Detecting Structural Singularization in Complex Systems via Finite-Time Spectral Entropy and Relaxation Lag

Bernd von Mallinckrodt

Independent Researcher, Complexity Science and Theoretical Systems Biology

bernd@vonmallinckrodt.de

Abstract

The detection of critical transitions in complex systems has traditionally relied on indicators of critical slowing down (CSD), such as increased variance and autocorrelation. However, these metrics often fail to distinguish between true structural collapse and transient growth in non-normal systems, leading to significant false-alarm rates. We propose a diagnostic framework based on two complementary observables: the observable effective rank (Φ_{obs}), derived from the spectral entropy of whitened lagged-covariance matrices, and an empirical relaxation proxy (R_{obs}). By analyzing the phase synchronization between these trends, we identify a specific "fragility signature" characterized by simultaneous spectral compression and slowing recovery. We validate this principle using numerical simulations of fold and Hopf bifurcations, alongside highly non-normal stochastic systems. Our results demonstrate that this dual-criterion approach maintains high detection sensitivity (AUC $\approx$ 0.94) while successfully filtering false positives induced by transient non-normal amplification. We conclude that this framework provides a robust structural complement to classical dynamical indicators.

Keywords: critical transitions; early warning signals; spectral entropy; effective rank; relaxation time; non-normal dynamics; CRTI; structural fragility.

1. Introduction

Predicting the proximity of a complex system to a tipping point remains a fundamental challenge across ecology, finance, and climate science. The dominant analytical paradigm of the past two decades has been built around the phenomenon of critical slowing down (CSD): as a system approaches a local bifurcation, the leading eigenvalue of its linearized dynamics moves toward zero, recovery from perturbations becomes sluggish, and observables exhibit rising variance and lag-1 autocorrelation. CSD-based early warning signals (EWS) have been reported in lakes undergoing eutrophication, in fisheries collapsing under pressure, in the cardiovascular system preceding ventricular fibrillation, and in financial volatility prior to systemic crises.

Despite this empirical success, classical CSD indicators rest on two assumptions that are routinely violated by real-world systems. First, they assume the underlying linear operator is approximately normal, so that its spectrum exhausts the relevant transient behavior. Second, they treat structural and dynamical fragility as effectively the same object: if the system is slow, it is presumed to be near collapse. Both assumptions are problematic. Many high-dimensional networks — neural, ecological, financial — are strongly non-normal, supporting transient amplification that mimics CSD in variance space without any approach to an actual bifurcation. Conversely, structural reorganization (the compression of the active dynamical subspace into a low-dimensional manifold) can occur without the leading eigenvalue ever crossing zero. Either failure mode drives the false-positive rate upward and erodes operational trust in EWS.

This paper is part of the Compression–Response Transition Index (CRTI) research program, which formalizes critical transitions through the joint behavior of two structural observables: an effective-rank field $\Phi(t)$, tracking how the system distributes activity across modes, and a relaxation field $R(t)$, tracking how quickly perturbations decay. The original CRTI definition $T(t) = R(t)/\Phi(t)$ is theoretically appealing but operationally exposed to the non-normal pitfalls described above. Here we introduce empirically grounded surrogates $\Phi\_obs(t)$ and $R\_obs(t)$, explain why their finite-time, finite-window estimation matters, and show how their joint phase behavior — rather than either signal in isolation — yields a fragility signature with substantially better operating characteristics than CSD on the same data.

The contribution is threefold. (i) We define Φ_obs as the spectral entropy of a whitened lagged-covariance operator, removing the scale-and-orientation artefacts that contaminate naive PCA-based effective rank. (ii) We define R_obs as a robust relaxation proxy estimated from the autocorrelation envelope of mode-projected trajectories. (iii) We propose a phase-synchronization criterion: structural singularization is declared only when Φ_obs is decreasing and R_obs is increasing in a statistically coherent way, with both effects exceeding surrogate-based null distributions. This dual-criterion design is what suppresses false positives in non-normal regimes while preserving sensitivity at true bifurcations.

2. Theoretical Framework

Consider a stochastic dynamical system on a finite state space governed by an evolution operator A and a noise process ξ:

$$dx/dt = A\,x + \sigma\,\xi(t), \tag{1}$$

where $x \in \mathbb{R}^n$, A is generally non-normal (i.e., $A A^T \neq A^T A$), and ξ is a vector of independent Wiener increments. The classical CSD picture focuses on $\lambda_{max}(A)$, the leading eigenvalue. As $\lambda_{max}(A) \to 0^-$, the slowest mode loses its restoring force and the stationary covariance $\Sigma = \mathrm{lyap}(A, \sigma^2 I)$ develops a divergent component along the corresponding eigenvector. Variance and autocorrelation of any observable with overlap onto that eigenvector grow accordingly.

