Brain Criticality — Phase Transition at the Edge of Chaos via the Latent Framework

Dr. Tamás Nagy

tnagyphd@gmail.com

draft • 2026-04-08

Executive summary

Experimental and modeling work on neural circuits often appeals to "criticality," but the hypothesis is hard to test when measurements are noisy and finite. This note stays deliberately modest: it uses a classical Curie–Weiss Ising scaffold for transparent mean-field intuition, then separates that story from a disordered, sparse-coupling Metropolis simulation that is closer in spirit to heterogeneous biological connectivity. On the formal side, a bundle of machine-checked real inequalities in Platonic organizes the narrative into order-parameter, subcritical-style, supercritical-style, critical-window, information-processing, and cross-domain blocks—without pretending those lemmas rederive Curie–Weiss on the page. On the numerical side, a small regression harness checks internal consistency (variance, spectral compressibility, a mutual-information proxy, and a coarse scan) and explicitly records a case where the largest on a temperature grid does not sit at the nominal mean-field scale. The aim is a reproducible calibration object for Latent-style compressibility diagnostics, not a claim about any particular cortical recording.

Abstract

Many neural systems are hypothesized to operate near a critical point between ordered and disordered dynamics, balancing sensitivity and stability. Testing this hypothesis on finite data requires metrics that survive estimation noise and connect to established statistical physics models.

The Latent framework contributes scalar diagnostics that remain comparable across experiments once measurement embeddings are fixed, complementing scaling-collapse analyses that demand large dynamic ranges.

This paper pairs a **mean-field Curie–Weiss scaffold** (§2) with a **Platonic lemma bundle** (§3) and **Monte Carlo diagnostics** on finite spin systems (§4). The Latent Number ρ used in §2.4–§4 is a spectral compressibility ratio built from sampled covariance eigenvalues (see numerical_validation.py); it is **not** the abstract real parameter ρ that appears in the inequality lemma critical_balance_rho_one inside the formal file. The effective dimension N^* counts directions needed for a fixed cumulative variance threshold in that same embedding.

Thirty-six kernel-verified theorems in elysium/fields/bio_brain_criticality/platonic.py (zero user axioms in the domain state) organize into six narrative groups: order-parameter constraints, subcritical-style bounds, supercritical-style bounds, critical-window inequalities, information-processing lemmas, and cross-domain bridges to Ising and epidemic imagery. **§3 states how those lemmas sit in the story; it is not a full mean-field derivation on the page.**

Numerical regression uses $N = 50$ **sparse random symmetric couplings** (not the all-to-all Hamiltonian of §2). Fourteen scripted checks pass; $\rho > 1$ holds on the three reference temperatures,

while susceptibility need not peak at the nominal T_c scale when coupling is heterogeneous—the scan exhibits a χ maximum on the sampled grid (see §4).

1. Introduction

1.1 Criticality as a Latent phenomenon

Criticality is fundamentally about scaling: correlation lengths diverge, susceptibilities peak, and fluctuations organize along special directions. Those directions are Latent in the literal sense: few collective modes capture most of the variance at criticality, even as microscopic dimension grows.

1.2 Ising scaffold

We use Curie–Weiss (all-to-all) Ising models for analytic transparency. The choice sacrifices biological realism in connectivity but preserves the core Latent story about emergent low-dimensional structure near T_c. Extensions to sparse random graphs are noted throughout.

1.3 Contributions

A Platonic-verified inequality bundle (36 targets) for the criticality narrative, reproducible Monte Carlo diagnostics with 14 regression checks, and cross-domain lemmas linking the same abstract quantities to Ising and epidemic imagery.

1.4 Why start with equilibrium Ising

Neural slices are nonequilibrium. Equilibrium Ising is nonetheless the standard sanity scaffold for criticality because its phase diagram is analytically controlled. Latent statistics computed here serve as calibration targets: any nonequilibrium extension must recover this skeleton in appropriate limits.

2. Mathematical Framework

2.1 Energy and measure

Spins $s_i \in \{\pm 1\}$, $i = 1, \dots, N$. Hamiltonian

$$H(s) = -\frac{J}{2N} \sum_{i \neq j} s_i s_j - h \sum_i s_i,$$

with coupling $J > 0$, external field h, and Curie–Weiss scaling $1/N$ ensuring extensive free energy. Inverse temperature $\beta = 1/T$ in natural units.

2.2 Order parameter

Magnetization $m(s) = \frac{1}{N} \sum_i s_i$ is the primary order parameter. Its distribution sharpens as $N \to \infty$, with nonzero spontaneous magnetization below T_c at $h = 0$.

