Paper ID: FRC 787.787 | Version: v1.2

Lambda-Flight v1.2: Operational and Latent Coherence Transport Models

Author: Hadi Servat

Affiliation: Independent Researcher

Contact: publish@fractalresonance.com

Website: fractalresonance.com

Abstract

FRC 787.787 v1.2 separates three objects previously called Lambda. Lambda_obs=Lambda0 ln C_obs is an operational transform of preregistered observables; Lambda_eq is an observation-derived target; Lambda_dyn is an optional latent reaction-diffusion state that relaxes toward that target. A transport model must select Lambda_obs or separately validated Lambda_dyn before scoring. Support for Lambda_dyn requires improvement over the observation-only model and an equal-feature non-Lambda latent baseline. Even successful residual prediction would not establish a fundamental field. The nested baselines, negative controls, six test cards, and numerical regime pilot are retained. Version 1.2 adds a locally frozen secondary-

> *data test of P3 on a public Rijke-tube benchmark. The composite is not promoted on that benchmark: its negative held-out point estimate is accompanied by low power, a class/protocol confound, and an imperfect equal-information control. The external test concerns predictive early warning only and does not test entropy-coherence reciprocity.*

1. Scope

FRC 787.787 is a framework paper for coherent transport in driven open physical systems. It is not a claim of a new fundamental force, a replacement for accepted equations, or a license to reinterpret famous experiments after the fact.

The framework asks one testable question: does a measured coherence field improve held-out prediction after the accepted baseline has already been given a fair chance?

If the answer is no, Lambda-Flight fails for that platform. If the answer is yes, the result is a nested residual law, not a discovery that ordinary dynamics were absent.

This is the consistency lock with the 100-series. FRC 100.003 treats microstate-distribution dynamics as one conditional collapse candidate. FRC 100.007 supplies a narrow rho-drift audit, not a replacement ontology. FRC 566.030 computes an exact system-only ledger curve without settling the open physical conjecture. FRC 787.787 therefore keeps every Lambda object operationally typed and baseline-tested.

2. Definitions

Let $\Omega \subset \mathbb{R}^d$ be a physical domain. The minimal measured fields are:

Symbol	Meaning
$\rho(x,t)$	conserved or approximately conserved density
$\psi(x,t)$	order parameter, mode amplitude, phase field, or envelope
$E(t)$	external drive or control parameter
$C_{\text{obs}}(x,t) > 0$	positive dimensionless coherence functional built from preregistered observables
$\Lambda_{\text{obs}}(x,t)$	instantaneous operational transform
$\Lambda_{\text{eq}}(x,t)$	observation-derived calibrated target

$\Lambda_{\rm dyn}(x,t)$	optional latent or coarse-grained dynamical state

The operational transform is

$$\Lambda_{\rm obs}(x,t) = \Lambda_0 \ln C_{\rm obs}(x,t).$$

A useful construction is multiplicative:

$$C_{\rm obs}(x,t) = \prod_{j=1}^{m} C_j(x,t)^{w_j}, \qquad w_j \geq 0, \qquad \sum_j w_j = 1.$$

Then

$$\frac{\Lambda_{\rm obs}(x,t)}{\Lambda_0} = \sum_{j=1}^{m} w_j \ln C_j(x,t).$$

The weights and normalization constants must be fixed before testing or fitted only on training data. A field built after seeing the event is a visualization, not evidence.

The target $\Lambda_{\rm eq} = \widehat{\mathcal{G}}_{\rm eq}[O,E]$ is fixed by a declared calibration or training rule. Where finite lag is tested, the optional latent state obeys

$$\tau_\Lambda(\partial_t + u\cdot\nabla)\Lambda_{\rm dyn} = \ell_\Lambda^2 \nabla^2 \Lambda_{\rm dyn} - \Lambda_{\rm dyn} + \Lambda_{\rm eq} + \xi_\Lambda.$$

$\Lambda_{\rm dyn}$ is a model state, not a directly observed or fundamental field. A fundamental Lambda field is a separate conjecture requiring independent assumptions and evidence.

