Tokamak Plasma Λ Diagnostics (POC v2.3): A Multi-Scale Structured Diagnostic Family for Zonal-Flow Organization, Mixing, and Energy-Conversion Bottlenecks

(Authors / affiliations omitted in POC draft)

January 18, 2026

Abstract

We propose a proof-of-concept (POC) study for early warning and mechanistic attribution of pedestal events (e.g. edge localized modes, large pedestal collapses, or comparable edge rearrangements) using a *structured diagnostic family* derived from the geometry of the $E \times B$ flow generator and from energy-conversion channels. The family consists of: (i) a derivative-free noncommutativity ladder $\{\Lambda_{\Delta t_j}\}$ based on commutators of generator samples at multiple lags, (ii) a projector-drift signal L_Π tracking robust subspace motion of the principal in-plane shear/strain direction (stable near degeneracy), (iii) an adiabatic-mixing parameter $\varepsilon_{\mathrm{ad}}$ comparing mixing/off-diagonal forcing to spectral-gap stiffness, (iv) a shear-coherence indicator q_S providing context for interpreting adiabatic persistence, (v) baseline-relative scaling deviation $\Delta\beta_\Lambda$ as a portable precursor feature (with explicit sign interpretation), (vi) scale-of-onset j_* with a well-defined convention in quiescent intervals, and (vii) energy-conversion proxies based on pressure–strain interaction (plus an optional kinetic attribution layer when distribution functions are available in simulation). We specify operational definitions, experimental embodiments (BES/probes), causal scale-dependent threshold calibration (rolling quantiles), numerical stabilization rules (floors tied to noise), and data-quality logging and gating via gradient-fit conditionedness κ (including spike detection). We make no performance claims; the goal is a reproducible pipeline that can validate or falsify predictive and interpretive value in specified regimes.

Contents

1 Scope, Motivation, and POC Deliverables

1.1 Motivation

Edge and pedestal dynamics in magnetized plasmas can exhibit abrupt transitions that are difficult to forecast with actionable lead time. Existing alarms are often:

- *too late* (trigger at/after onset),
- *too broad* (high false-positive burden),
- *too model-specific* (not robust across machines/regimes),
- or *mechanistically opaque*.

This POC tests a geometric hypothesis: prior to certain collapses, the local perpendicular-flow generator undergoes measurable *loss of approximate integrability* and/or *loss of coherent shear organization*. We seek operational signals that detect this robustly in noisy data and, when possible, connect it to physically meaningful energy-conversion channels.

1.2 Deliverables (reproducible and falsifiable)

1. A reproducible data-to-signal pipeline producing time series of all family members on a prescribed region-of-interest (ROI).
2. A pre-registered evaluation protocol: event labels, lead-time windows, alert definitions, and metrics (ROC/AUC, precision-recall, false-positive burden, and stability under ablations).
3. Cross-machine threshold calibration methodology using *causal*, *scale-dependent* baseline normalization.
4. A mechanistic attribution layer: separation of (a) geometric organization and mixing from (b) energy-conversion bottlenecks.
5. A falsification clause: explicit conditions under which the family fails to provide predictive value in tested regimes.

2 Physical and Mathematical Setup

2.1 Primary embodiment: perpendicular $E \times B$ advector

We focus on the perpendicular $E \times B$ drift as the advecting velocity:

$$\boldsymbol{v}_E \;=\; \frac{\boldsymbol{E} \times \boldsymbol{B}}{B^2}, \qquad \boldsymbol{E} \;=\; -\nabla\phi \quad \text{(electrostatic embodiment)}. \tag{1}$$

This choice does *not* assert $\boldsymbol{v}_E$ is the full plasma flow. It is a minimal experimentally accessible advector in many edge settings. Variants are listed in Section 11.

