The Persistence Functional: A Mathematical Measure of Attractor Resilience

Author: Robert Galida
Date: May 2026

Abstract

The attractor framework says that **persistence under disturbance** is the basic mark of reality.
To turn this idea into a formal science, we introduce the **persistence functional** $P(x)$.

$P(x)$ is a single number that measures:

- How deep a state is inside an attractor basin.
- How quickly it returns after a knock.

We define $P(x)$ for three different kinds of systems:

1. **Deterministic dissipative systems** – here P is linked to Lyapunov exponents and basin stability.
2. **Stochastic systems** – here P is linked to escape time and quasipotential.
3. **Information-theoretic systems** – here P is linked to negative free energy or mutual information.

The **recovery rate** $-\dot{P}/P$ is a universal sign of **critical**

slowing down – a warning that a system is about to tip.

We also discuss limitations: resilience may depend on direction (“anisotropic”), and multiple timescales may need **vector** or **tensor** persistence. We list open mathematical problems.
This paper is a **roadmap**, not a finished theory.

---

1. Introduction

In the attractor framework, **persistence under disturbance** is central. But we have not had a single number to say *how persistent* a state is.
The **persistence functional** $P(x)$ aims to fill that gap.

What $P(x)$ should do:

- $P(x)>0$ for states inside an attractor basin.
- For a **conservative attractor** (like a free electron), P is maximal (normalised to 1).
- For a **dissipative attractor**, P drops after a disturbance and then recovers.
  The recovery rate $-\dot{P}/P$ equals:
    - the negative of the largest Lyapunov exponent (for deterministic systems)
    - the inverse return time (for stochastic systems)
    - the rate of information loss (for informational systems)
- P falls as the system approaches a **bifurcation**, giving early warning.

We do **not** give one universal formula. Instead, we give a **family** of definitions, each suited to a different type of system, all united by the same purpose — measuring resilience.

---

2. Deterministic Dissipative Systems

Consider a smooth system $\dot{x}=f(x)$ with a stable attractor A and its basin $B(A)$.

A natural candidate for $P(x)$ uses a **Lyapunov function** $V(x)$ — a kind of energy that always decreases inside the basin ($\dot{V}<0$).

We define:$$P(x)=1-\frac{V(x)-V_A}{V_{\max}-V_A}$$

This gives $P=1$ on the attractor and $P\to 0$ at the basin boundary.
Near the attractor, the recovery rate is related to the **largest Lyapunov exponent** λ_1: $-\dot{P}/P\approx-\lambda_1$

When the system approaches a tipping point, $\lambda_1\to 0^-$, so the recovery rate slows down — this is **critical slowing down**.

Conclusion: For deterministic systems, P can be built from a Lyapunov function. The recovery rate equals the negative of the largest Lyapunov exponent.

---

3. Stochastic Systems

When noise is present, persistence is about how long it takes to escape from the basin.
The **mean first passage time** $\tau(x)$ — the average time to leave — is a natural measure.

We define: $$P(x)=\frac{\tau(x)}{\tau_{\max}}$$

where $\tau_{\max}$ is the value at the attractor.

For weak noise, $\tau(x)$ grows exponentially with the **quasipotential** $U(x)$ (Freidlin–Wentzell theory): $$\tau(x)\sim e^{U(x)/\epsilon}$$

So: $$P(x)\propto e^{-(U_{\max}-U(x))/\epsilon}$$

The recovery rate is the inverse of the return time. As a tipping point is approached, the return time diverges, and the recovery rate goes to zero. This again gives **critical slowing down** — rising variance and autocorrelation.

Conclusion: For stochastic systems, P is proportional to the mean exit time (or the exponential of the quasipotential). This connects persistence to large deviation theory.

---

4. Information-Theoretic Systems

For systems where information matters (neural, cognitive, social), we can define persistence using **mutual information** between past and future.