3. Coherence Factors

The factors C_j must be platform-specific and measurable. Examples include:

Factor	Example construction	Interpretation
Amplitude regularity	$C_A = \exp[-\,\mathrm{Var}_\ell(A)/A_0^2]$	smoothness of envelope amplitude
Phase order	$C_\phi = R_\ell + \epsilon$, with R_ℓ the local circular phase order	local phase locking

Spectral concentration	$C_s = (P_{\text{band}} + \epsilon)/(P_{\text{total}} + \epsilon)$	mode concentration
Defect suppression	$C_d = \exp[-n_d/n_0]$	low local defect density
Correlation persistence	$C_\xi = \exp[-\ell_0/\xi_{\text{corr}}]$	long correlation length
Dissipation regularity	$C_\sigma = \exp[-\sigma_{\text{exc}}/\sigma_0]$	reduced excess entropy-production proxy
Boundary matching	$C_b = \exp[-M_b^2/b_0^2]$, where M_b is boundary mismatch	compatibility with constraints

No row in this table is mandatory. A platform may require different factors. What is mandatory is that the factors are chosen before the held-out prediction and that simpler single-factor explanations are tested.

4. Operational Lambda-Flight

The residual transport field must be selected before scoring:

$$\Lambda_{\text{tr}} = \begin{cases} \Lambda_{\text{obs}}, & \text{observation-only residual model,} \\ \Lambda_{\text{dyn}}, & \text{separately validated finite-lag model.} \end{cases}$$

Let $q(x,t)$ be the measured density or intensity of a coherent structure. Define its center and radius:

$$X(t) = \frac{\int_\Omega x q(x,t)\,dx}{\int_\Omega q(x,t)\,dx}, \qquad R^2(t) = \frac{\int_\Omega \|x - X(t)\|^2 q(x,t)\,dx}{\int_\Omega q(x,t)\,dx}.$$

Let $K(t)$ be an internal coherence score such as phase order, spectral concentration, or inverse defect density. A structure is a Lambda-Flight candidate over $[0,T]$ only if:

1. $K(t)$ stays above a preselected threshold $K_{\min}$ for most of $[0,T]$.
2. $R(t)$ grows more slowly than the matched diffusive or dissipative baseline.
3. The residual velocity after known baseline forces are removed is predicted by the precomputed $\nabla\Lambda_{\text{tr}}$ field on held-out data.

The third condition is the load-bearing one. Persistent motion alone is not Lambda-Flight.

5. Minimal Closure

The minimal continuum closure is a nested extension of ordinary transport:

$$\partial_t \rho + \nabla \cdot (\rho u) = \nabla \cdot (D_\rho \nabla \rho - \chi \rho \nabla \Lambda_{\mathrm{tr}}) + S_\rho .$$

The flux is

$$J_\rho = \rho u - D_\rho \nabla \rho + \chi \rho \nabla \Lambda_{\mathrm{tr}} .$$

This equation must be read conservatively. The term $\chi \rho \nabla \Lambda_{\mathrm{tr}}$ is a residual closure. It is permitted only when:

1. the accepted baseline transport model has been specified;
2. its residuals are measured;
3. $\nabla \Lambda_{\mathrm{tr}}$ predicts those residuals out of sample;
4. simpler equal-access alternatives do not explain the same gain.

If these conditions are not met, χ is set to zero and the model reduces to the baseline.

Support for Λ_{dyn} requires improvement over both the observation-only Λ_{obs} model and an equal-feature non-Lambda latent baseline. Otherwise $\Lambda_{\mathrm{tr}} = \Lambda_{\mathrm{obs}}$ or the residual term is removed.

