PROVISIONAL PATENT APPLICATION (Umbrella)

Systems and Methods for Predictive Diagnostics Using Non-Commutativity of Dynamical Generators

Ricky J. Mathews

November 8, 2025

1 CROSS-REFERENCE TO RELATED APPLICATIONS

This application relates to multiple domain-specific embodiments, including:

- (i) **Predictive Adaptive Mesh Refinement (AMR) in CFD** based on a Lagrangian curvature functional;
- (ii) **Qtrace** for robotics pose/sensor fault prediction;
- (iii) **getQore** for quantum error detection and validation; and
- (iv) **GeoSpec** for tectonic stress-regime change indication.

This umbrella application claims the **general computational framework** independent of field of use, while separate provisionals may claim domain-specific embodiments.

2 FIELD OF THE INVENTION

The invention concerns **predictive diagnostics in dynamical systems**. More particularly, it concerns computational methods and systems that forecast impending instability, fine-scale structure formation, model divergence, or error bursts by detecting **non-commutativity** between a local generator of dynamics and its time evolution.

3 BACKGROUND

Many engineered and natural systems operate near regimes where dynamics abruptly change: turbulence onsets, estimator divergence in robotics, decoder failures in quantum error correction, or stress-regime shifts in geophysics. Existing triggers are typically **reactive** and domain-specific (e.g., gradient or residual thresholds, ad-hoc detectors), offering **limited lead time** and poor portability. There is a longstanding need for a **first-principles, representation-agnostic**, and **computationally tractable** predictor that generalizes across domains and deploys in real-time.

4 SUMMARY OF THE INVENTION

The invention provides a **representation-agnostic diagnostic** built from an operator/tensor $A(t)$ that locally generates or summarizes the system's dynamics (e.g., velocity gradient, EKF Jacobian, syndrome transition operator, strain-rate tensor). Let $\dot{A}(t)$ denote a **causal** estimate of its time derivative or material derivative. Define the **non-commutativity functional**

$$\Lambda(t) = \| \, [A(t), \, \dot{A}(t)] \, \|_F, \quad [X, Y] = XY - YX$$

optionally normalized by $\|A\| \, \|\dot{A}\|$ or a local scale. Sustained increases of Λ anticipate regimes where a single smoothly evolving eigenbasis no longer explains local dynamics, a precursor to fine-scale structure, estimator blow-ups, error bursts, or hazard shifts. Thresholding or ranking Λ in **percentiles** yields **predictive alarms** that drive pre-emptive actions (e.g., mesh refinement, controller gating, decoder retuning).

Key features: (a) **first-principles**; (b) **portable** across domains; (c) **computationally efficient** with streaming implementations; (d) **lead-time** sufficient for actuation in practice.

5 BRIEF DESCRIPTION OF THE DRAWINGS

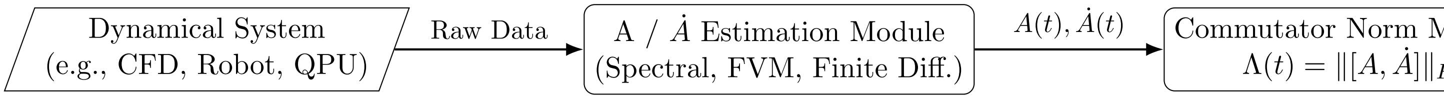

Figure 1: System architecture for the umbrella diagnostic (data source → A/$\dot{A}$ estimation → commutator norm → alarm).

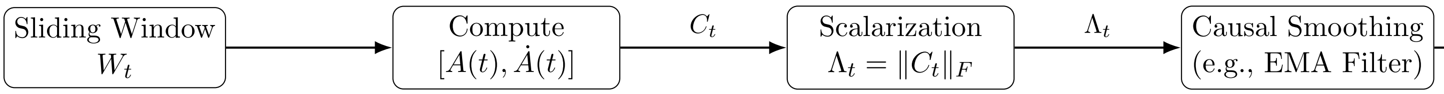

Figure 2: Sliding-window computation and causal smoothing pipeline for streaming data.