2.2 Local perpendicular generator and decomposition

Let $\nabla_\perp$ denote derivatives in a local perpendicular plane (approximated as (x, y) in a flux-tube or slab coordinate patch). Define the in-plane flow-gradient generator:

$$\boldsymbol{A} \;:=\; \nabla_\perp \boldsymbol{v}_E \in \mathbb{R}^{2\times 2}, \qquad \boldsymbol{S} \;:=\; \frac{1}{2}(\boldsymbol{A} + \boldsymbol{A}^\top), \qquad \boldsymbol{\Omega} \;:=\; \frac{1}{2}(\boldsymbol{A} - \boldsymbol{A}^\top). \tag{2}$$

Here $\boldsymbol{S}$ is the in-plane symmetric shear/strain-rate tensor and $\boldsymbol{\Omega}$ is the in-plane rotation tensor.

2.3 $E \times B$-advective material derivative

Define the $E \times B$-advective material derivative:

$$D_t \;\equiv\; \partial_t + \boldsymbol{v}_E \cdot \nabla_\perp. \tag{3}$$

This is a modeling/diagnostic choice appropriate when diagnosing structures transported primarily by $\boldsymbol{v}_E$. Alternative derivatives are discussed in Section 11.

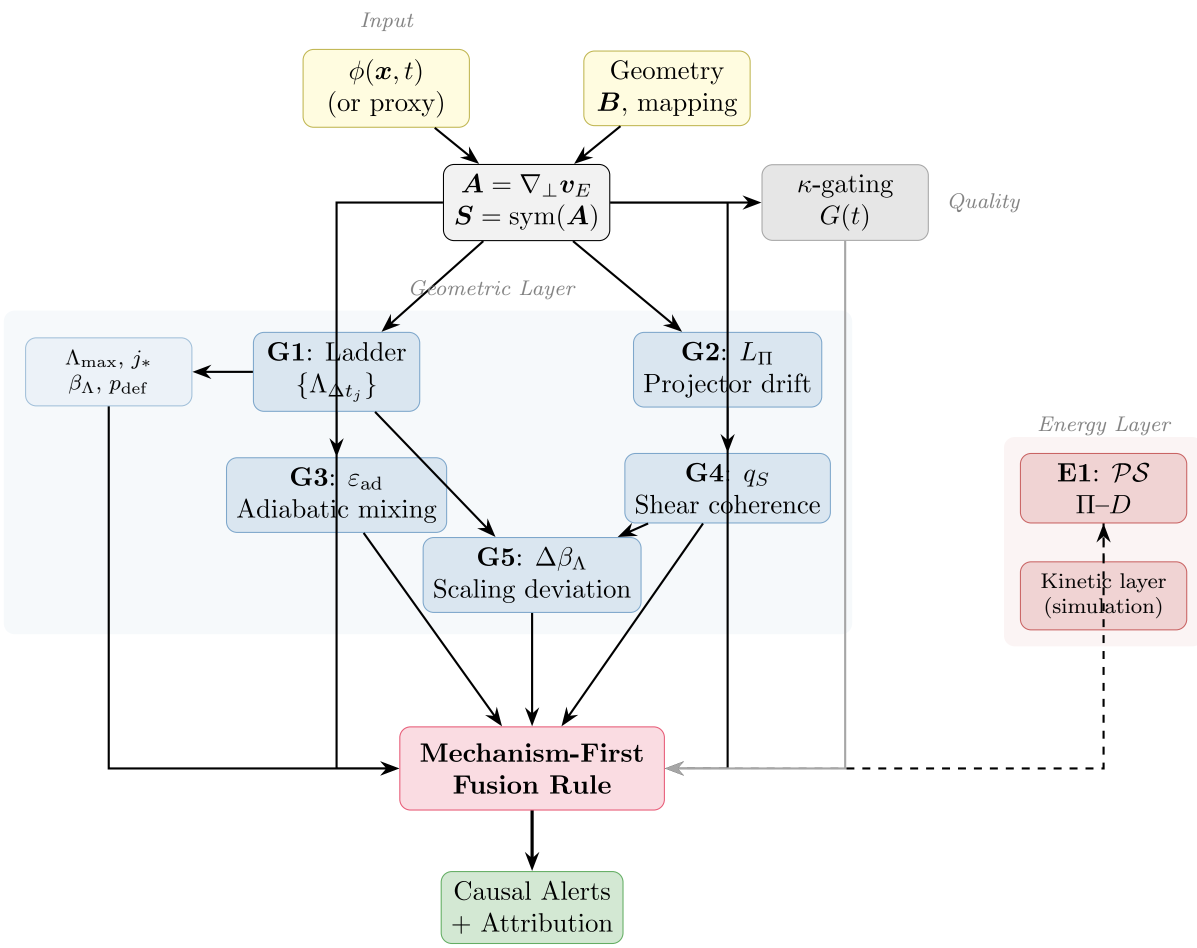

Figure 1: **Diagnostic family architecture.** Input fields (potential ϕ or proxy, geometry) feed into the generator computation. The geometric layer (G1–G5, blue) extracts structural organization signals; the energy layer (E1, red) provides conversion diagnostics (dashed arrows indicate simulation-only components). The κ-quality channel (gray) gates all signals. The mechanism-first fusion rule combines channels for causal alerts with mechanistic attribution.

3 Structured Diagnostic Family

3.1 Overview and design intent

We define a family of signals intended to be used jointly. Figure 1 illustrates the overall architecture and information flow of the diagnostic family.

The family components are:

- **(G1) Multi-scale noncommutativity ladder** $\{\Lambda_{\Delta t_j}\}$: derivative-free measures of failure of approximate simultaneous diagonalizability across multiple lags.