6. Variational Discipline

A phenomenological Lyapunov-compatible form is

$$F[\rho, \psi, \Lambda_{\mathrm{dyn}}] = \int_\Omega \Big[f_\rho(\rho) + \frac{\kappa_\psi}{2} |\nabla \psi|^2 + U(\psi) + \frac{\kappa_\Lambda}{2} |\nabla \Lambda_{\mathrm{dyn}}|^2 + \frac{a_\Lambda}{2} (\Lambda_{\mathrm{dyn}} - \Lambda_{\mathrm{eq}})^2 - \eta \Lambda$$

The original formal draft treated this as a free-energy proof. The corrected status is weaker and cleaner: it is a Lyapunov-compatible template under stated assumptions. If Λ_{eq} depends on ρ or ψ, the full functional derivatives include additional terms. Those terms must be retained in a platform-specific derivation. Without that derivation, the variational section is a discipline check, not a theorem of monotonic free-energy decay.

7. Linear Stability

For the two-field model

$$\partial_t \rho = D\nabla^2 \rho - \chi \nabla \cdot (\rho \nabla \Lambda_{\mathrm{dyn}}),$$

$$\tau \partial_t \Lambda_{\mathrm{dyn}} = \ell^2 \nabla^2 \Lambda_{\mathrm{dyn}} - \Lambda_{\mathrm{dyn}} + a(\rho - \rho_0),$$

linearizing around $(\rho, \Lambda_{\mathrm{dyn}}) = (\rho_0, 0)$ with perturbations $e^{\sigma t + ik\cdot x}$ gives

$$(\sigma + Dk^2)(\tau\sigma + 1 + \ell^2 k^2) - \chi a \rho_0 k^2 = 0.$$

Coherence feedback can overcome diffusive smoothing only when

$$\chi a \rho_0 > D(1 + \ell^2 k^2).$$

This is a candidate instability threshold. It becomes a physical claim only after a, χ, and the coherence channel are estimated without test leakage.

8. Flight Numbers

Let L be the structure size, $\Delta\Lambda_{\mathrm{tr}}$ the selected transport-field contrast, and D the effective diffusivity. The coherence transport number is

$$Pe_\Lambda = \frac{v_\Lambda L}{D} \sim \frac{\chi \Delta\Lambda_{\mathrm{tr}}}{D}.$$

For the finite-lag Λ_{dyn} model only, let V be the observed or baseline-predicted structure speed. The relaxation number is

$$St_\Lambda = \frac{\tau_\Lambda V}{L}.$$

The regime map is:

Regime	Expected behavior	Kill condition
$Pe_\Lambda < 1$	diffusion or baseline transport dominates	claimed flight persists with shuffled or zero Lambda
$Pe_\Lambda > 1$, $St_\Lambda \ll 1$	adiabatic coherence-guided transport possible	no held-out residual prediction

$Pe_\Lambda > 1$, $St_\Lambda \sim 1$	lag, overshoot, hysteresis	hysteresis is already explained by baseline
$St_\Lambda \gg 1$	Lambda decouples from moving structure	no degradation relative to low-St_Λ case

The estimates must be made from training data or direct calibration:

Quantity	Required estimate
χ	residual mobility from training trials only
$\Delta\Lambda_{\rm tr}$	held-out-blind contrast of the preselected transport field
D	matched diffusion or spreading coefficient
τ_Λ	response lag from independent relaxation or sweep data
V	baseline speed or training-set structure speed
L	preselected structure scale

9. Nested Baselines

Lambda-Flight is evaluated as a nested residual model. The model stack is:

Level	Model	Purpose
B0	persistence baseline	recent trajectory continues
B1	accepted platform equation	PDE, TDDFT, phase-field, Kuramoto, reaction-diffusion, or other domain baseline
B2	best single observable	amplitude, phase order, defect density, energy, or one coherence factor
B3	equal-feature non-Lambda latent model	same observables but no Lambda geometry or relaxation law
B4	typed Lambda model	$B1 + \Lambda_{\rm tr}, \nabla\Lambda_{\rm tr}$ and, only for the lag model, τ_Λ

Support requires B4 to beat B0-B3 on held-out data by a preselected metric with uncertainty estimates. A visible Lambda map is not support. A post-hoc fit is not support. A famous experiment that can already be explained by B1 is not support.

