tail}(\epsilon) \geq 0.2, \qquad fixed\_point \quad otherwise.$$

The primary qualitative boundary for this benchmark was the first threshold crossing of A_{tail} through 0.2 along the ordered ϵ-sweep, obtained by linear interpolation between adjacent sweep samples.

An auxiliary equation-based reference was also computed for FitzHugh-Nagumo. For equilibrium (v_*, w_*), the Jacobian is

$$J(v_*, w_*; \epsilon) = \begin{pmatrix} (1-v_*)/\epsilon & -1/\epsilon \\ 1 & -b \end{pmatrix}.$$

The frozen implementation reported the smallest positive candidate

$$\epsilon_H = \frac{1-v_*^2}{b}$$

for which $\det J(v_*, w_*; \epsilon_H) > 0$, corresponding to the trace-zero condition used as an auxiliary Hopf reference [2]. This quantity was not the primary qualitative boundary and did not enter any structural score computation.

For the autonomous forced van der Pol benchmark, the oracle layer used the lifted phase variable ϕ to define a stroboscopic section over the final 70% of the post-transient trajectory. One point in the physical phase plane (x_1, x_2) was sampled at each 2π-wrap of ϕ, yielding a stroboscopic point cloud $P_\lambda \subset \Re^2$. From this cloud, the frozen oracle computed: (i) a cluster count obtained by connectivity under a distance tolerance equal to 0.05 times the tail phase-plane scale, clipped below by 10^{-3}; (ii) a normalized spread, defined as stroboscopic RMS spread divided by the RMS tail radius of the physical trajectory; and (iii) a peak-period coefficient of variation computed from the x_1-signal over the same tail window. The saved oracle label was

$$complex\_response$$

whenever any of the following held:

$$cluster\ count > 4, \qquad period\ CV > 0.03, \qquad normalized\ spread > 0.75.$$

The primary qualitative boundary for this benchmark was the first threshold crossing of the stroboscopic cluster-count curve through the frozen threshold value 4, again obtained by linear interpolation between adjacent sweep samples.

Two points are important for interpretation. First, the oracle layer is benchmark-specific and does not define the meaning of the structural score. Second, the structural detector and the oracle layer are related only by downstream comparison. Agreement or separation between their extracted boundaries is therefore evaluation evidence, not a built-in consequence of the detector design.

3.4 Boundary Extraction, Robustness, Specificity, and Ablation

The structural boundary is extracted from the structural score by a single-change-point detector based on two-segment linear regression. Given sweep samples $(\lambda_i, S_i)_{i=1}^{N}$, define

$$k^* = \arg\min_k \left[\min_{a_L, b_L} \sum_{i \le k} (S_i - a_L \lambda_i - b_L)^2 + \min_{a_R, b_R} \sum_{i > k} (S_i - a_R \lambda_i - b_R)^2 \right],$$

subject to a minimum segment size constraint. The reported structural boundary is then

$$\lambda_{struct} = \lambda_{k^*},$$

that is, the last sampled parameter in the left segment. This choice intentionally ties the structural boundary to an actual sweep value rather than to an interpolated location.

For the main runs, the frozen implementation also reports an algorithmic sensitivity band rather than a statistical confidence interval. The structural boundary was recomputed over a grid of admissible detector settings: smoothing windows $\{1, 3, 5\}$ and admissible minimum segment sizes drawn from $\{1, 2, 3, \lfloor N/6 \rfloor\}$. The minimum and maximum candidate boundaries across this scan define the reported sensitivity interval. By contrast, robustness runs with replicated sweeps used bootstrap resampling of whole sweep curves to obtain replicate-based boundary intervals.

Qualitative boundaries were extracted from oracle curves by first threshold crossing with linear interpolation. If $Q(\lambda_i)$ and $Q(\lambda_{i+1})$ bracket the relevant threshold θ, the interpolated qualitative boundary is

$$\lambda_{qual} = \lambda_i + \frac{\theta - Q(\lambda_i)}{Q(\lambda_{i+1}) - Q(\lambda_i)} (\lambda_{i+1} - \lambda_i).$$

Lead distance was then computed as

$$\Delta_{lead} = \lambda_{qual} - \lambda_{struct}.$$

Under the sweep orientation used in this study, positive lead distance means that the trajectory-only structural boundary occurs earlier along the sweep than the oracle-defined qualitative boundary.

Robustness was organized into three tiers. Core robustness covered replicate initial conditions, relative full-state observation noise, coarser output sampling, shorter post-transient windows, and solver cross-checks. Supplemental robustness covered scalar-coordinate delay embeddings with and without added observation noise. Stress robustness covered a higher-noise full-state perturbation. Each case generated replicated sweep curves, from which mean metric curves, mean structural score curves, and replicate-based bootstrap boundary intervals were computed. Core cases were accepted when all of the following held: absolute boundary shift from the baseline ≤ 0.08, aligned Spearman correlation of the structural score with baseline ≥ 0.75, and boundary-interval width ≤ 0.25. Supplemental cases used the same architecture but looser thresholds: boundary shift ≤ 0.12, aligned Spearman correlation of the primary component $time\_$

pr with baseline ≥ 0.60, and interval width ≤ 0.25. Stress cases used thresholds 0.20, 0.35, 0.45, respectively. For supplemental and stress cases, sign alignment of boundary-driving metrics to the baseline was permitted when this improved rank agreement; the same aligned orientation was then used for boundary extraction. This alignment rule was used only for the harder supplemental/stress settings and did not alter the main full-state detector.