- **(G2) Projector drift** L_Π: robust principal-subspace motion (stable under eigenvector sign flips and near degeneracy).
- **(G3) Adiabatic-mixing parameter** $\varepsilon_{\mathrm{ad}}$: mixing force versus spectral-gap stiffness (configuration-dependent interpretation).
- **(G4) Shear coherence indicator** q_S: contextualizes adiabatic persistence by quantifying anisotropy/coherence.
- **(G5) Baseline-relative scaling deviation** $\Delta\beta_\Lambda$: portable feature capturing shifts in multi-scale ladder scaling (including sign structure).
- **(E1) Energy-conversion proxies**: pressure–strain interaction and deviatoric conversion ("Π–D"), with an optional kinetic attribution layer in simulation.

Each component has a distinct failure mode; agreement among components strengthens interpretability.

3.2 (G1) Multi-scale noncommutativity ladder (experimental primary)

Figure 2 illustrates the multi-scale ladder concept.

Definition 3.1 (Discrete noncommutativity surrogate)**.** Let Δt be a sampling interval and $\boldsymbol{A}(t,\boldsymbol{x})$ be as in (2). Define

$$\Lambda_{\Delta t}(t,\boldsymbol{x}) \;:=\; \left\|[\boldsymbol{A}(t,\boldsymbol{x}),\boldsymbol{A}(t-\Delta t,\boldsymbol{x})]\right\|_F\,. \tag{4}$$

Definition 3.2 (Multi-scale commutator ladder and derived features)**.** Let Δt_0 be the native sampling interval and let $\Delta t_j = 2^j\Delta t_0$ for integers $j \in \{0,\ldots,J\}$. Define

$$\Lambda_{\Delta t_j}(t,\boldsymbol{x}) \;:=\; \left\|[\boldsymbol{A}(t,\boldsymbol{x}),\boldsymbol{A}(t-\Delta t_j,\boldsymbol{x})]\right\|_F\,. \tag{5}$$

We define ladder features:

$$\Lambda_{\max}(t,\boldsymbol{x}) := \max_{0\le j\le J} \Lambda_{\Delta t_j}(t,\boldsymbol{x}), \tag{6}$$

$$\beta_\Lambda(t,\boldsymbol{x}) := \mathrm{slope}\Big(\log\Lambda_{\Delta t_j} \;\text{ vs }\; \log\Delta t_j\Big), \tag{7}$$

and a scale-of-onset statistic j_* defined below.