Let $I_{\text{past,future}}$ be the **predictive information**. Then: $$P(t)=I(\text{past;future at time } t)\quad\text{or}\quad P=e^{-\text{surprisal}}$$

The decay of $P(t)$ over time measures **memory loss**.
Landauer’s principle connects information loss to entropy production: $$\dot{P}/P\le-\frac{\dot{S}}{k_B\ln 2}$$

Alternatively, in the **free energy principle** (Friston), the negative free energy $-F$ acts like a Lyapunov function. We can set: $$P=e^{-F/kT}\quad\text{or}\quad P=-F$$

Then $-\dot{P}/P$ is the rate of free energy minimisation, which slows near bifurcations.

Conclusion: For information-theoretic systems, P can be defined via mutual information decay or negative free energy, linking persistence to entropy production and predictive coding.

---

5. Unifying Recovery Rate and Critical Slowing Down

Across all types of systems, the **recovery rate** $\lambda_{\text{rec}} = -\dot{P}/P$ (just after a small disturbance) is a universal indicator:

- **Deterministic dissipative:** $\lambda_{\text{rec}} = |\lambda_1|$ (absolute value of the largest Lyapunov exponent)
- **Stochastic:** λ_{rec} = inverse of the return time, related to the quasipotential's curvature
- **Information-theoretic:** λ_{rec} = rate of free energy minimisation or information loss

As the system approaches a bifurcation, $\lambda_{\text{rec}} \to 0$. This is **critical slowing down**.
It shows up as rising lag-1 autocorrelation and variance (Scheffer et al., 2009).
So P and its recovery rate give early warnings.

---

6. Normalisation for Conservative

Attractors

For a perfect **conservative attractor** (e.g., an electron in its ground state, no decay), the persistence functional should be constant and maximal: $P_{\text{cons}} = 1 \text{ for all times}$

No recovery rate is defined (or it is zero). This anchors the scale.

For **emergent approximate conservative systems** (like atomic clocks), P is very close to 1 and decays extremely slowly.

---

7. Limitations – Scalar Collapse and Anisotropic Resilience

A single scalar $P(x)$ may not be enough for systems where resilience is **anisotropic** – that is, recovery speed depends on the direction of the perturbation.
High-dimensional systems can have **multiple timescales** (fast and slow modes). A scalar average can miss important structure.

Future work may need:

- **Vector persistence** – a list of recovery rates along different directions.
- **Tensor persistence** – a metric that captures the full shape of the basin.
- **Persistence manifold** – the geometry of the basin in state space.

We accept this limitation. The scalar P is a useful first approximation for systems with isotropic resilience or for early-warning applications where a single number is enough.

For complex systems, a multidimensional generalisation is an open research problem.

---

8. Open Mathematical Problems

1. **Derive $P(x)$ from first principles** for a given class of systems (e.g., from a variational principle).
2. **Prove that $-\dot{P}/P=\langle\lambda_1\rangle$** for a wide class of dissipative systems.
3. **Extend the definition to systems with multiple attractors and chaotic basins** (where basin stability is fractal).
4. **Establish a rigorous relationship between P and the mutual information decay rate** for non-equilibrium processes.
5. **Formulate a universal persistence functional** that works across all regimes — or prove it’s impossible.
6. **Test the predictive power of P** in controlled experiments (e.g., ecological microcosms, neural cultures, social media sentiment).
7. **Develop vector/tensor persistence** for anisotropic resilience.

---

9. Conclusion

The persistence functional $P(x)$ gives a mathematical language for attractor resilience.

We have given **operational definitions** for three regimes:

- **Deterministic dissipative** → Lyapunov / basin stability

- **Stochastic** → escape time / quasipotential
- **Information-theoretic** → mutual information / free energy

The **recovery rate** $-\dot{P}/P$ unifies critical slowing down across all these domains.

We have explicitly noted **limitations** (scalar collapse, anisotropy) as open problems.

This paper is a **roadmap**, not a final theory. The framework now has a quantitative step.

Suggested citation: Galida, R. S. (2026). *The Persistence Functional: Towards a Mathematical Measure of Attractor Resilience (Reader-Friendly Version).* Fantasy Attractor.