2.3 Susceptibility

Define $\chi = N(\langle m^2 \rangle - \langle m \rangle^2)$ at $h = 0^+$. Divergence of χ as $N \to \infty$ marks criticality in mean-field theory. Finite-N peaks locate the crossover region studied numerically.

2.4 Latent embedding of spin configurations

Let $\Phi(s)$ map each configuration to features including $m(s)$, low-order Fourier modes across spin indices (trivial here due to exchangeability), and polynomial moments up to fixed degree. In the **numerical suite**, ρ is implemented as the ratio of the two largest eigenvalues of the activity covariance (larger ρ means faster spectral decay); N^* is the count of eigenmodes needed to reach 90% cumulative variance. This operational ρ should not be conflated with the formal parameter named rho in the Platonic lemmas, where the identifier is only a typed real variable satisfying the listed inequalities.

2.5 Exchangeability simplification

Because Curie–Weiss interactions are permutation invariant, any feature that breaks index symmetry artificially inflates N^*. We therefore restrict Φ to permutation-invariant polynomials up to degree four, sufficient to capture bimodality near T_c without overfitting noise.

2.6 Estimators

Expectations $\langle \cdot \rangle$ in §4 are Monte Carlo averages after burn-in, matching numerical_validation.py. Extended jackknife or batch debiasing is left to future revisions of that harness.

3. Formal Proof Bundle

Source of truth. All lemma targets are executed in elysium/fields/bio_brain_criticality/platonic.py. A full verify_all() pass reports **37/37** checks: **36** are the named narrative lemmas below, plus one logic-bootstrap declaration (Not) walked by the same harness. The domain metadata (elysium/fields/bio_brain_criticality/ELYSIUM_STATE.yaml) agrees with **0** user axioms at last verification. The lemmas are deliberately lightweight: chained linear and nonlinear real inequalities that encode monotonicity, nonnegativity, and bookkeeping for the criticality story—**not** a line-by-line Curie–Weiss analysis (that material stays in §2 as classical background).

Grouping (mirrors the file's section headers).

- **Order parameter (6).** Variance, susceptibility, correlation-length, avalanche, and branching proxies as typed nonnegativity or spike bounds.
- **Subcritical-style block (6).** Decay, finite correlation sums, low-entropy ordered signatures, sub-extensive means, spectral-gap proxies, refractory stabilization.
- **Supercritical-style block (6).** Growth lower bounds, runaway activity caps, high-entropy disordered signatures, short correlation lengths, saturation, burst probabilities.
- **Critical window (6).** Balance inequalities using a formal rho parameter, power-law tail upper bounds, mutual-information ceilings, scale-invariance windows, finite-size scaling exponents, mixing slowdown placeholders.
- **Information processing (6).** Entropy rate, channel capacity, Fisher information, Cramér–Rao style bounds, coding efficiency, metabolic cost per bit.

- **Cross-domain (6).** Parallel Ising criticality lemmas, susceptibility bridges, epidemic threshold analogy, percolation overlap, universality inequalities, triple-domain product bounds.

Logical ordering. Narratively, blocks 1–4 follow temperature intuition; block 5 layers information quantities; block 6 is largely independent modulo the shared real inequalities.

4. Numerical Validation

We run the bundled harness elysium/fields/bio_brain_criticality/numerical_validation.py on **sparse symmetric random Ising networks** with $N = 50$. Couplings J_{ij} are drawn $\mathcal{N}(0, 1/\sqrt{N})$, symmetrized, zero diagonal—so this section **does not** instantiate the all-to-all Hamiltonian of §2. The script compares three reference temperatures expressed as multiples of a nominal scale $T_c \equiv 1$: $2T_c$, T_c, and $0.5T_c$. (Historical variable names in the code call the high-temperature draw sub and the low-temperature draw super; those labels are **not** the standard statistical-mechanics usage of sub-/super-critical.)

Representative output from a reproducible run (8000 sweeps after 2000 burn-in per reference network):

Regime (temperature)	T/T_c	χ	ρ (spectral)	N^*/N
High-T reference	2.0	0.87	1.13	78%
Mid reference	1.0	0.71	1.02	70%
Low-T reference	0.5	0.33	1.48	54%

A coarse scan on $T \in [0.3, 3.0]$ (10 points, 3000 sweeps + 1000 burn-in) places the largest empirical χ near $T \approx 2.7$ under this random-coupling ensemble—**not** automatically at the §2 mean-field crossover. The 14 automated checks instead verify internal consistency: nonnegative variances, $\rho > 1$ at the three references, mutual-information proxy ≥ 0, correlation sanity, a χ peak exists on the grid, and the mid-T susceptibility stays above half the smaller wing value.