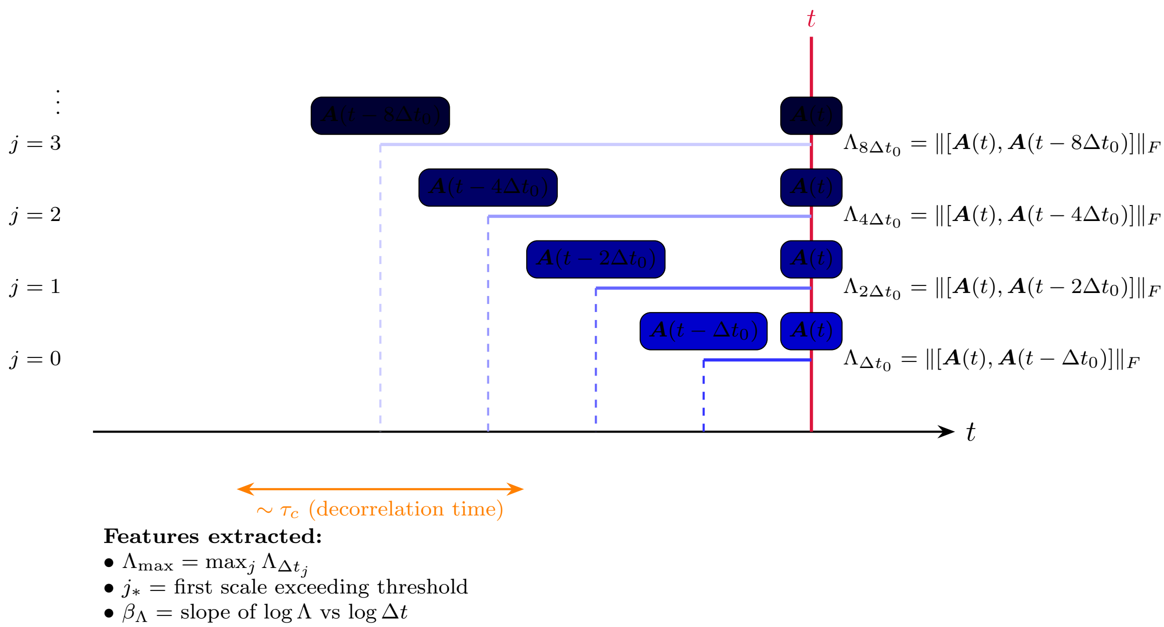

Figure 2: **Multi-scale commutator ladder concept.** At each scale j, the commutator $\Lambda_{\Delta t_j}$ measures failure of simultaneous diagonalizability between the generator $\boldsymbol{A}$ at time t and at lag $\Delta t_j = 2^j \Delta t_0$. Larger lags probe longer-timescale structural changes. The ladder yields features $\Lambda_{\max}$ (peak across scales), j_* (earliest anomalous scale), and β_Λ (scaling exponent).

Definition 3.3 (Scale-of-onset j_* with quiescent convention)**.** Fix per-scale thresholds $\{\theta_j\}_{j=0}^{J}$ (defined causally by Section 5). Define

$$j_*(t, \boldsymbol{x}) \;:=\; \min\{j \in \{0, \dots, J\} :\; \Lambda_{\Delta t_j}(t, \boldsymbol{x}) \geq \theta_j\}. \tag{8}$$

If $\Lambda_{\Delta t_j}(t, \boldsymbol{x}) < \theta_j$ for all $j \in \{0, \dots, J\}$, we set the convention

$$j_*(t, \boldsymbol{x}) := J + 1, \tag{9}$$

interpreted as "no scale exceeds threshold." Over an interval I, we track the defined/triggered fraction

$$p_{\text{def}}(I) \;:=\; \frac{1}{|I|}\, |\{t \in I :\; j_*(t, \boldsymbol{x}) \leq J\}|. \tag{10}$$

Remark 3.4 (Relation to correlation time and saturation)*.* As $\Delta t \to 0$, $\Lambda_{\Delta t}$ behaves like a derivative-free proxy for rapid generator evolution. As Δt becomes large relative to a local decorrelation time τ_c, $\boldsymbol{A}(t)$ and $\boldsymbol{A}(t - \Delta t)$ become weakly correlated, and $\Lambda_{\Delta t}$ may saturate into a "decorrelated commutator" regime. The POC includes explicit sensitivity studies of ladder features versus cadence and downsampling (Section 10).

