Persistence Geometry in Ecological Regime Shifts

A Resilience-Model Extension Distinguishing Boundary Proximity from Structural Trajectory

Charles Carroll
ORCID: 0009-0008-6411-7626

May 27, 2026

Abstract

Current ecological early warning signal (EWS) frameworks rely heavily on proximity-based indicators such as rising variance and lag-1 autocorrelation to estimate how close a system lies to a bifurcation boundary. While these methods are theoretically robust, empirical deployments frequently produce false positives, false negatives, and inconsistent transition timing across monitored ecosystems.

This study tests a mechanistically constrained resilience-model extension distinguishing boundary proximity from structural trajectory. We hypothesize that proximity metrics alone incompletely characterize persistence state because they measure only the system's position relative to instability, while structural trajectory variables measure how the system processes stress as it moves through state space.

A fully pre-registered, zero-trust computational pipeline was constructed and locked prior to empirical deployment. The architecture incorporates: (1) dynamic metadata verification through live Ecological Metadata Language (EML) manifest auditing, (2) functional guild translation to eliminate long-term taxonomic drift, and (3) retrospective baseline normalization with log-clipped geometric assembly to prevent future-data leakage.

The pipeline was deployed against the North Temperate Lakes Long-Term Ecological Research (NTL-LTER) repository using Lake Wingra (WI) and Allequash Lake (AL) as the first matched shallow-lake pair. Empirical analysis demonstrated a divergence between noisy proximity signals and steadily degrading structural trajectory metrics within the stressed system. However, formal Cox proportional hazards competition yielded an indeterminate result ($\Delta AIC \approx 0$), attributed to limited event geometry inherent to a single matched-pair timeline.

This outcome validates the inferential discipline of the pre-registration framework by demonstrating that the pipeline refuses to overfit constrained ecological datasets. The computational architecture is now prepared for Phase 2 pooled multi-lake survival analysis involving staggered transition timelines across additional shallow-lake networks.

Contents

1 Introduction

Ecological regime shifts represent one of the central challenges in resilience science and complex systems ecology [7, 9]. Many shallow lake systems exhibit abrupt transitions between relatively stable clear-water and turbid-water states under sustained nutrient loading, hydrological disturbance, or altered trophic dynamics [8, 3].

A major research effort within resilience ecology has focused on the development of early warning signals (EWS) intended to detect critical slowing down prior to transition [5, 6, 4]. These approaches typically monitor statistical indicators such as rising variance, lag-1 autocorrelation, skewness, and flickering behavior within ecological time series.

While EWS methods remain theoretically important, empirical deployment has revealed significant limitations. In many monitored systems, warning signals fluctuate inconsistently, trigger false alarms, or fail to provide coherent directional information prior to observed transitions [1, 2]. This raises a foundational question:

> Do proximity metrics fully characterize ecological persistence state?

The framework explored in this study proposes that they do not.

Specifically, this paper distinguishes between two independent dimensions of resilience dynamics:

(a) **Boundary Proximity**: where a system currently sits relative to a bifurcation boundary.

(b) **Structural Trajectory**: how the system processes stress and moves through state space over time.

Within this framework, proximity indicators such as lag-1 autocorrelation measure only the first dimension. Structural trajectory variables — specifically functional redundancy and repair/degradation dynamics — characterize the second.

The central hypothesis tested here is therefore:

> Systems at similar proximity to instability may possess substantially different persistence states because their structural trajectory geometry differs.

Importantly, this study does not attempt to replace existing EWS frameworks. Rather, it proposes a constrained mechanistic extension intended to explain residual variance left unresolved by proximity-only approaches.

2 Theoretical Framework: Proximity vs. Trajectory

2.1 Boundary Proximity

Traditional EWS methods estimate proximity to instability through signatures of critical slowing down [10]. In shallow lakes, rising autocorrelation and variance in chlorophyll-a dynamics are often interpreted as indicators that the system approaches a bifurcation threshold.

Within the present framework, this dimension is represented by:

$$B_{norm} \tag{1}$$

where B_{norm} denotes normalized proximity metrics derived primarily from lag-1 autocorrelation behavior.

2.2 Structural Trajectory

The second dimension concerns the system's capacity to process stress internally.

Two systems may appear equally close to instability according to EWS metrics while differing substantially in:

- functional redundancy,
- repair capacity,
- nutrient buffering,
- trophic stabilization,
- and stress redistribution pathways.