Test harness (14/14). Susceptibility ordering sanity, per-regime variance/spectral/MI/correlation probes, and the scan peak test—see the script for exact predicates.

Sampler. Single-spin-flip Metropolis with detailed balance at the quoted sweep counts.

Replication. Fix the NumPy RNG seeds in make_ising / simulate_ising to reproduce the table; increase sweeps if tighter error bars are required.

5. Cross-Domain Connections

Classical Ising. Mean-field theory is the explicit bridge; finite-N Latent coordinates capture departures from pure self-averaging.

Epidemic thresholds. Dense-graph SIR models exhibit threshold behavior analogous to spontaneous symmetry breaking. The formal bundle includes monotonicity lemmas that support **qualitative** dashboards comparing neural and epidemiological order parameters; quantitative peak alignment must be checked model-by-model (§4 shows a counterexample where χ and spectral ρ peak at different temperatures under disorder).

6. Discussion

The Latent perspective ties criticality to compressibility: collective modes can dominate covariance spectra even though the microstate space is exponential. $\rho > 1$ in the spectral diagnostic signals nontrivial slippage between the top two eigenvalues; the brain-criticality hypothesis, translated into this language, predicts structured N^* trajectories when a biological preparation is embedded in an informative Φ. The heterogeneous $N = 50$ toy model shows that χ and ρ need not peak together—an honest warning when mapping mean-field intuition to disordered connectivity.

Limitations: Curie–Weiss connectivity ignores spatial structure and biological constraints. Results are meant as a controlled Latent calibration, not as a claim about empirical cortical data.

Future work will repeat the pipeline on Erdős–Rényi and small-world graphs, add external drive (stimulus ensembles), and compare ρ trajectories during plasticity rules that move networks toward or away from criticality.

Non-claims. We do not analyze electrophysiological recordings here; all numbers are from simulated Ising systems.

Practical interpretation. If a biological preparation shows χ-like susceptibility peaks without a ρ shift, the system may be critical in mean-field sense but not compressible in the chosen measurement basis—hinting that the relevant Latent features are missing from Φ.

Ethics and data. This manuscript reports only simulated spin systems; no animal or human neural recordings are included.

Open bridge. When external drive breaks detailed balance, ρ may track nonequilibrium effective temperatures; formal theorems await an extension of the Latent calculus to steady-state large deviations.

During the preparation of this work the author used large language models to assist with manuscript drafting and organization. The author reviewed and edited the content and takes full responsibility for the publication.

References

1. Beggs, J. M. & Plenz, D. (2003). "Neuronal avalanches in neocortical circuits." *J. Neurosci.*
2. Chialvo, D. R. (2010). "Critical brain networks." *Physica A*
3. Tkacik, G. et al. (2015). "Thermodynamics and signatures of criticality in a network of neurons." *Nat. Phys.*
4. Ellis, R. S. (1985). *Entropy, Large Deviations, and Statistical Mechanics.* Springer.
5. Baxter, R. J. (1982). *Exactly Solved Models in Statistical Mechanics.* Academic Press.
6. Koch, C. (1999). *Biophysics of Computation.* Oxford University Press.
7. Newman, M. E. J. (2002). "Spread of epidemic disease on networks." *Phys. Rev. E*
8. Mézard, M., Parisi, G. & Virasoro, M. A. (1987). *Spin Glass Theory and Beyond.* World Scientific.
9. Stanley, H. E. (1971). *Introduction to Phase Transitions and Critical Phenomena.* Oxford University Press.

10. Goldenfeld, N. (1992). *Lectures on Phase Transitions and the Renormalization Group.* Westview Press.
11. Hohenberg, P. C. & Halperin, B. I. (1977). "Theory of dynamic critical phenomena." *Rev. Mod. Phys.*
12. Maass, W., Natschläger, T. & Markram, H. (2002). "Real-time computing without stable states." *Neural Comput.*
13. Shew, W. L., Yang, H., Yu, S., Roy, R. & Plenz, D. (2011). "Information capacity and transmission in a cortical network." *J. Neurosci.*
14. Touboul, J. & Destexhe, A. (2017). "Can power-law scaling and neuronal avalanches arise from stochastic interference?" *PLoS Comput. Biol.*
15. Wilting, J. & Priesemann, V. (2018). "Inferring collective dynamical states from widely unobserved systems." *Nat. Commun.*
16. Castellano, C., Fortunato, S. & Loreto, V. (2009). "Statistical physics of social dynamics." *Rev. Mod. Phys.*