10. Negative Controls

Every serious test should include at least these controls:

1. Shuffled Lambda labels across trials.
2. Time-reversed Lambda fields.
3. Spatially shifted Lambda fields.
4. Single-factor ablations.
5. Equal-feature latent baselines.
6. Training/test split by run, parameter regime, or time block.
7. No-refit prediction horizons fixed before scoring.

If the Lambda model survives only when these controls are omitted, it is not a pillar result.

11. Six Test Cards

The six predictions are now written as test cards. Each has a null and a kill clause.

Prediction	Test	Null	Kill
P1 coherence-gradient drift	residual $\dot{X}_{\text{res}}$ aligns with $\nabla\Lambda_{\text{tr}}$	residuals are baseline noise	no held-out alignment after controls
P2 basin residence	T_{res} scales with $\Delta\Lambda_{\text{tr}}$ in a rare-event regime	residence explained by energy or amplitude	single observable beats Lambda
P3 transition early warning	Lambda curvature, variance, or gap degrades before collapse	baseline observable warns equally early	no lead time on held-out transitions
P4 finite-relaxation hysteresis	lag grows with $\tau_\Lambda r$ under sweeps	standard response delay explains loop	no Lambda-specific scaling
P5 coherence-over-amplitude selection	higher Lambda structure persists when amplitude is matched	amplitude or energy selects structure	amplitude-only model wins
P6 Pe/St collapse	flight fails outside registered Pe/St window	failure is numerical or baseline instability	Pe/St does not stratify outcomes

11.1 Executed P3 benchmark: Rijke-tube early warning

A locally frozen secondary-data test applied P3 to the public Rijke-tube thermoacoustic transition data used by Bury et al. [16], originally collected by Pavithran and Sujith [17]. The analysis used the pinned reproduction repository at commit `d56b5df879efbd557dad38aca1281cd00dc404c9` [18]. It compared accepted variance/autocorrelation indicators (B1), the best single coherence factor (B2), a factor-tau model (B3), and the fixed logarithmic composite added to B1 (B4). Splits, horizons, features, model selection, and the promotion rule were frozen before the FRC coherence features were computed. The checksum record establishes local chronology rather than an externally timestamped preregistration, and the included post-analysis amendment discloses the later strict-JSON repair.

P3 status: NOT PROMOTED on this benchmark. At the frozen primary horizon of -0.50 s, row-level held-out AUC was 0.722 for B1 and 0.694 for B4, giving

$$\Delta\,\mathrm{AUC} = \mathrm{AUC}_{B4} - \max(\mathrm{AUC}_{B1}, \mathrm{AUC}_{B2}, \mathrm{AUC}_{B3}) = -0.028.$$

After aggregating the two segments from each null run, group-level AUC was 0.667 for B1 and 0.611 for B4. The exact nine-group label permutation gave $p = 0.357$. The stratified group-bootstrap stability interval for the row-level increment was $[-0.250, 0.111]$. That interval includes zero and gains above the registered promotion threshold of 0.05, so the negative point estimate is low-power evidence and does not establish absence of an effect. B4 first reached AUC ≥ 0.75 at -0.25 s, the same frozen horizon as B1, and therefore supplied no registered lead-time advantage.