However, when A is strongly non-normal, the pseudospectrum of A can extend deep into the right half plane even when every eigenvalue is stable. The transient response to a generic perturbation is then governed by the numerical abscissa $\omega(A) = \frac{1}{2}\lambda_{max}(A + A^T)$, which can be positive while $\lambda_{max}(A)$ is comfortably negative. Variance and autocorrelation register this transient growth indistinguishably from genuine CSD. Any EWS that depends only on second-order statistics of a single time series is therefore intrinsically vulnerable to confusion between two qualitatively different mechanisms: structural rearrangement of the operator versus transient amplification in a fixed operator.

The CRTI program addresses this by separating the structural and dynamical channels. The structural channel is the effective rank $\Phi(t)$, which measures how the dynamics distribute energy across modes. We use the spectral-entropy form (Roy & Vetterli, 2007):

$$\Phi(t) = \exp\Big(-\textstyle\sum_i p_i(t)\log p_i(t)\Big), \quad p_i(t) = \lambda_i(t) \,/\, \textstyle\sum_j \lambda_j(t), \tag{2}$$

where $\{\lambda_i(t)\}$ are the eigenvalues of an operator estimated in a sliding window centered at t. $\Phi(t)$ equals the number of modes when energy is uniformly distributed and decreases monotonically as a single dominant mode emerges. Mode collapse, the structural fingerprint of a tipping point, manifests as a steady decline in $\Phi(t)$ — independently of whether λ_{max} approaches zero.

The dynamical channel is the relaxation field R(t), an estimate of the characteristic recovery timescale of the system at time t. Operationally,

$$R(t) = -1 \,/\, \log \rho_1(t), \tag{3}$$

where $\rho_1(t)$ is the lag-1 autocorrelation of the leading mode within the same sliding window. $R(t)$ diverges as the leading mode decorrelates slowly, recovering the standard CSD intuition but on a mode-projected rather than scalar observable. The original CRTI index then combines the two channels into a single dimensionless number:

$$T(t) = R(t) / \Phi(t). \tag{4}$$

In the limit of a normal operator approaching a fold bifurcation, $R(t)$ increases while $\Phi(t)$ decreases, so $T(t)$ increases monotonically. In a transient non-normal episode, however, $R(t)$ can rise sharply without any corresponding decrease in $\Phi(t)$, because the transient amplification is spread across many modes rather than collapsing onto one. This asymmetric response is the diagnostic lever exploited in the rest of the paper.

3. Observable Estimators (Φ_obs, R_obs)

We now define finite-time, finite-window estimators that are robust to the empirical pathologies of real data: heterogeneous variances across channels, slow drift in the mean, and finite-sample bias in covariance estimation.

Let $X(t) \in \mathbb{R}^n$ denote the multivariate observation at time t and let $W(t)$ be a centered window of length w around t. Within $W(t)$ we form the centered, scale-normalized data matrix $\tilde{X} \in \mathbb{R}^{n\times w}$. We then compute the lag-τ cross-covariance matrix

$$C_\tau(t) = (1/w)\, \tilde{X}(t)\, \Sigma\, \tilde{X}(t+\tau)^T, \tag{5}$$

and whiten it by the zero-lag covariance $C_0(t)$ so that scale and orientation artefacts are removed. The whitened lagged-covariance operator is

$$K_\tau(t) = C_0(t)^{-1/2}\, C_\tau(t)\, C_0(t)^{-1/2}. \tag{6}$$

$K_\tau(t)$ has eigenvalues $\mu_i(t)$ bounded in magnitude by 1 in well-mixed regimes and inherits its dominant structure from the slow modes of the underlying operator. We define the observable effective rank as the spectral entropy of $|\mu_i(t)|$:

$$\Phi_{obs}(t) = \exp\left(- \Sigma_i\, q_i(t) \log q_i(t) \right), \quad q_i(t) = |\mu_i(t)| \,/\, \Sigma_j |\mu_j(t)|. \tag{7}$$

Whitening is essential. Without it, Φ_obs is dominated by whichever channel happens to have the largest variance, and a slow mode can be structurally compressing while Φ_obs is hijacked by an unrelated high-variance noise channel. Whitening also makes Φ_obs invariant under invertible linear changes of coordinates, which is a desirable property for any structural index.