Specificity was evaluated by pairing each main sweep with a benchmark-matched nuisance sweep that varied a secondary parameter while holding the benchmark family fixed. Let

$$R_j^{main} = ptp(M_j^{main}), \qquad R_j^{nuis} = ptp(M_j^{nuis})$$

denote raw peak-to-peak spans of the core structural metrics $j \in \{time\_pr,\ occupancy\_gap\}$. The nuisance sweep passed specificity if (i) the oracle layer produced no qualitative boundary, (ii) oracle labels remained constant across the nuisance sweep, and (iii)

$$\max_j \frac{R_j^{nuis}}{R_j^{main}} \leq 0.35.$$

The use of raw spans rather than standardized scores was intentional: specificity is meant to rule out small nuisance drifts masquerading as benchmark-scale structural change.

Ablation compared the full structural score with two single-metric variants, `single_time_pr` and `single_occupancy_gap`. The ablation suite was designed as a criticism-hardening check rather than as a universal optimality claim. The full atlas passed ablation when all of the following held: (i) its lead distance relative to the primary qualitative boundary was positive, (ii) its lead distance was not materially worse than the median single-metric lead distance up to a tolerance of 0.05, and (iii) its structural boundary remained within the band spanned by the single-metric boundaries, expanded by the same tolerance.

Finally, this study included explicit reproducibility, integrity, and contract checks. Every saved sweep bundle wrote results, configuration, provenance, and manifest artifacts, together with selected trajectories where requested. Manifest files recorded file sizes and SHA-256 hashes. At study level, the frozen pipeline aggregated run summaries, specificity and ablation reports, an integrity report, and a contract report. The executable study contract declared the admissible structural sources, the default boundary-driving metrics, forbidden core quantities, required benchmark families and sweep roles, and the acceptance requirements on qualitative boundaries, positive lead distance, robustness, specificity, ablation, and integrity. These checks do not enlarge the scientific claim. They do something simpler: they constrain what counts as a valid realization of the frozen evidence bundle reported by this study.

4. Results

Results are reported benchmark by benchmark. Throughout this section, structural boundaries are extracted from the trajectory-only core, whereas qualitative boundaries are extracted independently from the separated oracle validation layer.

4.1 FHN Main Result

On the FitzHugh-Nagumo main sweep, the trajectory-only atlas exhibited a coherent ordered evolution across ϵ. The time-weighted tangent participation ratio increased from 1.38496 at $\epsilon = 0.02$ to 1.88957 at $\epsilon = 0.50$, while the occupancy gap decreased from 0.11337 to -0.25002. When these two boundary-driving indicators were combined into the structural score, the resulting curve rose rapidly at small ϵ and then transitioned to a slower-growth regime. The extracted structural boundary was

$$\epsilon_{struct} = 0.1158.$$

Under the admissible smoothing and segment-size settings used for the sensitivity scan, the candidate structural boundaries remained in the interval

$$\epsilon_{struct} \in [0.1158, 0.145].$$

Oracle validation placed the corresponding qualitative boundary much later in parameter space. Using the prespecified tail-amplitude rule for the v-component, the amplitude threshold was crossed at

$$\epsilon_{qual} = 0.49719.$$

The associated lead distance was therefore

$$\Delta_{lead}^{FHN} = \epsilon_{qual} - \epsilon_{struct} = 0.49719 - 0.1158 = 0.38139.$$

The ordering $\epsilon_{struct} < \epsilon_{qual}$ is not marginal in this benchmark. At the extracted structural boundary $\epsilon = 0.1158$, the oracle tail amplitude remained 3.73831, far above the threshold value 0.2. Even near the right edge of the sweep, at $\epsilon = 0.47$, the tail amplitude was still 2.13543, and the threshold crossing occurred only after interpolation between $\epsilon = 0.47$ and $\epsilon = 0.50$. The structural boundary is therefore not behaving as a late proxy for the final loss of oscillation under this oracle. Within this benchmark, the trajectory-only atlas registers a reproducible change in trajectory organization well before the oracle-defined qualitative boundary.