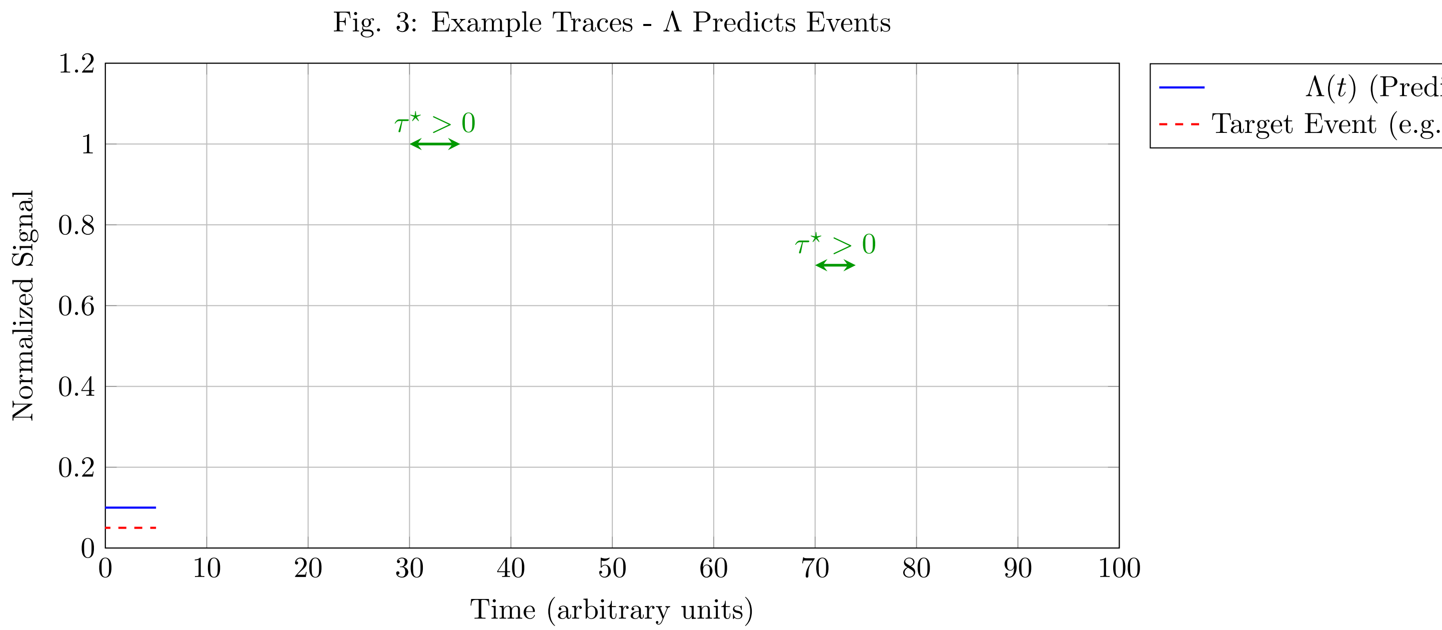

Figure 3: Example traces showing the predictive functional $\Lambda(t)$ (blue) rising with a clear lead time ($\tau^\star$) before the target event (red, dashed).