This second dimension is represented through:

$$\Pi_t = B_{norm}^{w_1} R_{norm}^{w_2} PD_{norm}^{w_3} \tag{2}$$

where:

- R_{norm} = normalized functional redundancy,
- PD_{norm} = normalized repair/degradation ratio,
- $w_1 = w_2 = w_3 = 0.333$ (locked equal weights).

The log-transformed form used operationally is:

$$\ln(\Pi_t) = 0.333 \ln(B_{norm}) + 0.333 \ln(R_{norm}) + 0.333 \ln(PD_{norm}) \tag{3}$$

The theoretical claim is therefore not that additional variables improve prediction by brute force, but that trajectory variables measure a separate structural dimension not fully captured by proximity metrics alone.

3 Pre-Registration and Inferential Discipline

Prior to empirical deployment, the complete computational architecture was frozen under release tag:

```
v1.0.0
```

The following constraints were locked before ingestion of real ecological records:

- Equal predictor weighting.
- Morphological filtering rules.
- Functional guild translation architecture.
- Missing-data exclusion criteria.
- Baseline normalization windows.
- Survival-analysis comparison structure.
- AIC success/failure thresholds.

The framework therefore prohibits:

- post-hoc parameter tuning,
- adaptive weighting,
- retrospective normalization,
- and selective pair inclusion.

4 Methods

4.1 Data Sources

Primary empirical deployment utilized the North Temperate Lakes Long-Term Ecological Research (NTL-LTER) repository.

Core packages included:

- `knb-lter-ntl.29` — Physical and Chemical Limnology
- `knb-lter-ntl.37` — Zooplankton Community Composition
- `knb-lter-ntl.115` — External Phosphorus Loading

4.2 Zero-Trust EDI Manifest Verification

A namespace-robust ingestion auditor dynamically queried live Ecological Metadata Language (EML) manifests through the EDI PASTA API.

The system:

1. identified latest repository revisions,
2. parsed XML manifests dynamically,
3. extracted live entity tokens,
4. verified schema integrity,
5. rejected incompatible datasets automatically.

This architecture prevents silent schema drift from contaminating downstream analyses.

4.3 Morphological Exclusion Criteria

Systems were automatically excluded if:

- mean depth ≥ 3.0 m,
- surface area mismatch ≥ 20
- critical-decade completeness less than 75 percent,
- or consecutive missing-year gaps exceeded 1 year.

4.4 Functional Guild Translation

To eliminate observer bias and taxonomic drift, species-level records were translated into broad stabilizing functional guilds.

Primary zooplankton guilds:

- G1: Large-bodied macro-filter feeders
- G2: Medium omnivorous filter feeders
- G3: Small micro-grazers

Primary macrophyte guilds:

- M1: Submerged structural macrophytes
- M2: Floating-leaved canopy species
- M3: Free-floating degradation indicators

Conservative fallback logic prevented inflation of redundancy metrics under ambiguous taxonomic records.

4.5 Normalization Architecture

All variables were normalized against a retrospective baseline window comprising the first 10 years of each lake record.

No future data were available to the normalizer.

All normalized variables entering the logarithmic assembly stage were clipped to:

$$10^{-3} \tag{4}$$

in order to prevent undefined logarithmic collapse during extreme ecological deterioration.

4.6 Survival Analysis

A Cox proportional hazards framework was used to compare:

1. proximity-only models (B_{norm}),
2. composite persistence geometry models ($\ln(\Pi_t)$).

Control systems were right-censored at the analytical horizon.

Model competition used Akaike Information Criterion (AIC) comparisons.

Pre-registered thresholds:

- $\Delta AIC < -2$: Composite success
- $|\Delta AIC| \leq 2$: Indeterminate
- $\Delta AIC > 2$: Composite failure

5 Adversarial Validation

Prior to live deployment, the ingestion engine underwent adversarial testing using synthetic malformed lake profiles.

The engine successfully excluded:

- deep-lake violations,
- surface-area mismatches,
- excessive data fragmentation,
- consecutive missing-year failures.

This validation established that the architecture enforced pre-registered constraints automatically before model fitting.

6 Phase 1 Empirical Deployment

6.1 Candidate Pair Selection

The first retained matched pair consisted of:

- Lake Wingra (WI) — stressed target system
- Allequash Lake (AL) — low-stress control system

Excluded systems included:

- Lake Mendota — excluded for depth
- Trout Lake — excluded for depth
- Fish Lake — excluded for insufficient completeness

6.2 Empirical Divergence Between Proximity and Trajectory

Within Lake Wingra, the trajectory variables exhibited steady directional deterioration:

- collapse of large-bodied grazer redundancy,
- sustained decline in repair/degradation performance,
- directional structural erosion.