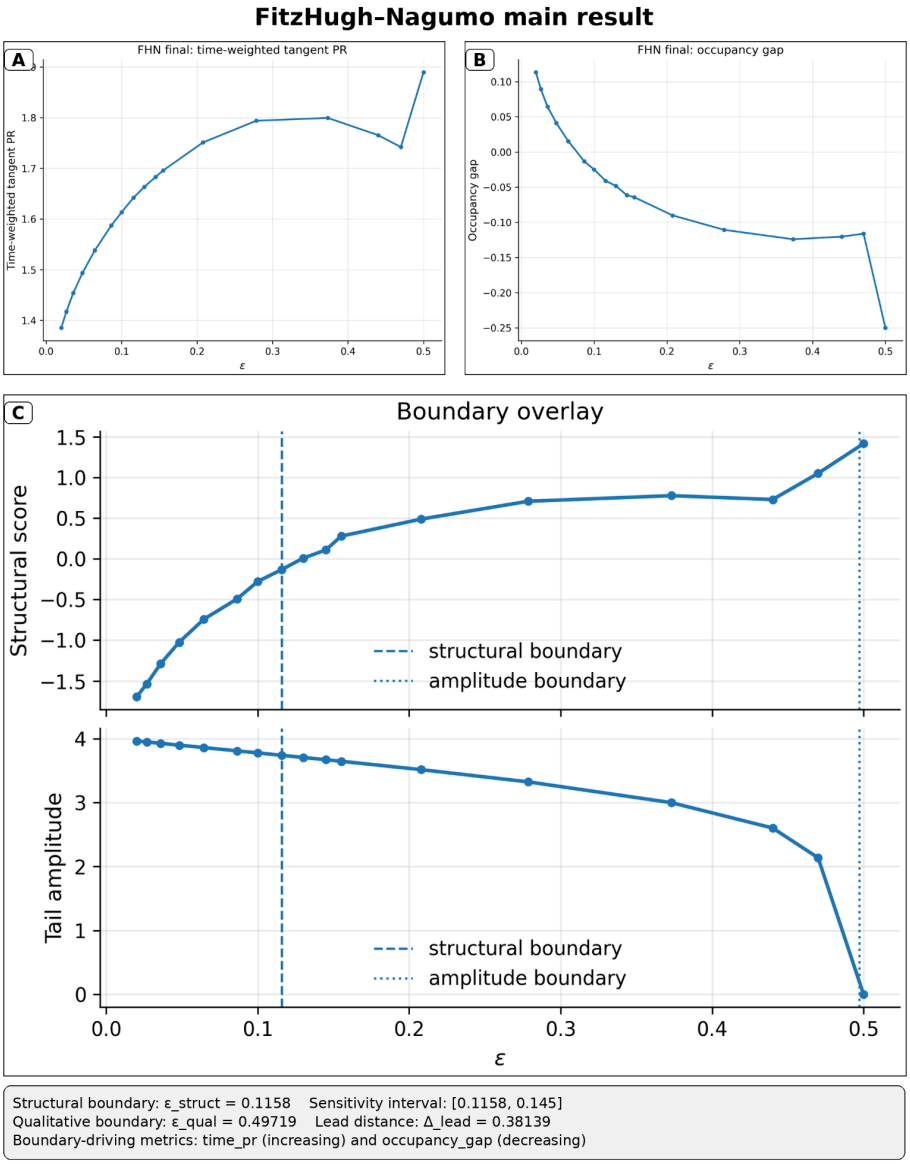

FIGURE 2: FitzHugh-Nagumo flagship result

4.2 FHN Nuisance / Specificity

Specificity was evaluated in FitzHugh-Nagumo by fixing $\epsilon = 0.08$ and sweeping the nuisance parameter I over the narrow interval $0.495 \leq I \leq 0.505$. Under the oracle validation layer, the oscillatory label remained constant across the entire nuisance sweep, and no qualitative boundary was extracted. This is the first part of the specificity result: the benchmark-specific qualitative regime did not change under the nuisance variation.

The trajectory-only metrics did vary slightly over this nuisance interval, and the change-point extractor accordingly returned a formal structural split at

$$I = 0.50125.$$

That split should not be overinterpreted. The boundary extractor is defined to return a single split on any finite ordered curve, including nuisance sweeps with small but nonzero drift. Specificity in this study is therefore assessed by oracle constancy together with the scale of nuisance variation relative to the main-sweep structural signal, not by the mere presence or absence of a nuisance change-point.

On that scale-based criterion, the nuisance sweep remained small. The raw nuisance span of $time\_pr$ was 0.0022846, compared with 0.5046106 on the main-ϵ sweep, giving a nuisance/main ratio of 0.00453. For $occupancy\_gap$, the raw nuisance span was 0.0024999, compared with 0.3633877 on the main sweep, giving a ratio of 0.00688. Both ratios are far below the prespecified tolerance of 0.35. The nuisance sweep therefore remained structurally minor relative to the benchmark-scale transition while the oracle layer stayed qualitatively unchanged.

A practical note is important here. Because the structural score is standardized within each sweep, a nuisance-score curve can appear visually steep on its own axis even when its raw underlying excursions are extremely small. For specificity, the raw metric spans are therefore the relevant quantities.

4.3 VdP Main Result

The autonomous forced van der Pol benchmark showed the same qualitative ordering. Across the forcing-amplitude sweep, the time-weighted tangent participation ratio increased from 1.34169 at $A = 0.2$ to 1.82390 at $A = 1.6$, while the occupancy gap decreased from -0.21140 to -0.69211. The corresponding structural score increased smoothly across the sweep and yielded the structural boundary

$$A_{struct} = 1.025.$$

For this benchmark, the main-run sensitivity scan was degenerate at the same value, so the reported structural sensitivity band collapsed to

$$A_{struct} \in [1.025, 1.025].$$

The prespecified oracle for this benchmark was the stroboscopic cluster-count boundary. Under that rule, the qualitative boundary occurred at

$$A_{qual} = 1.25,$$

where the cluster count first reached the threshold value 4. The resulting lead distance was

$$\Delta_{lead}^{VdP} = A_{qual} - A_{struct} = 1.25 - 1.025 = 0.225.$$

Again, the separation is clearly resolved rather than marginal. At the structural boundary $A = 1.025$, the oracle cluster count was still 10, well above the threshold that defines the primary qualitative boundary. The oracle curve remained in a many-cluster regime through $A = 1.15$, dropped to 5 at $A = 1.2$, and reached 4 only at $A = 1.25$. Thus, although the oracle complexity measures fluctuate in the pre-boundary region, the extracted qualitative boundary remains later than the structural boundary by a positive amount. Within this second benchmark, the trajectory-only atlas again identifies an earlier structural boundary than the oracle-defined qualitative boundary.