3.3 (G2) Projector drift: robust subspace motion

Let $\boldsymbol{S}$ have eigenvalues $\lambda_1 \geq \lambda_2$ with corresponding (rank-1) spectral projectors $\Pi^{(1)}$ and $\Pi^{(\perp)} = \mathrm{Id} - \Pi^{(1)}$ when eigenvalues are distinct.

Definition 3.5 (Projector drift L_Π)**.** Define the projector drift magnitude:

$$L_\Pi(t, \boldsymbol{x}) \;:=\; \left\| \Pi^{(1)}(t, \boldsymbol{x}) - \Pi^{(1)}(t - \Delta t_0, \boldsymbol{x}) \right\|_F. \tag{11}$$

Remark 3.6 (Why projectors are the default). Eigenvectors are gauge-dependent (sign flips) and unstable near degeneracy. Projectors are gauge-invariant and admit stable generalizations when eigenvalues are clustered (Section 6.1).

Remark 3.7 (Interpretation via zonal-flow health). In tokamak edge physics, transport barriers and confinement improvement are strongly associated with persistent zonal flows and sheared $E \times B$ layers. The geometric diagnostics (G1–G4) can therefore be interpreted as quantitative measures of *zonal-flow structural health*: sustained low L_Π and high shear coherence (large q_S) indicate coherent organization, while elevated drift/mixing and reduced coherence indicate shear-layer tumbling or breakup consistent with loss of suppression and imminent collapse.

3.4 (G3) Adiabatic-mixing parameter: mixing force vs gap stiffness

Definition 3.8 (Spectral gap). Define the in-plane spectral gap of $\boldsymbol{S}$:

$$\delta(t,\boldsymbol{x}) := \lambda_1(t,\boldsymbol{x}) - \lambda_2(t,\boldsymbol{x}) \ \geq 0. \tag{12}$$

Definition 3.9 (Adiabatic-mixing parameter $\varepsilon_{\mathrm{ad}}$). Define

$$\varepsilon_{\mathrm{ad}}(t,\boldsymbol{x}) \ := \ \frac{\|(D_t\boldsymbol{S})_{\mathrm{mix}}\|_F}{\delta(t,\boldsymbol{x})^2 + \delta_0^2}, \tag{13}$$

where $\delta_0 > 0$ is a stability floor and $(D_t\boldsymbol{S})_{\mathrm{mix}}$ is the off-diagonal mixing block.

3.5 (G4) Shear coherence / anisotropy indicator

Definition 3.10 (Shear coherence indicator q_S). Define

$$q_S(t,\boldsymbol{x}) \ := \ \frac{\delta(t,\boldsymbol{x})}{\|\boldsymbol{S}(t,\boldsymbol{x})\|_F + \delta_0}. \tag{14}$$

Large q_S indicates a well-separated principal direction (coherent anisotropic shear); small q_S indicates near-isotropy where orientation is ill-conditioned.

Figure 3 provides a guide for joint interpretation of $(q_S, \varepsilon_{\mathrm{ad}})$.