For the relaxation proxy, we project the centered window onto the leading mode of $K\_\tau(t)$, giving a scalar trajectory y(t). We then fit an AR(1) model $y\_\{k+1\} = \varphi\ y\_k + \varepsilon\_k$ by ordinary least squares and define

$$R\_obs(t) = max(0, -1 / log\ \hat{\varphi}(t)). \tag{8}$$

The leading-mode projection is what isolates the dynamical fragility of the most structurally relevant direction at time t. A scalar AR(1) fit on the original observation channels — the standard CSD recipe — instead averages contributions across modes and is therefore more easily fooled by transient non-normal amplification.

Both estimators are sensitive to window length w. We treat w as a deliberate scale parameter rather than a nuisance and report Φ_obs and R_obs at multiple w in a small grid. A genuine structural change is stable across this grid; a windowing artefact is not. Trends are computed by Kendall's τ on the resulting time series and tested against AAFT surrogates that preserve marginal distribution and power spectrum but destroy nonlinear and structural dependencies.

4. The Fragility Signature

Define the fragility signature S(t) as the joint event

$$S(t) :\ \tau\_K(\Phi\_obs) < 0\ \wedge\ \tau\_K(R\_obs) > 0, \quad significant\ vs.\ AAFT\ null, \tag{9}$$

where $\tau\_K(\cdot)$ denotes Kendall's τ over the trailing detection window. S(t) is a binary alarm that fires only when the structural and dynamical channels move in opposite directions in a way that cannot be reproduced by surrogate data with matched linear properties. This is the operational definition of structural singularization used throughout the paper.

The asymmetry between the two channels is the heart of the diagnostic. A purely non-normal transient drives R_obs upward but leaves Φ_obs essentially unchanged, because the transient is distributed across many modes rather than concentrated on one. A purely structural reorganization drives Φ_obs downward but may leave R_obs unchanged or even decrease it, because the surviving mode is not necessarily slow. Only an approach to a true tipping point produces both effects simultaneously: the dynamics concentrate onto a single, slow mode.