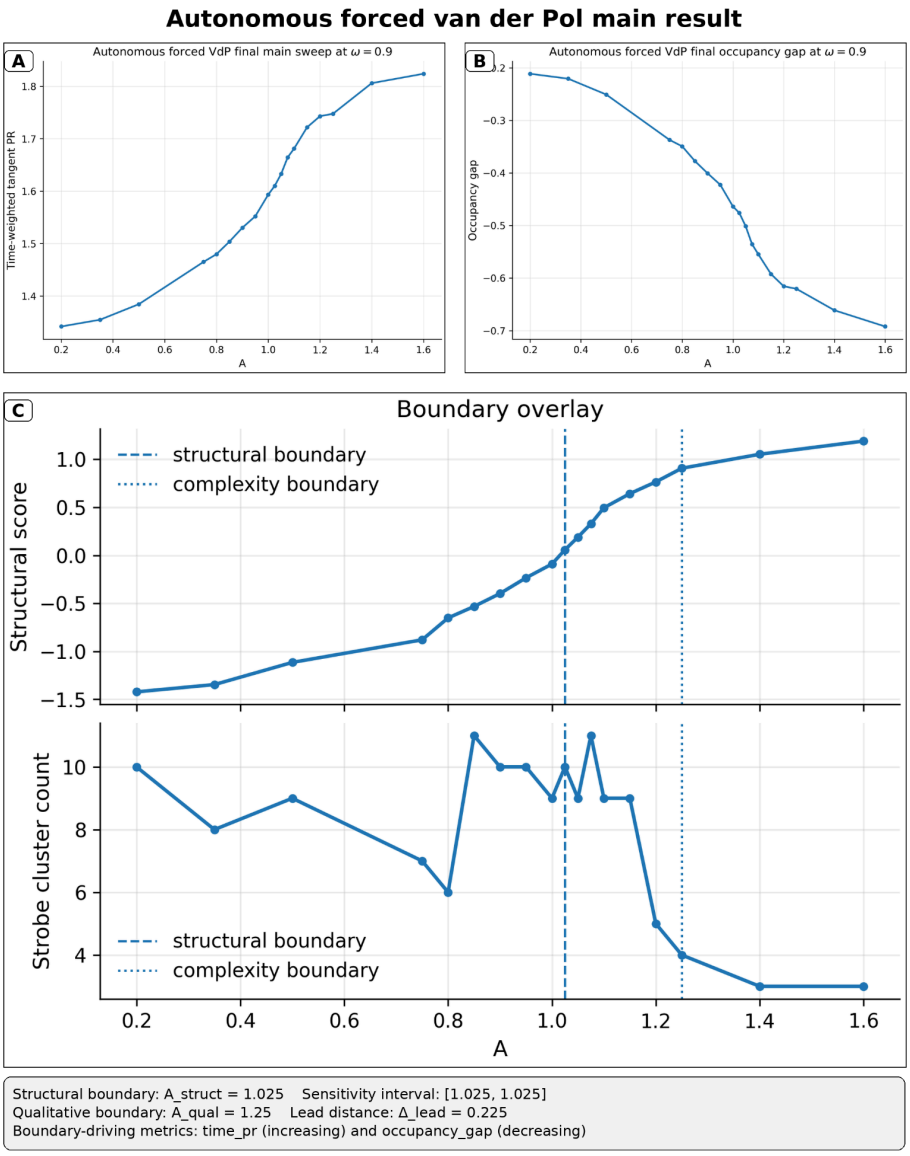

FIGURE 3: Autonomous forced van der Pol benchmark result

4.4 VdP Nuisance / Specificity

For the van der Pol nuisance analysis, the forcing amplitude was fixed at $A = 0.8$ and the forcing frequency ω was varied over $0.86 \leq \omega \leq 0.94$. As in the FitzHugh-Nagumo nuisance case, the oracle layer remained qualitatively constant across the sweep and no primary qualitative boundary was extracted.

The nuisance structural score still produced a formal split,

$$\omega = 0.9,$$

but, as above, that split is not the specificity criterion. The relevant question is whether the nuisance-induced structural excursion is small relative to the main benchmark-scale excursion while the oracle layer stays constant.

That condition was satisfied. The raw nuisance span of $time\_pr$ was 0.03399, compared with 0.48221 on the main amplitude sweep, giving a nuisance/main ratio of 0.07049. For $occupancy\_gap$, the raw nuisance span was 0.04188, compared with 0.48071 on the main sweep, giving a ratio of 0.08713. Both ratios remain well below the prespecified 0.35.

The nuisance sweep therefore stayed structurally limited relative to the main amplitude sweep, while the oracle layer did not register a qualitative transition.

As in the FitzHugh-Nagumo nuisance case, the standalone nuisance structural-score curve should be interpreted in light of within-sweep standardization. The specificity result rests on the raw-span ratios and oracle constancy, not on visual comparison of score amplitudes across separate sweeps.

4.5 Robustness and Ablation Summary

Across both main benchmarks, all prespecified core robustness cases satisfied their acceptance criteria. In FitzHugh-Nagumo, the largest core boundary shift relative to the robustness baseline was 0.0251, and the minimum aligned Spearman correlation of the structural score with baseline across core cases was 0.9559. In the autonomous forced van der Pol benchmark, the largest core boundary shift was 0.05, and the aligned Spearman correlation of the structural score with baseline was 1.0 in every core case. Coarser sampling, shorter windows, observation noise, and solver cross-checks therefore did not overturn the basic ordering of structural boundary before qualitative boundary in either benchmark.