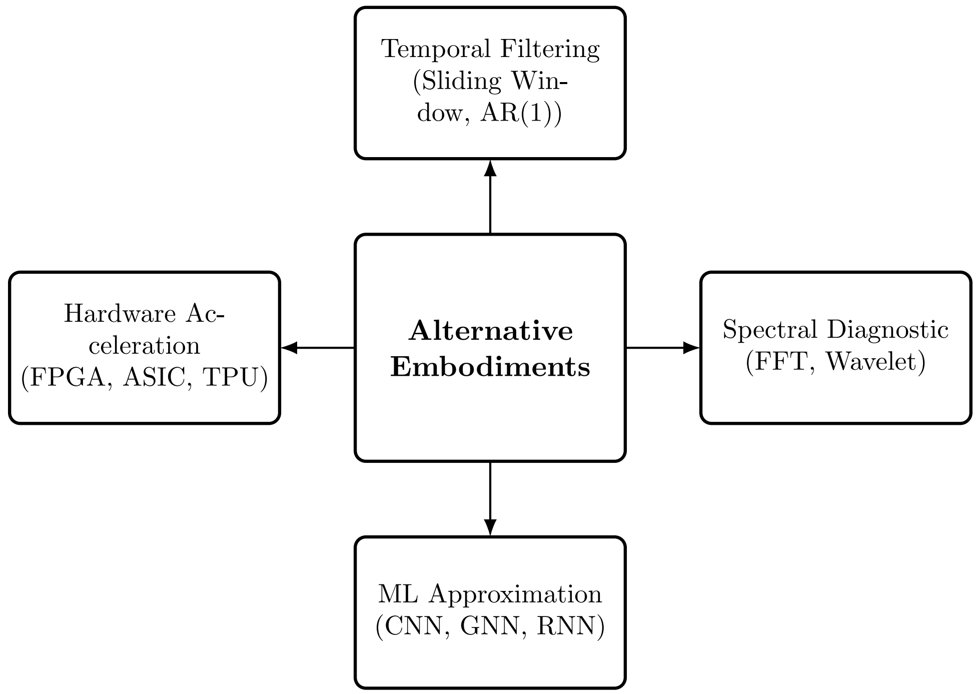

Figure 4: Alternative embodiments of the core non-commutativity diagnostic.

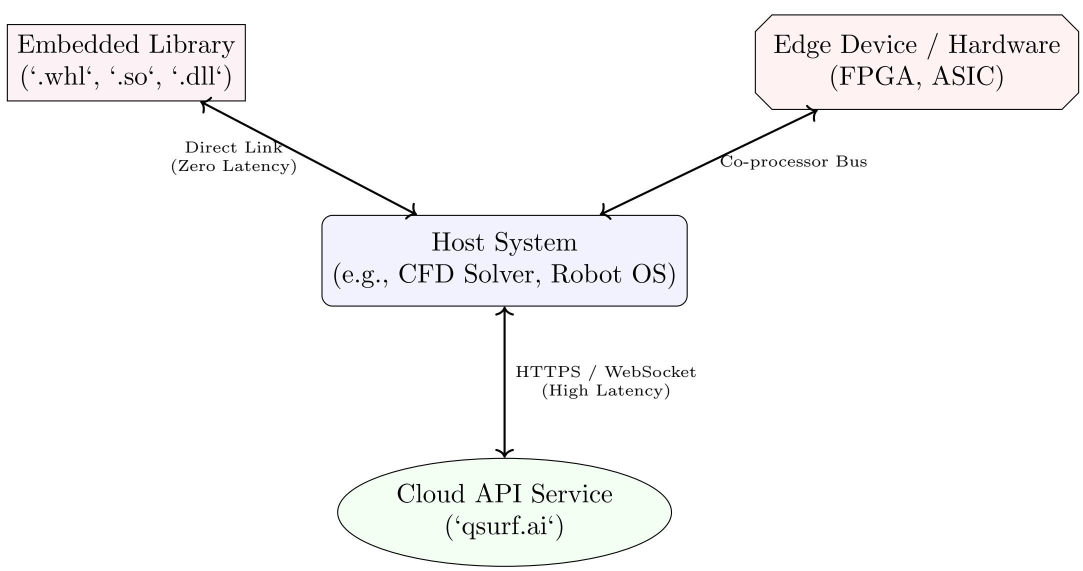

Figure 5: Deployment topologies, including embedded libraries, edge hardware, and cloud API services.

6 DETAILED DESCRIPTION OF THE INVENTION

6.1 Mathematical Framework

Let the system supply measurements from which an operator/tensor $A(t)$ is computable. Examples include, without limitation:

- **Fluid/CFD**: velocity gradient tensor $A = \nabla u$; material derivative DA/Dt.
- **Robotics**: estimator Jacobian $F_k = \partial f/\partial x$, observability Gramian, or angular-rate operator $\Omega \in \mathfrak{so}(3)$.
- **Quantum**: syndrome transition operator T_t, Pauli transfer matrix for noise channel $\mathcal{E}_t$.
- **Geophysics**: geodetic strain-rate tensor $E(t)$ from GNSS/INSAR inversions.

