The Bhasin Invariance Principle

Entropy-Rate Conservation in Coarse-Grained Dynamical Systems

Ankur Bhasin
Bhasin Research Unit for Hyperphysics (BRUH)

Abstract

We state a minimal, falsifiable operational principle asserting entropy-rate conservation across coarse-grained dynamical systems. Multichannel time series are mapped into a frozen discrete coherence–information state process governed by a single time-homogeneous generator. Across stable operating regimes, the entropy rate of this generator concentrates near a domain-invariant target value, while architectural differences manifest only in relaxation structure. The principle is stated in irreducible form and is intended to be attacked directly. If false, it should fail under the explicit breakdown conditions listed here.

1 State representation

Let $x(t) \in \mathbb{R}^d$ denote a multichannel observed time series sampled at times $t_n = n\Delta t$. Define a fixed feature map

$$X_n = \Phi(x[t_n - T,\, t_n]) \in \mathbb{R}^m, \tag{1}$$

where X_n represents a coarse-grained coherence–information state derived from a sliding observation window. The mapping Φ is fixed once and then held unchanged across datasets and domains.

Frozen discretization

Let $Q : \mathbb{R}^m \to \{1, 2, \ldots, K\}$ denote a fixed discretizer, yielding a discrete state process

$$S_n = Q(X_n). \tag{2}$$

The pair (Φ, Q) is frozen after calibration and reused unchanged.

2 Frozen generator

The principle asserts that, to leading order, S_n is governed by a time-homogeneous first-order Markov generator with transition matrix P:

$$\mathbb{P}(S_{n+1} = j \mid S_n = i, \mathcal{F}_n) = P_{ij}, \qquad \sum_j P_{ij} = 1, \tag{3}$$

where $\mathcal{F}_n$ denotes the full past history. Conditioning on deeper history does not materially improve prediction under the frozen construction.

Define a scalar instability or risk score

$$R_n = \mathcal{R}(S_n; P), \tag{4}$$

where $\mathcal{R}$ is a fixed functional of the current state and generator.

3 Invariant quantity

Let π denote the stationary distribution of P. The entropy rate of the generator is

$$H(P) = -\sum_i \pi_i \sum_j P_{ij} \log P_{ij}. \tag{5}$$

Invariant claim. Across domains and datasets, the frozen construction yields generators whose entropy rate concentrates near a domain-invariant constant H^*:

$$|H(P) - H^*| \leq \varepsilon_H, \tag{6}$$

while other structural quantities (for example spectral gap, relaxation time, or dwell distributions) vary in architecture-dependent ways.

4 Breakdown and falsification

The principle is falsified if any of the following occur:

- Surrogate data destroying temporal or cross-channel structure reproduces the same invariance.
- A single frozen construction fails to generalize across domains without retuning.
- Conditioning on deeper history materially improves prediction beyond first order.
- Decision thresholds require dataset-specific adjustment to maintain performance.

5 Operational implication

Given a downstream observable failure sequence Y_n, the principle implies that a fixed rule based on R_n can flag an upcoming transition prior to Y_n with nonzero lead time, and that this advantage vanishes under surrogate destruction.

Disclosure scope

This document states the principle only. Feature constructions, discretization schemes, risk functionals, calibration protocols, and code implementations are intentionally not disclosed here and will appear, if at all, in separate application or prediction papers.

Status

This invariance principle constitutes the core law underlying the broader VUH framework but stands independently of any particular domain realization.