The supplemental robustness cases based on scalar observation and delay embedding also satisfied their prespecified criteria in both systems, but the details reinforce their supplemental status. In the van der Pol benchmark, the supplemental cases remained closely aligned with the full-state baseline. In FitzHugh-Nagumo, the scalar-delay cases required the permitted sign alignment of the $time\_pr$ channel, and the largest supplemental boundary shift relative to the robustness baseline was 0.0393. These results matter because they show that some structural information persists under harder observation conditions, but they are not the basis of the paper's central claim.

The stress full-state noise cases also satisfied their acceptance criteria. The FitzHugh-Nagumo stress case shifted the boundary by 0.0251 relative to the robustness baseline, while the van der Pol stress case shifted the boundary by 0.025 with a boundary-interval width of 0.20, still below the stress tolerance. Thus, even under the designated stress settings, the detector did not reverse the boundary ordering observed in the main sweeps.

The ablation results were benchmark-dependent in an informative way. In FitzHugh-Nagumo, the full two-metric atlas produced

$$\epsilon_{struct} = 0.1158, \qquad \Delta_{lead} = 0.38139,$$

while the single-metric variants gave

$$\epsilon_{struct} = 0.1000 \quad for\ occupancy\_gap, \qquad \epsilon_{struct} = 0.1551 \quad for\ time\_pr.$$

The full-atlas boundary therefore lay between the single-metric boundaries, and its lead distance remained competitive within the prespecified tolerance. In the autonomous forced van der Pol benchmark, the ablation outcome was simpler: the full atlas and both single-metric variants all returned

$$A_{struct} = 1.025, \qquad \Delta_{lead} = 0.225.$$

This should be interpreted conservatively. In that benchmark, the two selected trajectory-only metrics are mutually consistent rather than strongly complementary.

At study level, the prespecified specificity, ablation, integrity, and contract checks were all satisfied. Both required benchmark families were present with both required sweep roles, both main runs exhibited positive lead distance together

with accepted robustness, and the archived study-level artifacts and manifests validated successfully. These checks do not enlarge the scientific claim, but they do support the narrower statement that the reported evidence is the output of a frozen and internally validated pipeline.