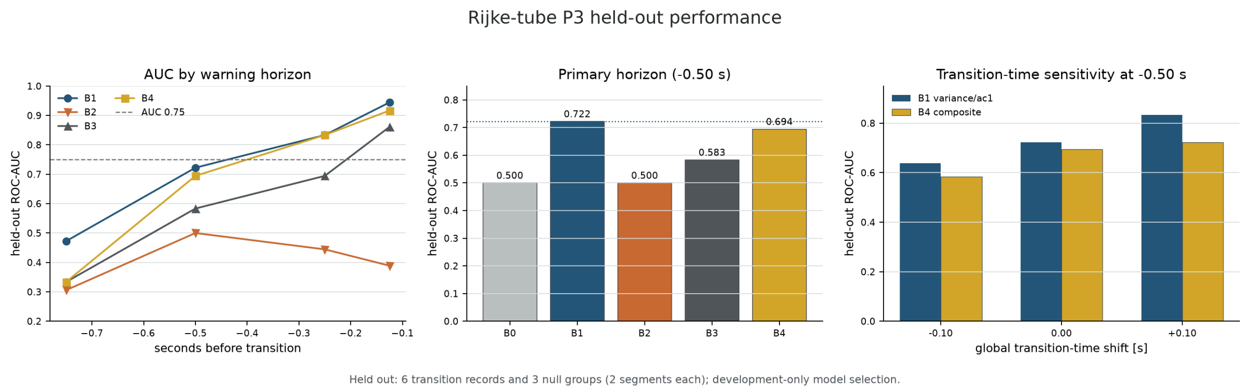

Held-out Rijke-tube P3 performance by warning horizon, primary comparison, and transition-time sensitivity.

Three interpretation limits are load-bearing. First, transition records came from forced ramps while controls came from quasi-static experiments; class and experimental protocol are therefore confounded, so the benchmark is not a matched-ramp estimate of transition forecasting. Second, B3 was not fully equal-information: B4 formed a joint rolling composite before temporal reduction, while B3 received separately reduced factor taus. Third, the small held-out cohort and unstable development folds limit power. These limitations prevent both promotion and decisive rejection. The result is a registered benchmark-specific non-promotion, not a closed mechanism or a rejection of P3 on every platform.

The test concerns predictive early warning only. It does not measure entropy-coherence reciprocity, validate a thermodynamic ledger, support a fundamental Lambda field, or alter the canonical k^* stance. The protocol, source record, provenance amendment, independent review, feature tables, predictions, scripts, tests, and figure are included under `external/rijke-p3/`.

12. Anti-Precursor-Hunting Rule

FRC 787.787 must not be used to claim that a new external result confirms FRC because it contains a delay, a coherent packet, a transition, or a transport anomaly. That is precursor hunting.

The rule is:

1. The Lambda observables must be defined before scoring the external result.
2. The accepted baseline must be implemented or quoted as the primary null.
3. The Lambda model must make a separable prediction that the baseline does not already determine.
4. The prediction must be scored out of sample.

The attosecond-STM gate is the example. A measured tunneling delay is not automatically τ_Λ. The disciplined question is whether field-to-response delay tracks a pre-registered surface coherence variable better than the TDDFT response baseline. If TDDFT already contains the dependence, there is no FRC paper. If a separable Lambda relaxation law beats TDDFT out of sample, the result can become a separate confirmation-of-existing-prediction paper.

13. Worked Platform Protocol

A controlled first platform should be easier than attosecond STM. Suitable candidates are coupled oscillator lattices, driven electrical resonator arrays, nonlinear optical cavities, reaction-diffusion media, and phase-field simulations with known ground truth.

A minimal protocol:

1. Choose the accepted platform equation and fit it on training data.
2. Define C_j factors from measured present/past observables only.

3. Fix normalization and weights on training data.
4. Compute Λ_{obs} and, if registered, Λ_{dyn}; select Λ_{tr}; then compute its gradients and flight numbers.
5. Predict held-out trajectory, transition time, residence time, or collapse class.
6. Score B0-B4 and all negative controls.
7. Publish the failure if B4 does not beat the baselines.

14. Numerical Pilot

The release package includes a deterministic toy pilot:

`pilot/pilot_lambda_flight.py`

The pilot simulates a conserved one-dimensional packet trying to follow a moving coherence basin. It uses a conservative upwind finite-volume flux for density transport and a semi-implicit Fourier update for Lambda relaxation. The sweep varies $Pe_\Lambda \in \{0.35, 1.20, 3.00\}$ and $St_\Lambda \in \{0.10, 0.60, 1.50\}$.

The code labels $Pe_\Lambda \geq 1$ and $St_\Lambda \leq 0.6$ as its internal regime window. Because that label is built into the toy model, the pilot is only a solver and regime sanity check; it does not independently validate Prediction 6 or a physical flight threshold.