3.6 (G5) Baseline-relative scaling deviation

Definition 3.11 (Baseline-relative scaling deviation). Let $\beta_\Lambda(t,\boldsymbol{x})$ be the ladder slope defined in (7). Let $\beta^{\mathrm{ROI}}_{\Lambda,\mathrm{base}}(t)$ denote a causal baseline estimate. Define

$$\Delta\beta_\Lambda(t) \ := \ \beta^{\mathrm{ROI}}_{\Lambda}(t) - \beta^{\mathrm{ROI}}_{\Lambda,\mathrm{base}}(t). \tag{15}$$

Remark 3.12 (Sign of $\Delta\beta_\Lambda$ as a mechanistic indicator). The sign of $\Delta\beta_\Lambda$ may carry interpretive content: $\Delta\beta_\Lambda > 0$ corresponds to *steepening* (structural coherence building preferentially at longer lags), whereas $\Delta\beta_\Lambda < 0$ corresponds to *flattening* (rapid decorrelation invading longer scales). Which sign correlates with collapse is an empirical question.

3.7 (E1) Energy-conversion diagnostics

Definition 3.13 (Fluid pressure–strain interaction). For a species s with bulk velocity $\boldsymbol{u}_s$ and pressure tensor $\boldsymbol{P}_s$, define:

$$\mathcal{PS}_s \ := \ -\boldsymbol{P}_s : \boldsymbol{D}_s, \qquad -\boldsymbol{P}_s : \boldsymbol{D}_s = -p_s(\nabla\cdot\boldsymbol{u}_s) - \boldsymbol{\Pi}_s : \boldsymbol{D}_s. \tag{16}$$

The second term $(-\boldsymbol{\Pi}_s : \boldsymbol{D}_s)$ is the deviatoric/shear-driven conversion ("Π–D").

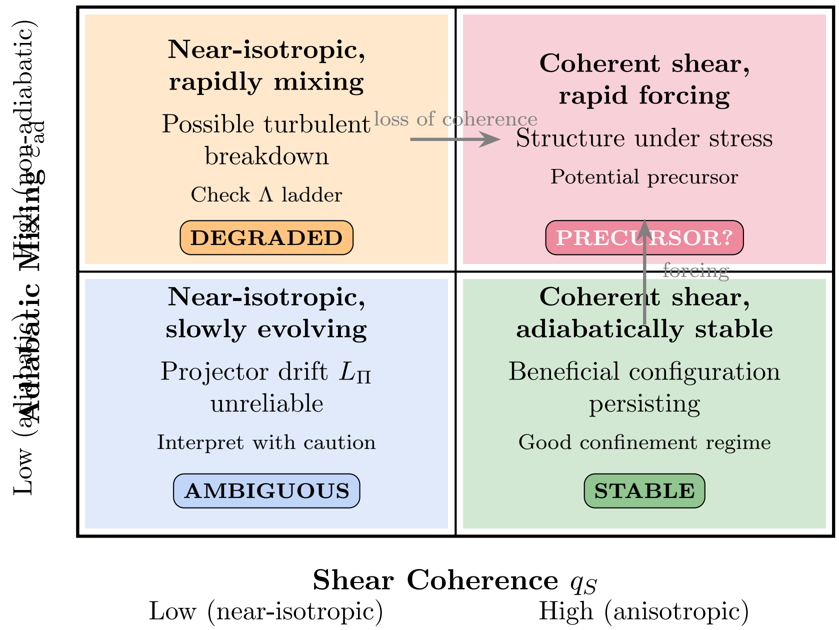

Figure 3: **Signal interpretation matrix for** $(q_S, \varepsilon_{\text{ad}})$**.** High q_S indicates well-separated principal shear direction; high ε_{ad} indicates rapid off-diagonal forcing relative to spectral gap. The upper-right quadrant (high q_S, high ε_{ad}) represents coherent structure under rapid forcing—a candidate precursor signature. Arrows indicate hypothesized transition paths toward collapse.

4 Experimental Embodiment

4.1 Design goals and quality control

The experimental embodiment includes explicit data-quality gating via the gradient-fit conditionedness $\kappa(t)$.