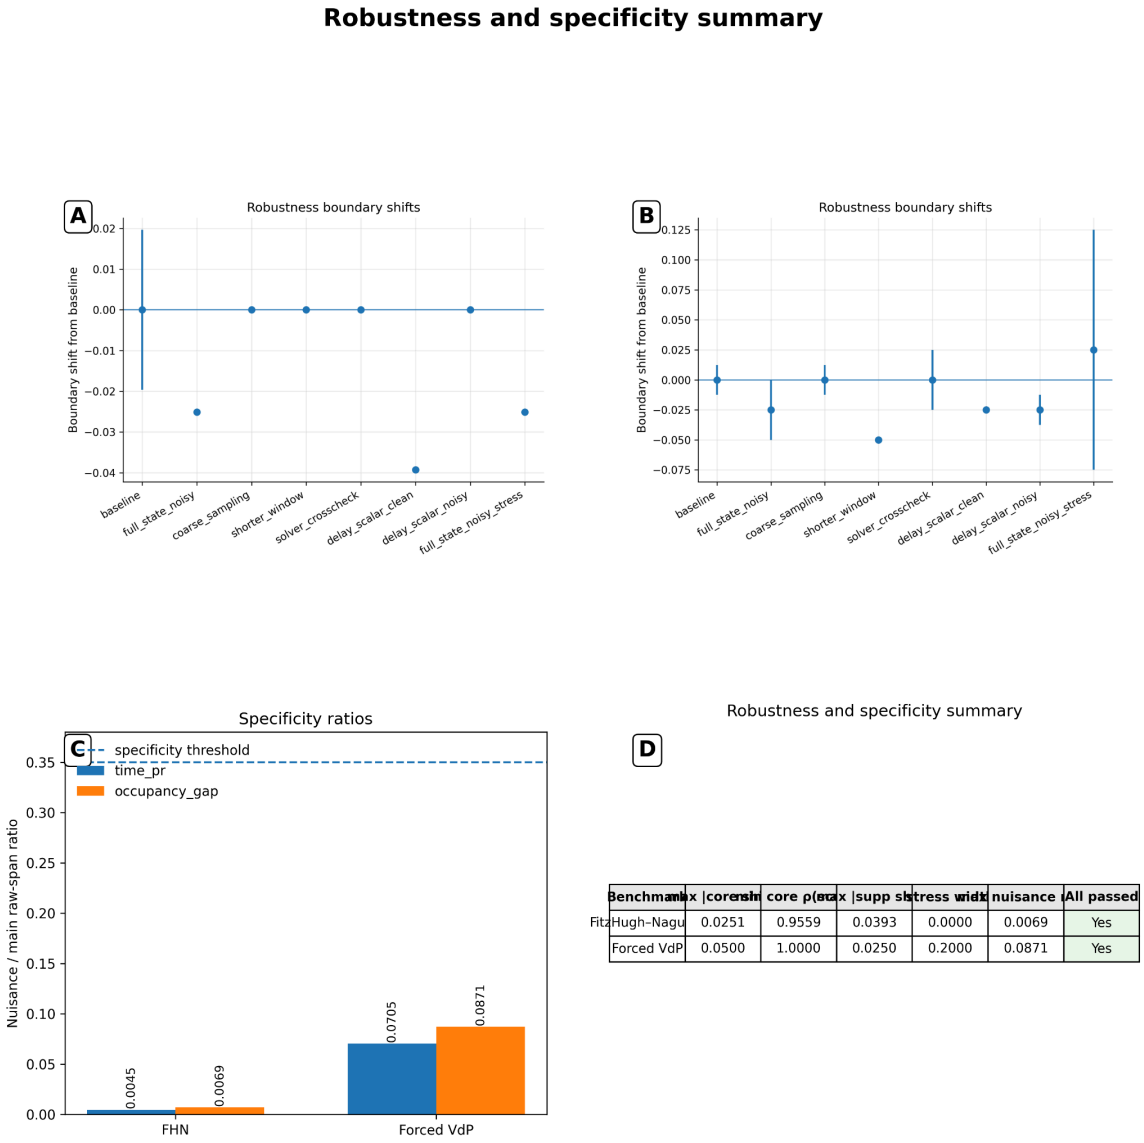

FIGURE 4: Robustness and specificity summary

Ablation summary

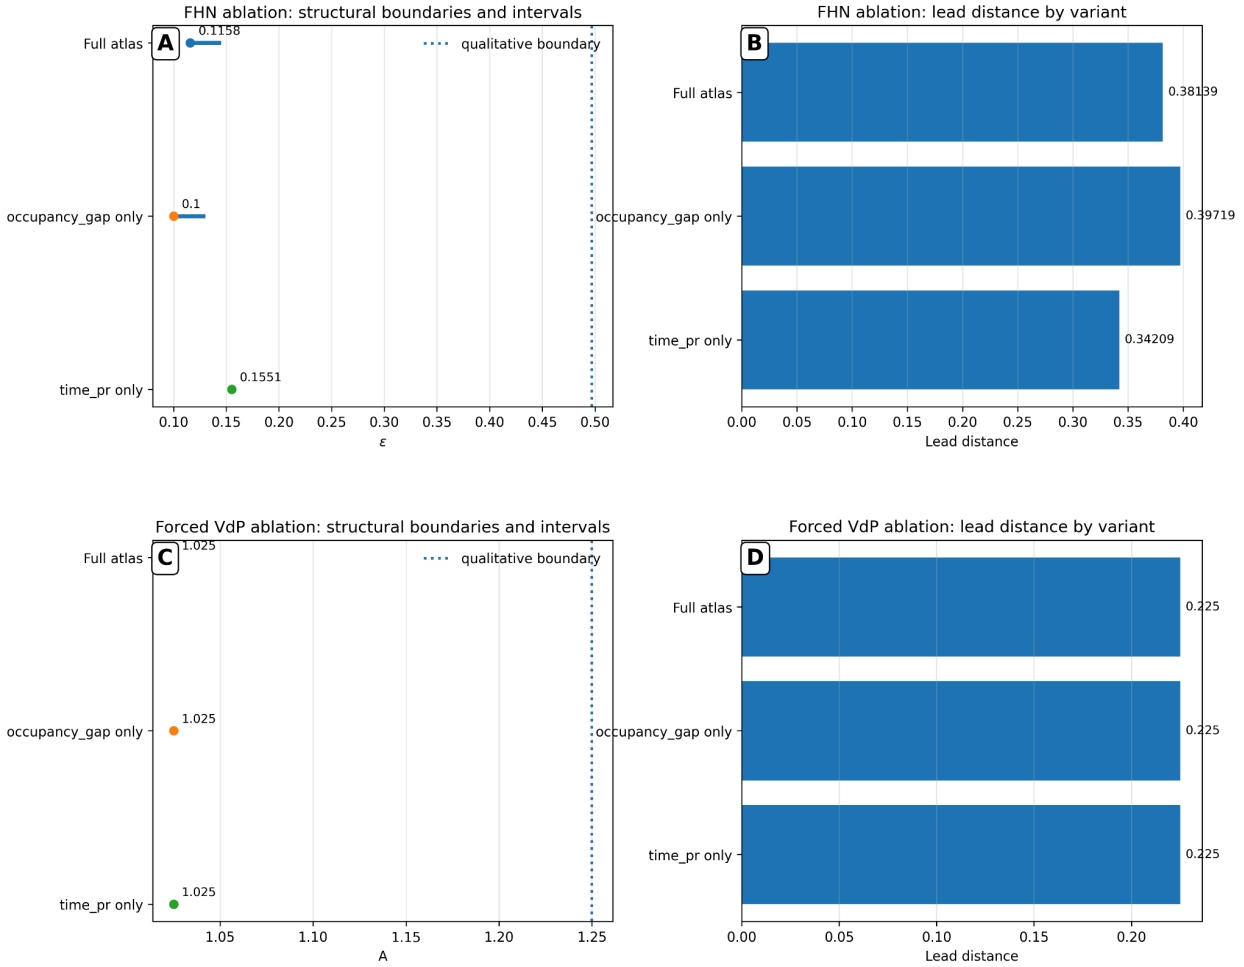

FIGURE 5: Ablation summary

Benchmark	FitzHugh-Nagumo	Autonomous forced van der Pol
Control parameter	A	epsilon
Structural boundary	1.025	0.1158
Structural sensitivity interval	[1.025, 1.025]	[0.1158, 0.145]
Qualitative boundary kind	stroboscopic cluster-count threshold crossing	tail-amplitude threshold crossing
Qualitative boundary	1.25	0.4971902598
Lead distance	0.225	0.3813902598
Core robustness	TRUE	TRUE
Supplemental robustness	TRUE	TRUE
Stress robustness	TRUE	TRUE
Specificity	TRUE	TRUE
Ablation	TRUE	TRUE
Integrity/contract	TRUE	TRUE