Metric	Best registered case	Worst case
Pe_Λ	3.00	0.35
St_Λ	0.10	1.50
median lag / L	0.448	0.499
final basin mass	0.594	0.450
final width ratio	1.388	1.709
tracking score	0.350	0.205

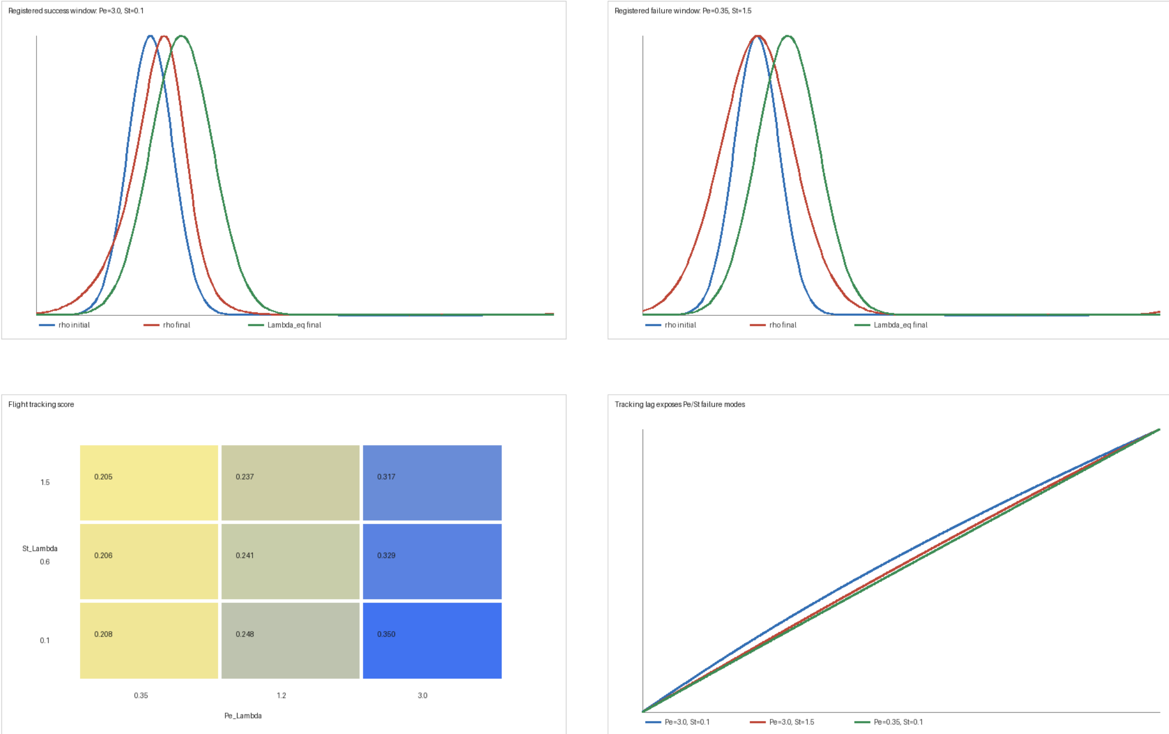

Lambda-flight toy pilot. A conserved packet follows a moving coherence basin best in the high-Pe_Λ, low-St_Λ corner. The pilot is a numerical regime check, not empirical confirmation.

The pilot gives the right limited lesson: higher Pe_Λ improves capture, larger St_Λ weakens tracking, and low Pe_Λ cannot be rescued by plotting Lambda after the fact. The separation is modest, so the paper should not oversell it.

15. Publication Standard

FRC 787.787 becomes a pillar paper only if its public version contains:

1. the measured-Lambda consistency lock;
2. a strict nested baseline hierarchy;
3. anti-precursor-hunting language;
4. pre-registration cards for each prediction;
5. negative controls;
6. a reproducible pilot or worked platform protocol;
7. clear kill clauses.