Definition 4.1 (κ-gating rule (level and spike))**.** Define a gated-valid indicator $G(t) \in \{0, 1\}$ by:

$$G(t) = 1 \text{ iff } \kappa(t) \le \kappa_{\max} \quad \text{and} \quad \Delta\kappa(t) := \kappa(t) - \kappa(t - \Delta t_0) \le \Delta\kappa_{\max}.$$

When $G(t) = 0$, we exclude a symmetric neighborhood $[t - w, t + w]$ in signal evaluation.

5 Causal Threshold Calibration

Definition 5.1 (Scale-dependent causal baseline normalization)**.** Let $X_j(t)$ be a ladder-scale ROI-aggregated signal and $W_{b,j}$ be a baseline window. Define

$$Q_{\alpha,j}(t) := \text{Quantile}_\alpha\big(\{X_j(s) : s \in [t - W_{b,j}, t)\}\big). \tag{17}$$

6 Numerical Stabilization

6.1 Degeneracy handling via cluster projectors

When $\delta(t, \boldsymbol{x})$ is small relative to δ_0, individual eigendirections are not reliable. We interpret L_Π and q_S cautiously and/or report them only when $\delta \ge c\,\delta_0$.

7 Simulation-First Validation Strategy

Figure 4 illustrates the three-tier validation architecture.

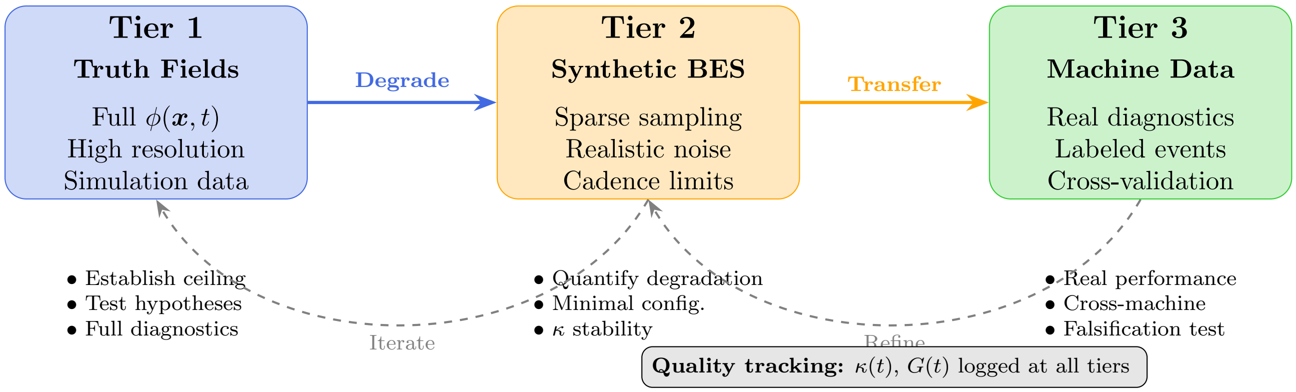

Figure 4: **Three-tier validation architecture.** Tier 1 uses full-resolution simulation truth to establish performance ceiling and test mechanistic hypotheses. Tier 2 applies synthetic diagnostic sampling (BES geometry, noise, cadence) to quantify degradation. Tier 3 validates on real machine data with labeled events. Feedback loops allow iteration based on findings at each tier.

7.1 Tier 1 candidate environments

Tier 1 requires time-resolved access to ϕ (or E_r), sufficient spatial resolution to estimate $\nabla_\perp \boldsymbol{v}_E$, and event-like episodes. Candidate environments include reduced fluid/gyrofluid edge turbulence simulations (e.g., BOUT++) and edge-capable gyrokinetic simulations.

8 Event Labeling and Evaluation Protocol

8.1 Metrics

We report: ROC/AUC and precision-recall curves, lead-time distribution, false-positive burden, stability under ablations, sign statistics for $\Delta\beta_\Lambda$, and $p_{\text{def}}^{\text{ROI}}$.