TABLE 2: Quantitative summary of main results and validation checks

A compact cross-benchmark summary is therefore as follows: in both tested deterministic systems, the trajectory-only core extracted a structural boundary earlier along the sweep than the oracle-defined qualitative boundary, and none of the prespecified robustness, specificity, ablation, or integrity checks overturned that ordering.

5. Discussion

What the benchmarks show: in both tested systems, the structural boundary occurs earlier along the sweep than the oracle-defined qualitative boundary.

What the lead distance means here: the trajectory-only atlas registers structural reorganization before the benchmark-specific oracle declares the qualitative transition.

What the checks do and do not show: robustness, specificity, ablation, and integrity reduce the chance that this ordering is an artifact of one unchecked pipeline, but they do not establish generality beyond the tested deterministic systems.

This study supports a restrained conclusion. On both benchmark families, the trajectory-only core extracted a structural boundary that preceded the corresponding qualitative boundary from the separated oracle validation layer. In FitzHugh-Nagumo, that separation was large on the tested sweep. In the autonomous forced van der Pol benchmark, the separation was smaller but remained clearly positive. Because the two benchmarks use different control parameters and different oracle definitions, the numerical magnitudes of the two lead distances should not be compared as if they were on a common scale. The more stable cross-benchmark statement is the ordering itself: in both tested systems, the trajectory-only structural boundary occurred earlier along the sweep than the oracle-defined qualitative boundary.

A central interpretive point is the distinction between structural and qualitative change. In this paper, a structural boundary is defined operationally by the behavior of a fixed trajectory-only atlas and its derived structural score. It is therefore not, by construction, a numerical estimate of a canonical bifurcation parameter, nor a substitute for model-aware local-stability or continuation analysis [2]. The qualitative boundary plays a different role. It is benchmark-specific, oracle-defined, and used only as an external reference. Positive lead distance therefore has a precise interpretation in this paper: the trajectory-only atlas registers structural reorganization before the separated oracle declares the benchmark-specific qualitative transition.

That distinction is also why the detector/oracle separation matters. If model-aware or oracle quantities were allowed to enter the detector, agreement or separation between the two boundaries would be difficult to interpret. Here, the structural boundary is computed from trajectory-only quantities, while the qualitative boundary is computed separately and compared only afterward. This does not make the detector universally valid, but it does make the comparison meaningful. The observed lead distances are not an artifact of feeding oracle information back into the structural score.

The robustness, specificity, ablation, and integrity checks help narrow the interpretation further. Robustness shows that the main ordering is not overturned by the finite-data perturbations, solver cross-checks, and observation changes designated in the frozen study contract. Specificity shows that small nuisance sweeps do not reproduce the benchmark-scale structural excursion while the oracle layer remains qualitatively constant. Ablation shows that the reported boundary is not merely a fragile artifact of an arbitrary score construction. Integrity and contract checks ensure that the reported evidence matches the declared pipeline. None of these checks proves generality. They do something simpler: they reduce the chance that the reported ordering is an artifact of one unchecked pipeline.

The two benchmarks also illustrate that the same qualitative conclusion can arise through somewhat different internal score behavior. In FitzHugh-Nagumo, the single-metric ablations produced distinct structural boundaries, and the full-atlas

boundary lay between them while retaining a large positive lead distance. This suggests that, in that benchmark, the selected trajectory-only indicators carry complementary information. In the autonomous forced van der Pol system, by contrast, the full atlas and the two single-metric variants yielded the same structural boundary. That outcome is easier to interpret: the chosen indicators appear to track the same structural change in that benchmark. The contrast between the two systems matters because it shows that the proposed atlas need not behave identically across benchmarks in order to support the same main result.

The scalar-observation and delay-embedding results should remain secondary in the interpretation of this paper. Those supplemental runs indicate that some structural information survives under harder observation models [4,5]. However, the main claim of this paper is not scalar detection. The main claim is restricted to multivariate/full-state trajectories. That restriction is important because scalar reconstruction introduces additional choices, like observation function, delay, embedding dimension, and noise sensitivity, that are not central to the full-state trajectory-only result reported here.

Several limitations therefore bound the present claims. First, the evidence is restricted to two deterministic ODE benchmark families. Second, the structural regime is defined operationally through the selected atlas, standardization rule, smoothing rule, and boundary extractor; a different trajectory-only atlas could produce a different structural boundary. Third, the reported boundaries are finite-grid objects: the structural boundary is tied to sampled parameter values, while the qualitative boundary depends on the chosen oracle thresholds and interpolation rules. Fourth, the main study treats ordered one-parameter sweeps rather than unconstrained observational time series. Fifth, the qualitative boundaries themselves are benchmark-specific validation rather than universal transition definitions. These limitations do not weaken the reported evidence, but they do fix its scope.