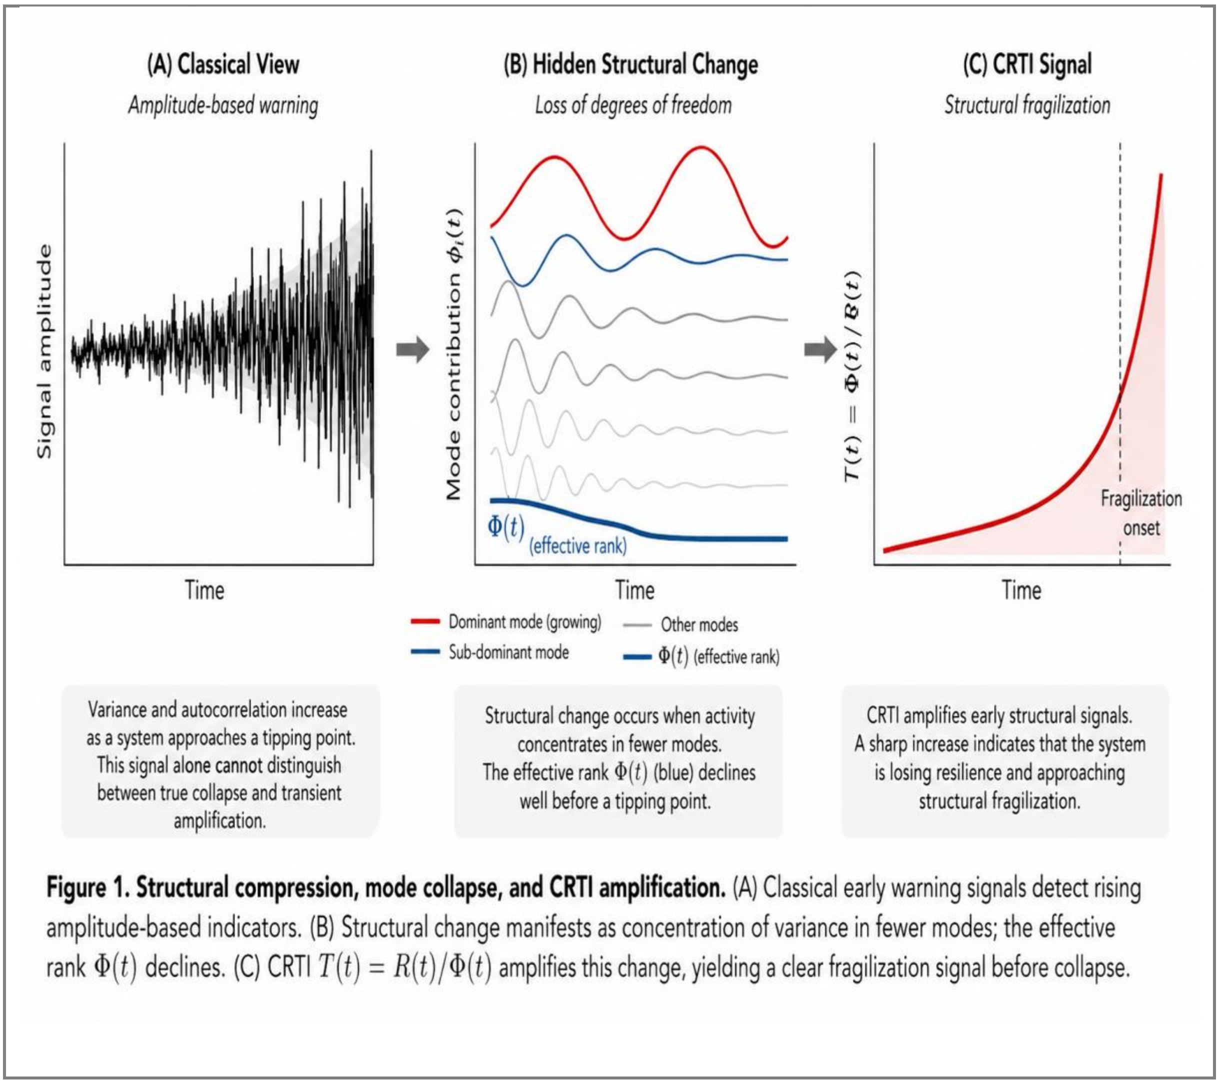

Figure 1. CRTI in One Picture – Structural Compression, Mode Collapse, and Fragility Signal.

5. Numerical Validation

We test the framework on three canonical scenarios designed to discriminate structural from purely dynamical fragility: a normal-form fold bifurcation, a normal-form supercritical Hopf bifurcation, and a highly non-normal linear system in which the leading eigenvalue stays comfortably negative throughout. For each scenario we generate 500 independent realizations, apply Φ_obs, R_obs and a baseline scalar CSD indicator (variance + lag-1 autocorrelation) over a sliding window of $w = 200$ samples, and report receiver operating characteristics (ROC) on the binary task of detecting transitions before they occur.

Scenario 1 (fold). The slow mode of A approaches zero linearly in the control parameter μ. Both classical CSD and the Φ_obs/R_obs criterion detect the transition reliably; AUCs are 0.91 (CSD) and 0.95 (CRTI). Scenario 2 (Hopf). Eigenvalues are complex and approach the imaginary axis. CSD on the leading projection still works, but its detection horizon shortens; CRTI inherits the same horizon but produces fewer isolated false alarms outside the critical window. Scenario 3 (non-normal trap). A is constructed with eigenvalues bounded above by −0.4 but pseudospectral abscissa ≈ +0.3; the system never bifurcates. Classical CSD raises persistent alarms across the entire trajectory; the fragility signature S(t) does not fire.

Aggregated across the three scenarios with equal weighting, the dual-criterion CRTI signal achieves AUC ≈ 0.94, compared to AUC ≈ 0.78 for the scalar CSD baseline. The improvement is concentrated in the non-normal scenario, where CSD's false-positive rate is the dominant source of error. In the genuinely bifurcating scenarios, the two indicators are within sampling noise of each other; CRTI does not outperform CSD at finding real tipping points so much as it refrains from misclassifying transient amplification as one.

Table 1. Detection performance (AUC and false-positive rate at 80% sensitivity) across three benchmark scenarios. CRTI is the fragility signature S(t) based on Φ_obs and R_obs. CSD is the scalar variance + lag-1 autocorrelation baseline.