Estimate $\dot{A}$ **causally** by finite differences, convective/material derivative, or model-based evolution. Compute

$$\Lambda(t) = \big\|[A(t),\, \dot{A}(t)]\big\|_F, \quad \Lambda_N(t) = \frac{\Lambda(t)}{(\|A\|_F + \epsilon)(\|\dot{A}\|_F + \epsilon)}.$$

Rationale: If $[A, \dot{A}] = 0$, A and $\dot{A}$ share eigenvectors and local dynamics are integrable in a single frame. When $[A, \dot{A}] \neq 0$, the eigenframe rotates relative to deformation/transition, signaling multi-scale structure formation or model mismatch. Empirically, sustained growth in Λ **precedes** event onsets across domains.

6.2 Streaming Algorithm (Causal)

1. **Windowing**: In a sliding window W_t, estimate A and $\dot{A}$.
2. **Commutator**: Compute $C_t = [A, \dot{A}]$ with batched einsum or sparse kernels.
3. **Scalarization**: $\Lambda_t = \|C_t\|_F$ (optionally Λ_N).
4. **Smoothing**: Causal EMA or median filter to suppress shot noise.
5. **Decision**: Trigger when Λ_t exceeds **percentile-based** threshold(s) within W_t; support **multi-level** alarms.
6. **Action**: Invoke domain actuator (refine mesh; gate controller; retune decoder; escalate monitoring).

Complexity: For dense 3×3 or moderate operators, per-step cost is $O(1)$ per site/cell. For large operators (e.g., syndrome matrices), use sparse or low-rank structure; spectral fingerprints (FFT/multitaper/wavelets) maintain $O(n \log n)$ per window.

6.3 Normalization & Robustness

- Normalize by magnitudes ($\|A\|\,\|\dot{A}\|$) to improve portability.
- Gate alarms by contextual scalars (e.g., speed, SNR, gate fidelity).
- Use **percentile thresholds** learned online to avoid brittle fixed numbers.
- Provide **hysteresis** to reduce chatter.

6.4 Alternative Embodiments

- **Spectral diagnostic**: Map streams to an operator representation and compute spectral fingerprints; alarm on spectral non-stationarity; equivalently, compute commutator in frequency-domain subspaces.
- **Temporal filtering**: Sliding window, causal AR(1) residualization, or changepoint methods on Λ.
- **Hardware acceleration**: Implement commutator-norm kernels on FPGA/ASIC for real-time embedded use.
- **ML approximation**: Train regressors (CNN/GNN/RNN) to approximate Λ from raw fields, bypassing explicit $\dot{A}$ when latency matters.

6.5 Representative Domain Embodiments (Non-limiting)

- **(E1) CFD/AMR:** $A = \nabla u$, $\dot{A} = DA/Dt$; demonstrated lead time of 0.15-0.21 T_K before enstrophy spikes with $\rho > 0.7$ correlation; refine cells with highest Λ percentiles.
- **(E2) Robotics/Qtrace:** $A = F_k$ or Ω; alarm when Λ rises, predicting filter divergence, slip, or de-sync; initiate re-localization or sensor reweighting.
- **(E3) Quantum/getQore:** $A = T_t$ (syndrome transition) or $\mathcal{E}_t$ (noise channel); validated on Google Willow data achieving $R^2 = 0.9999$ correlation with error rates; rises in Λ precede logical error bursts; pre-emptively adjust decoder or scheduling.
- **(E4) Geophysics/GeoSpec:** $A = E(t)$ strain-rate field; Λ drift indicates stress-regime change; use as hazard-regime indicator (not single-event predictor).

6.6 Deployment Models

- **Embedded library** (C/C++/Python) linked into host solver or runtime.
- **Edge device** for robotics/industrial monitoring.
- **Cloud/API service** ("equation-as-a-service") that returns alarms and ranked indices.
- **Batch pipelines** for archival backtests.

6.7 Enablement & Best Mode

The best mode presently known uses (i) batched einsum commutators for 3×3 tensors and (ii) sparse/low-rank updates for large transition operators, with causal EMA smoothing and percentile thresholds. Domain gates (e.g., motion magnitude, fidelity) reduce false positives. FFT/multitaper alternatives provide robust spectral features with $O(n \log n)$ cost per window where appropriate.