Meanwhile, proximity metrics fluctuated noisily through repeated autocorrelation spikes and resets.

This produced the first empirical indication that:

> Boundary proximity and structural trajectory are not equivalent ecological dimensions.

7 Results

7.1 Cox Model Competition

The formal likelihood comparison produced:

Guild	Description	Functional Role
G1	Large-bodied macro-filter feeders	Primary algal suppression
G2	Medium omnivorous filter feeders	Intermediate grazing pressure
G3	Small micro-grazers	Micro-algal and bacterial grazing
M1	Submerged structural macrophytes	Sediment stabilization and refuge
M2	Floating-leaved canopy plants	Habitat structure
M3	Free-floating surface mats	Degradation indicator

Table 1: Functional guild structure used for redundancy normalization.

Resulting:

$$\Delta AIC = -0.00006 \tag{5}$$

This falls entirely within the pre-registered indeterminate window.

7.2 Interpretation

The indeterminate outcome resulted from limited event geometry.

With only a single transition event available, both models converged on the same critical collapse horizon.

The survival engine therefore lacked sufficient independent temporal structure to separate smooth trajectory deterioration from proximity dynamics statistically.

Importantly, the framework did not inflate weak evidence into apparent success.

8 Discussion

8.1 Scientific Significance

The central contribution of Phase 1 is not confirmation of the persistence geometry hypothesis.

Rather, the contribution is:

1. successful operationalization of a fully pre-registered resilience-model extension,
2. demonstration of a measurable empirical divergence between proximity and trajectory,
3. and validation that the framework behaves conservatively under limited event structure.

The system encountered an empirical boundary condition and halted correctly.

This behavior is scientifically significant because it demonstrates that the architecture prioritizes inferential integrity over narrative interpretation.

8.2 Limitations

Primary limitations include:

- single matched-pair deployment,

- limited event geometry,
- constrained statistical power,
- and ecological domain restriction.

The framework has therefore not earned cross-domain generalization at this stage.

8.3 Phase 2 Scaling

Future deployment will expand the pooled survival matrix through:

- Florida shallow-lake systems,
- Danish shallow-lake repositories,
- additional staggered eutrophication histories.

The pooled Cox architecture will utilize:

```
strata = PAIR_ID
```

in order to preserve local ecological constraint environments.

9 Conclusion

This study documents the successful transformation of a broad systems-level persistence hypothesis into a constrained empirical resilience-model program.

The Phase 1 deployment demonstrated:

- stable computational architecture,
- successful adversarial validation,
- robust ingestion discipline,
- measurable divergence between proximity and trajectory behavior,
- and conservative inferential performance under limited event geometry.

The framework now advances to Phase 2 pooled ecological scaling under fully locked constraints.

Its future validity will be determined entirely by empirical performance across independent collapse histories.

Appendix A: Core Pipeline Logic

```
ln_Persistence =
0.333 * ln(B_norm)

* 0.333 * ln(R_norm)
* 0.333 * ln(PD_norm)
```

Appendix B: Failure Criteria

The framework is considered operationally unsupported if:

- pooled multi-event survival analysis fails to outperform proximity-only models,
- trajectory variables remain statistically indistinguishable from proximity metrics,
- or operationalizations fail reproducibility testing across independent ecological repositories.

References

[1] Boettiger, C., & Hastings, A. (2013). Tipping points: From patterns to predictions. *Nature*, 493, 157–158.

[2] Burthe, S. et al. (2016). Do early warning indicators consistently predict nonlinear change? *Journal of Applied Ecology*, 53(3), 666–676.

[3] Carpenter, S. R. et al. (1999). Management of eutrophication for lakes subject to potentially irreversible change. *Ecological Applications*, 9(3), 751–771.

[4] Carpenter, S. R. et al. (2011). Early warnings of regime shifts: A whole-ecosystem experiment. *Science*, 332(6033), 1079–1082.

[5] Dakos, V. et al. (2008). Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38), 14308–14312.

[6] Drake, J. M., Griffen, B. D. (2010). Early warning signals of extinction in deteriorating environments. *Nature*, 467, 456–459.

[7] Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4, 1–23.

[8] Scheffer, M. et al. (1993). Alternative equilibria in shallow lakes. *Trends in Ecology Evolution*, 8(8), 275–279.

[9] Scheffer, M. et al. (2001). Catastrophic shifts in ecosystems. *Nature*, 413, 591–596.

[10] Scheffer, M. et al. (2009). Early-warning signals for critical transitions. *Nature*, 461, 53–59.