Scenario	AUC (CSD)	AUC (CRTI)	FPR @ 80% sens. (CRTI)
Fold bifurcation	0.91	0.95	0.06
Hopf bifurcation	0.86	0.92	0.09
Non-normal (no bifurc.)	0.58	0.93	0.04
Aggregate	0.78	0.94	0.06

6. Discussion

Two empirical features of the results deserve emphasis. First, the gain from the dual-criterion construction is not uniform across scenarios: in the fold and Hopf cases CRTI is at most marginally better than scalar CSD, and the case for the framework rests almost entirely on the non-normal scenario. This is the intended behavior. The framework is designed to suppress a specific failure mode of CSD, not to outperform CSD at its own task. Reported aggregate AUCs are sensitive to the relative weighting of bifurcating and non-bifurcating scenarios in the benchmark, and we encourage readers to interpret the per-scenario numbers in Table 1 as primary.

Second, the construction has nontrivial limitations. (i) Φ_obs and R_obs are window-dependent; while the multi-scale strategy alleviates this, no estimator of mode structure is fully scale-free. (ii) The whitening step in Eq. (6) requires $C\_0(t)$ to be well-conditioned; in regimes where the system already lives on a low-dimensional manifold, regularization becomes load-bearing and conclusions about further compression should be drawn cautiously. (iii) The AAFT surrogate null preserves linear structure, so any genuine nonlinear stationary system will register as "significant" relative to it; this is a feature for tipping-point detection but a caveat for absolute claims of nonstationarity. (iv) The framework as stated is observational. It does not, by itself, license a causal claim that the system is approaching a bifurcation; it licenses the claim that the joint structural-dynamical signature consistent with such an approach is present.

These caveats also clarify the relationship between the present estimators and the original CRTI quantity $T(t) = R(t)/\Phi(t)$. The ratio form is theoretically clean but operationally fragile: it amplifies noise wherever Φ_obs is small, and it loses the diagnostic asymmetry between the structural and dynamical channels by collapsing them into a single number. The phase-synchronization criterion preserves that asymmetry and is the recommended operational quantity for empirical applications.

7. Conclusion

We have argued that the long-running difficulty of distinguishing structural fragility from transient non-normal amplification — a difficulty that is the principal source of false alarms in classical early warning indicators — is best resolved not by inventing yet another scalar metric but by separating the structural and dynamical channels of the system and requiring them to move coherently. The observable effective rank Φ_obs and the leading-mode relaxation proxy R_obs instantiate this separation in a way that is invariant under linear changes of coordinates, robust to finite-sample artefacts, and testable against surrogate nulls. Their phase-synchronized decline-and-rise is the empirical fragility signature of structural singularization.

Two extensions are immediate. The first is a non-stationary surrogate framework that conditions on slow trends rather than treating them as signal, which would tighten the false-positive control further in geophysical and macroeconomic applications. The second is the explicit incorporation of mechanism classes (M1/M2/M3 in the broader CRTI framework) into the alarm logic, allowing the same observables to discriminate not only whether a transition is approaching but also which kind. Both are deferred to subsequent work.

8. References

Boers, N. (2021). Observation-based early-warning signals for a collapse of the Atlantic Meridional Overturning Circulation. Nature Climate Change, 11(8), 680–688.

Carpenter, S. R., & Brock, W. A. (2006). Rising variance: a leading indicator of ecological transition. Ecology Letters, 9(3), 311–318.

Dakos, V., Carpenter, S. R., Brock, W. A., Ellison, A. M., Guttal, V., Ives, A. R., Kéfi, S., Livina, V., Seekell, D. A., van Nes, E. H., & Scheffer, M. (2012). Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. PLoS ONE, 7(7), e41010.

Hennequin, G., Vogels, T. P., & Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. Neuron, 82(6), 1394–1406.

Kuehn, C. (2011). A mathematical framework for critical transitions: bifurcations, fast–slow systems and stochastic dynamics. Physica D, 240(12), 1020–1035.

Mallinckrodt, B. von (2026). The Compression–Response Transition Index: a structural framework for critical transitions. Zenodo preprint.

Murphy, B. K., & Miller, K. D. (2009). Balanced amplification: a new mechanism of selective amplification of neural activity patterns. Neuron, 61(4), 635–648.

Roy, O., & Vetterli, M. (2007). The effective rank: a measure of effective dimensionality. In 15th European Signal Processing Conference (EUSIPCO) (pp. 606–610). IEEE.

Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., & Sugihara, G. (2009). Early-warning signals for critical transitions. Nature, 461(7260), 53–59.

Schreiber, T., & Schmitz, A. (2000). Surrogate time series. Physica D, 142(3–4), 346–382.

Trefethen, L. N., & Embree, M. (2005). Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press.