The phrase structural degradation should also be understood in that operational sense. This paper does not prove that a unique reduced model ceases to be valid at the extracted structural boundary. Rather, it shows that the chosen trajectory-only atlas detects a reproducible loss of the earlier trajectory organization before the final oracle-defined qualitative transition. That is the appropriate level of claim for the present evidence. It is compatible with later theoretical work on reduced descriptions, but it does not depend on such a theory.

The next steps are straightforward. One is to test the same detector/oracle separation principle on a broader deterministic benchmark set, especially systems with more ambiguous or multiple qualitative transitions. A second is to study alternative trajectory-only atlases and boundary extractors while preserving the same separation between core detector and oracle validation. A third is to analyze more systematically why the two selected boundary-driving indicators move early in the benchmarks studied here. A fourth is to expand the partial-observation setting beyond the present scalar-delay supplement. Each of these extensions would refine the present picture, but none is needed for the narrower conclusion already supported by the frozen study suite.

6. Conclusion

This study introduced a minimal, non-ML framework for trajectory-only structural regime detection in deterministic dynamical systems and evaluated it under a frozen, oracle-separated benchmark suite. The main result is simple: on both tested systems, the structural boundary occurred before the oracle-defined qualitative boundary. In FitzHugh-Nagumo, $\epsilon_{\mathrm{struct}} = 0.1158$, $\epsilon_{\mathrm{qual}} = 0.49719$, and $\Delta_{\mathrm{lead}} = 0.38139$. In the autonomous forced van der Pol benchmark, $A_{\mathrm{struct}} = 1.025$, $A_{\mathrm{qual}} = 1.25$, and $\Delta_{\mathrm{lead}} = 0.225$.

The contribution of this study is narrow but concrete. It does not propose a new Lyapunov method, a generic bifurcation detector, or a learned transition classifier. It shows, instead, that a fixed trajectory-only atlas can identify reproducible structural boundaries from multivariate/full-state trajectories in the tested deterministic systems, and that those boundaries can precede independently defined qualitative boundaries. Core robustness, supplemental robustness, specificity, ablation, integrity, and contract checks all passed in the frozen paper suite. Within the tested scope, the evidence supports the

central claim of this paper: trajectory-only structural boundaries can precede independently defined qualitative boundaries in deterministic dynamical systems.

Code and supplementary notes availability. The code for this study is archived separately under its existing Zenodo release with DOI: https://doi.org/10.5281/zenodo.19465798. Supplementary notes, including extended robustness tables, exact sweep grids, and archived artifact details, are available in the project repository at: https://github.com/guptaaryanr/RegimeAtlas.

References

[1] Strogatz, Steven H. Nonlinear Dynamics and Chaos : With Applications to Physics, Biology, Chemistry, and Engineering. Second edition., CRC Press, 2018. Studies in Nonlinearity.

[2] Kuznet͡sov, I͡U. A. Elements of Applied Bifurcation Theory. Fourth edition., Springer, 2023. Applied Mathematical Sciences ; 112.

[3] Abarbanel, H. D. I. and Institute for Nonlinear Science. Analysis of Observed Chaotic Data. Springer, 1996.

[4] Takens, Floris. “Detecting Strange Attractors in Turbulence.” Dynamical Systems and Turbulence, Warwick 1980, edited by David Rand and Lai-Sang Young, vol. 898, Lecture Notes in Mathematics, Springer, 1981, pp. 366–381.

[5] Sauer, Tim, James A. Yorke, and Martin Casdagli. “Embedology.” Journal of Statistical Physics, vol. 65, 1991, pp. 579–616.

[6] FitzHugh, Richard. “Impulses and Physiological States in Theoretical Models of Nerve Membrane.” Biophysical Journal, vol. 1, no. 6, 1961, pp. 445–66, https://doi.org/10.1016/S0006-3495(61)86902-6.

[7] Nagumo, J., et al. “An Active Pulse Transmission Line Simulating Nerve Axon.” Proceedings of the IRE, vol. 50, no. 10, 1962, pp. 2061–70, https://doi.org/10.1109/JRPROC.1962.288235.

[8] Virtanen Pauli, et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python.” Nature Methods, vol. 17, no. 3, 2020, pp. 261–72, https://doi.org/10.1038/s41592-019-0686-2.

[9] Hairer, E., et al. Solving Ordinary Differential Equations II : Stiff and Differential-Algebraic Problems. 2nd ed. 1996., Springer Berlin Heidelberg : Imprint: Springer, 1996. Springer Series in Computational Mathematics, 2198-3712 ; 14.

[10] van der Pol, Balth. “LXXXVIII. On ‘Relaxation-Oscillations.’” The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science, vol. 2, no. 11, 1926, pp. 978–92, https://doi.org/10.1080/14786442608564127.

[11] Cartwright, M. L., and J. E. Littlewood. “On Non-Linear Differential Equations of the Second Order: I. the Equation Y¨ − k(1-Y2)Y˙ + y = Bλk Cos(Λl + α), k Large.” Journal of the London Mathematical Society, vols. s1-20, no. 3, 1945, pp. 180–89, https://doi.org/10.1112/jlms/s1-20.3.180.

[12] Rosenstein, Michael T., et al. “A Practical Method for Calculating Largest Lyapunov Exponents from Small Data Sets.” Physica. D, vol. 65, no. 1, 1993, pp. 117–34, https://doi.org/10.1016/0167-2789(93)90009-P.