7 CLAIMS (Illustrative – umbrella scope)

1. **(Core method).** A computer-implemented method for predictive diagnostics in a dynamical system, comprising: (a) computing, from observed data, a time-varying operator or tensor $A(t)$ that characterizes local dynamics; (b) computing a causal estimate of $\dot{A}(t)$; (c) computing a commutator $[A(t), \dot{A}(t)]$; (d) computing a scalar non-commutativity functional

$\Lambda(t) = \|[A, \dot{A}]\|$; (e) comparing $\Lambda(t)$ to a decision threshold to generate an alarm prior to a target event selected from fine-scale structure formation, estimator divergence, error burst, or regime shift.

2. The method of claim 1 wherein A is selected from: a velocity-gradient tensor, an estimator Jacobian or observability Gramian, a syndrome transition operator or noise-channel representation, or a geodetic strain-rate tensor.
3. The method of claim 1 further comprising normalizing Λ by $\|A\| \, \|\dot{A}\|$ or a local scale.
4. The method of claim 1 wherein the threshold is percentile-based and optionally multi-level with hysteresis.
5. The method of claim 1 wherein the threshold selection comprises: (i) computing a statistical distribution of Λ values over a sliding window; (ii) setting the alarm threshold at a percentile between the 70th and 95th percentile of said distribution; (iii) dynamically updating said threshold as the distribution evolves.
6. The method of claim 1 wherein $\dot{A}$ is a material derivative, finite difference, or model-based predictor.
7. The method of claim 1 further comprising causal smoothing of Λ using exponential or median filters.
8. The method of claim 1 wherein the alarm triggers a pre-emptive actuator chosen from: mesh refinement; controller gating or re-localization; decoder retuning or schedule changes; or escalation of monitoring.
9. The method of claim 1 implemented as a streaming algorithm with per-step computational complexity bounded by $O(1)$ per site for small tensors or $O(n \log n)$ per window for spectral embodiments.
10. A system comprising processors and memory programmed to perform any of claims 1–8, provided as an embedded library, edge runtime, cloud service, or hardware accelerator.
11. A non-transitory computer-readable medium storing instructions that, when executed, cause a machine to perform any of claims 1–8.
12. **(Spectral alternative).** The method of claim 1 wherein the diagnostic further comprises computing a spectral fingerprint of data mapped into an operator representation via FFT, wavelet, or multitaper analysis, and generating the alarm from non-stationarity in the spectral features or a commutator computed in said representation.
13. **(Domain embodiments).** The method of claim 1 wherein the target event is, respectively: (i) dissipation/enstrophy burst in CFD; (ii) estimator divergence or slip in robotics; (iii) logical error burst in QEC; or (iv) stress-regime change in geophysics.
14. A method of providing computational diagnostic services comprising: (a) receiving dynamical system data from a remote client via a network interface; (b) computing said Λ functional on server infrastructure; (c) returning predicted alarm states or ranked indices to said client; (d) billing said client based on computational resources consumed or number of predictions generated.

8 ADVANTAGES

- **First-principles** and portable across domains.
- **Predictive** alarms enabling pre-emptive action.
- **Computationally efficient** with streaming and spectral variants.
- **Robust** via normalization, gating, and percentile thresholds.

9 INDUSTRIAL APPLICABILITY

Applicable to aerospace, energy, automotive, robotics, quantum hardware, HPC/cloud simulation platforms, and national lab workflows where early warning reduces cost or risk.

10 IMPLEMENTATION NOTES

- Provide APIs that accept domain-specific data and return Λ maps and ranked indices.
- Persist artifact logs for auditability (thresholds, percentiles, latencies).
- For safety-critical loops, include fallback to reactive detectors with OR-logic.

11 FILING NOTES (Provisional)

This umbrella filing establishes priority for the general non-commutativity diagnostic independent of field of use. Domain-specific provisionals (CFD AMR, robotics Qtrace, quantum getQore) may be filed in parallel and later combined in one or more non-provisional applications claiming priority to this umbrella and the respective domain provisionals.