9 Mechanistic Hypotheses and Falsifiability

Assumption 9.1 (Working hypotheses to test)**.**

1. **(H1)** Multi-scale precursor: ladder features exhibit pre-event structure at specific scales.
2. **(H1b)** Sign structure: the sign of $\Delta\beta_\Lambda$ exhibits a consistent bias in pre-event windows.
3. **(H2)** Robustness: L_Π remains informative when quality gating via $G(t)$ is satisfied.
4. **(H3)** Contextual adiabaticity: ε_{ad} interpreted jointly with q_S correlates with vulnerability.
5. **(H4)** Energy bottleneck alignment: geometric signals co-occur with pressure–strain changes.

Remark 9.2 (Falsification clause). If no consistent pre-event structure emerges, or false-positive burden is prohibitive, or results are not stable under ablations, then the approach is not validated in the tested regimes.

10 Planned Experiments and Visualization

Figure 5 shows the planned visualization layout for diagnostic signals over an event.

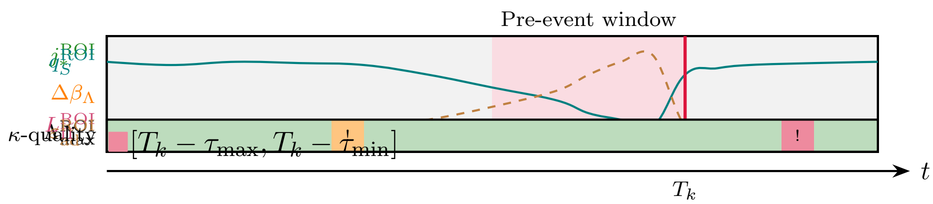

Figure 5: **Conceptual diagnostic timeline over a single event.** Panel 1: $\Lambda_{\max}^{\mathrm{ROI}}$ (solid blue) rises in pre-event window while $\Delta\beta_\Lambda$ (dashed orange) drops (flattening). Panel 2: j_*^{ROI} (green steps) decreases as finer scales trigger; L_Π^{ROI} (purple) increases. Panel 3: q_S^{ROI} (teal) drops indicating loss of coherence; $\varepsilon_{\mathrm{ad}}^{\mathrm{ROI}}$ (dashed brown) spikes. Bottom strip: κ-quality channel with flagged intervals (yellow/red). Event time T_k marked in red. *Note: This is a conceptual illustration; actual signal behavior is determined empirically.*

10.1 Ablations (required)

- Ladder sensitivity to cadence and downsampling
- L_Π and q_S sensitivity to gap thresholding
- Gradient estimation method comparison
- Proxy dependence: density-based vs direct ϕ
- ROI dependence: pedestal-only vs broader edge patches
- Baseline window sensitivity

10.2 Baselines (comparators)

Compare against amplitude-based alarms, decorrelation-time proxies, and established precursor phenomenology (magnetic precursors, inter-ELM oscillations).

11 Variants and Extensions

Extensions include electromagnetic E, alternative advectors in D_t, and 3D generalizations. An optional cross-scale coherence diagnostic (G6) tracking the correlation matrix $C_{ij} = \mathrm{Corr}(\Lambda_{\Delta t_i}, \Lambda_{\Delta t_j})$ is noted for future work.

Acknowledgments

(To be added.)

References

[1] P. H. Diamond, S.-I. Itoh, K. Itoh, and T. S. Hahm, "Zonal flows in plasma—a review," *Plasma Phys. Control. Fusion* **47**, R35 (2005).

[2] G. R. McKee *et al.*, "Experimental characterization of coherent, radially-sheared zonal flows in the DIII-D tokamak," *Phys. Plasmas* **10**, 1712 (2003).

[3] Y. Yang *et al.*, "Energy transfer, pressure tensor, and heating of kinetic plasma," *Phys. Plasmas* **24**, 072306 (2017).

[4] A. Diallo *et al.*, "Observation of edge instability limiting the pedestal growth in tokamak plasmas," *Phys. Rev. Lett.* **115**, 175001 (2015).

[5] B. D. Dudson *et al.*, "BOUT++: A framework for parallel plasma fluid simulations," *Comput. Phys. Commun.* **180**, 1467 (2009).