This v1.2 package contains those elements at the framework level and records one completed external P3 benchmark against an accepted variance/autocorrelation baseline. That benchmark did not meet the promotion gates and was underpowered to establish absence of an effect. The manuscript therefore remains a working preprint rather than a pillar result.

16. Conclusion

Lambda-Flight is a narrow claim about residual coherent transport. The operational transform

$$\Lambda_{\text{obs}} = \Lambda_0 \ln C_{\text{obs}}$$

turns measured coherence factors into a geometry that can be tested against motion, stability, residence, hysteresis, and collapse. An optional Λ_{dyn} adds finite lag only if it beats both the observation-only model and an equal-feature latent baseline. The framework is useful only when the selected transport field improves held-out prediction beyond accepted baselines and simpler alternatives.

That restriction is the point. If Lambda does not beat the registered baselines, the scoped claim is not promoted on that benchmark; the Rijke-tube P3 result records one such case without turning low power into decisive rejection. If Lambda survives nested tests elsewhere, 787.787 can become a disciplined bridge between the 100-series coherence program and concrete non-equilibrium physical systems.

Even a successful residual model would not establish a fundamental Lambda field.

17. Version note

Version 1.2 preserves the v1.1 separation of Λ_{obs}, Λ_{eq}, and Λ_{dyn}, the preselection of Λ_{tr}, and the canonical k^* stance. It adds the reviewed Rijke-tube P3 secondary-data result as a benchmark-specific non-promotion with explicit low-power, protocol-confound, equal-information, and predictive-only limits.

References

1. L. Onsager, "Reciprocal relations in irreversible processes. I," Physical Review 37, 405-426 (1931).
2. L. Onsager, "Reciprocal relations in irreversible processes. II," Physical Review 38, 2265-2279 (1931).
3. P. C. Hohenberg and B. I. Halperin, "Theory of dynamic critical phenomena," Reviews of Modern Physics 49, 435-479 (1977).

4. G. Nicolis and I. Prigogine, Self-Organization in Non-Equilibrium Systems. Wiley (1977).
5. H. Haken, Synergetics: An Introduction. Springer (1983).
6. Y. Kuramoto, Chemical Oscillations, Waves, and Turbulence. Springer (1984).
7. H. Risken, The Fokker-Planck Equation. Springer (1989).
8. M. C. Cross and P. C. Hohenberg, "Pattern formation outside of equilibrium," Reviews of Modern Physics 65, 851-1112 (1993).
9. A. J. Bray, "Theory of phase-ordering kinetics," Advances in Physics 43, 357-459 (1994).
10. J. W. Cahn and J. E. Hilliard, "Free energy of a nonuniform system. I. Interfacial free energy," Journal of Chemical Physics 28, 258-267 (1958).
11. S. M. Allen and J. W. Cahn, "A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening," Acta Metallurgica 27, 1085-1095 (1979).
12. A. Pikovsky, M. Rosenblum, and J. Kurths, Synchronization: A Universal Concept in Nonlinear Sciences. Cambridge University Press (2001).
13. R. Zwanzig, Nonequilibrium Statistical Mechanics. Oxford University Press (2001).
14. S. H. Strogatz, Nonlinear Dynamics and Chaos. Westview Press (1994).
15. M. C. Cross and H. Greenside, Pattern Formation and Dynamics in Nonequilibrium Systems. Cambridge University Press (2009).
16. T. M. Bury, R. I. Sujith, I. Pavithran, M. Scheffer, T. M. Lenton, M. Anand, and C. T. Bauch, "Deep learning for early warning signals of tipping points," Proceedings of the National Academy of Sciences 118, e2106140118 (2021). DOI: 10.1073/pnas.2106140118.
17. I. Pavithran and R. I. Sujith, "Effect of rate of change of parameter on early warning signals for critical transitions," Chaos 31, 013116 (2021). DOI: 10.1063/5.0025533.
18. T. M. Bury et al., `deep-early-warnings-pnas`, public reproduction repository, commit `d56b5df879efbd557dad38aca1281cd00dc404c9`, https://github.com/ThomasMBury/deep-early-warnings-pnas.