The Symmetry of a Mind: How Tracking the World Shapes the Brain's Dynamics

An exploration of structured dynamics in algorithmic agents.

Based on the work of Giulio Ruffini (TN0364, 2023)

Ruffini 2025, Entroy, https://www.mdpi.com/1099-4300/27/1/90

High-dimensional data isn't random. It lives on surprisingly simple, low-dimensional structures.

Natural data, like images or neural recordings, exists in extraordinarily high-dimensional spaces (e.g., millions of pixels, billions of neurons). Empirically, this data does not fill the space. It consistently clusters along smooth, low-dimensional structures called "manifolds."

This is the Manifold Hypothesis. It implies a deep, hidden organizing principle is at work.

Why? What force compels this dramatic reduction in complexity?

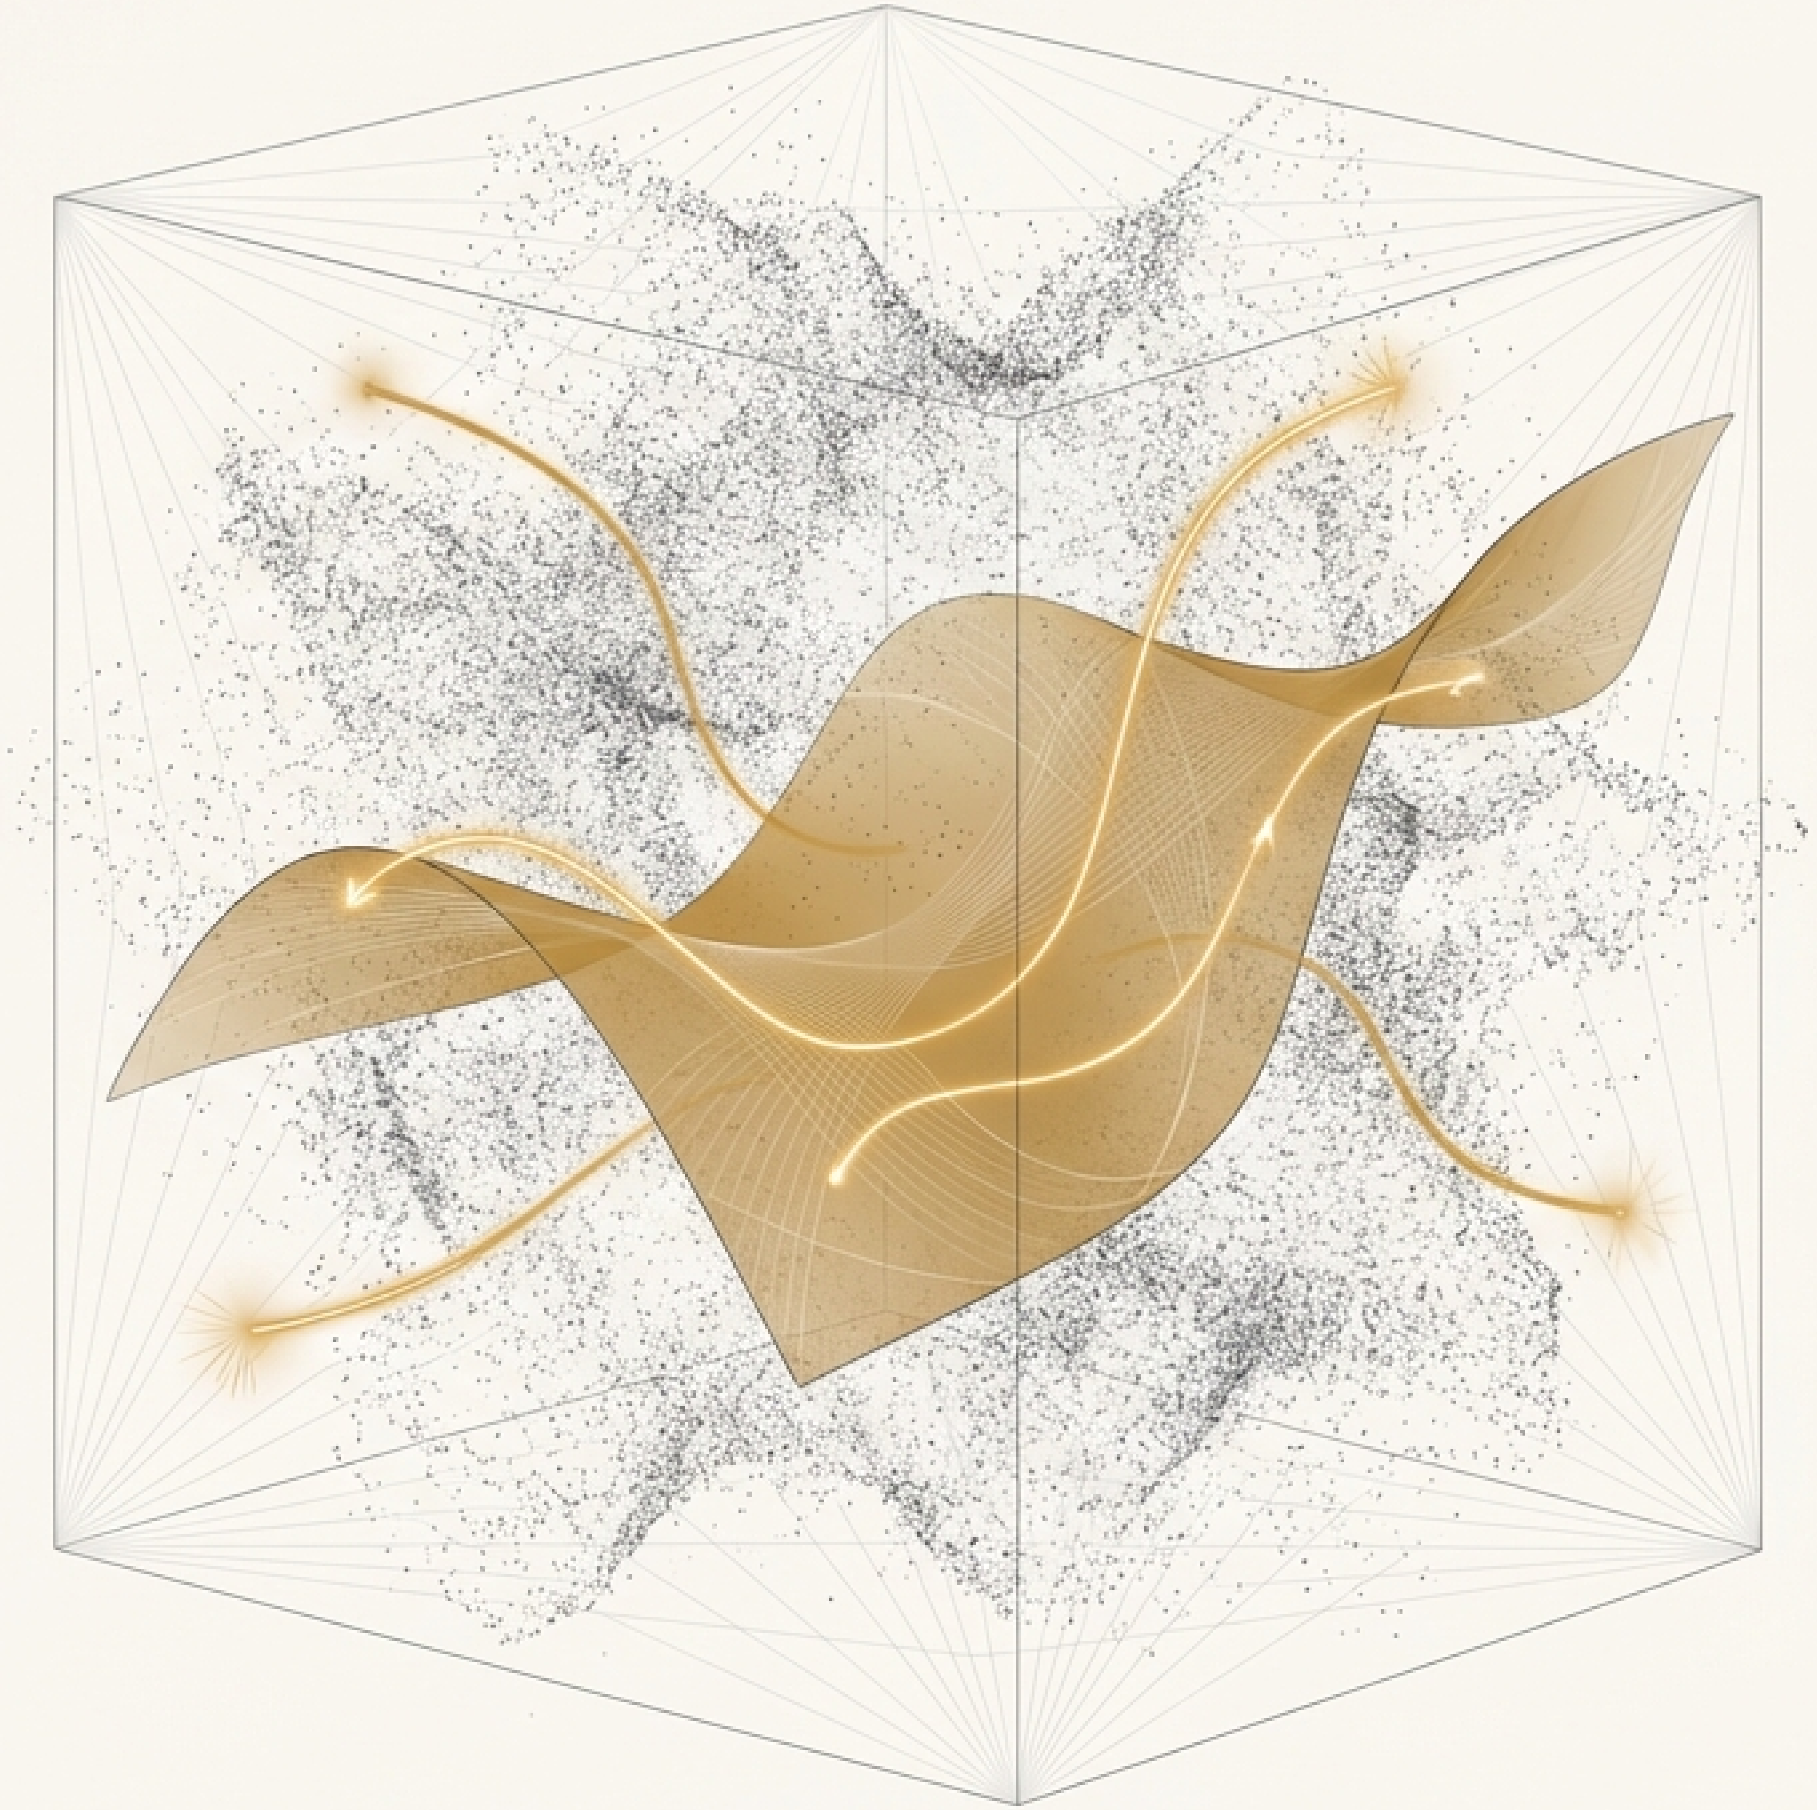

The agent's core function is to find **simple models** that **compress** complex world data.

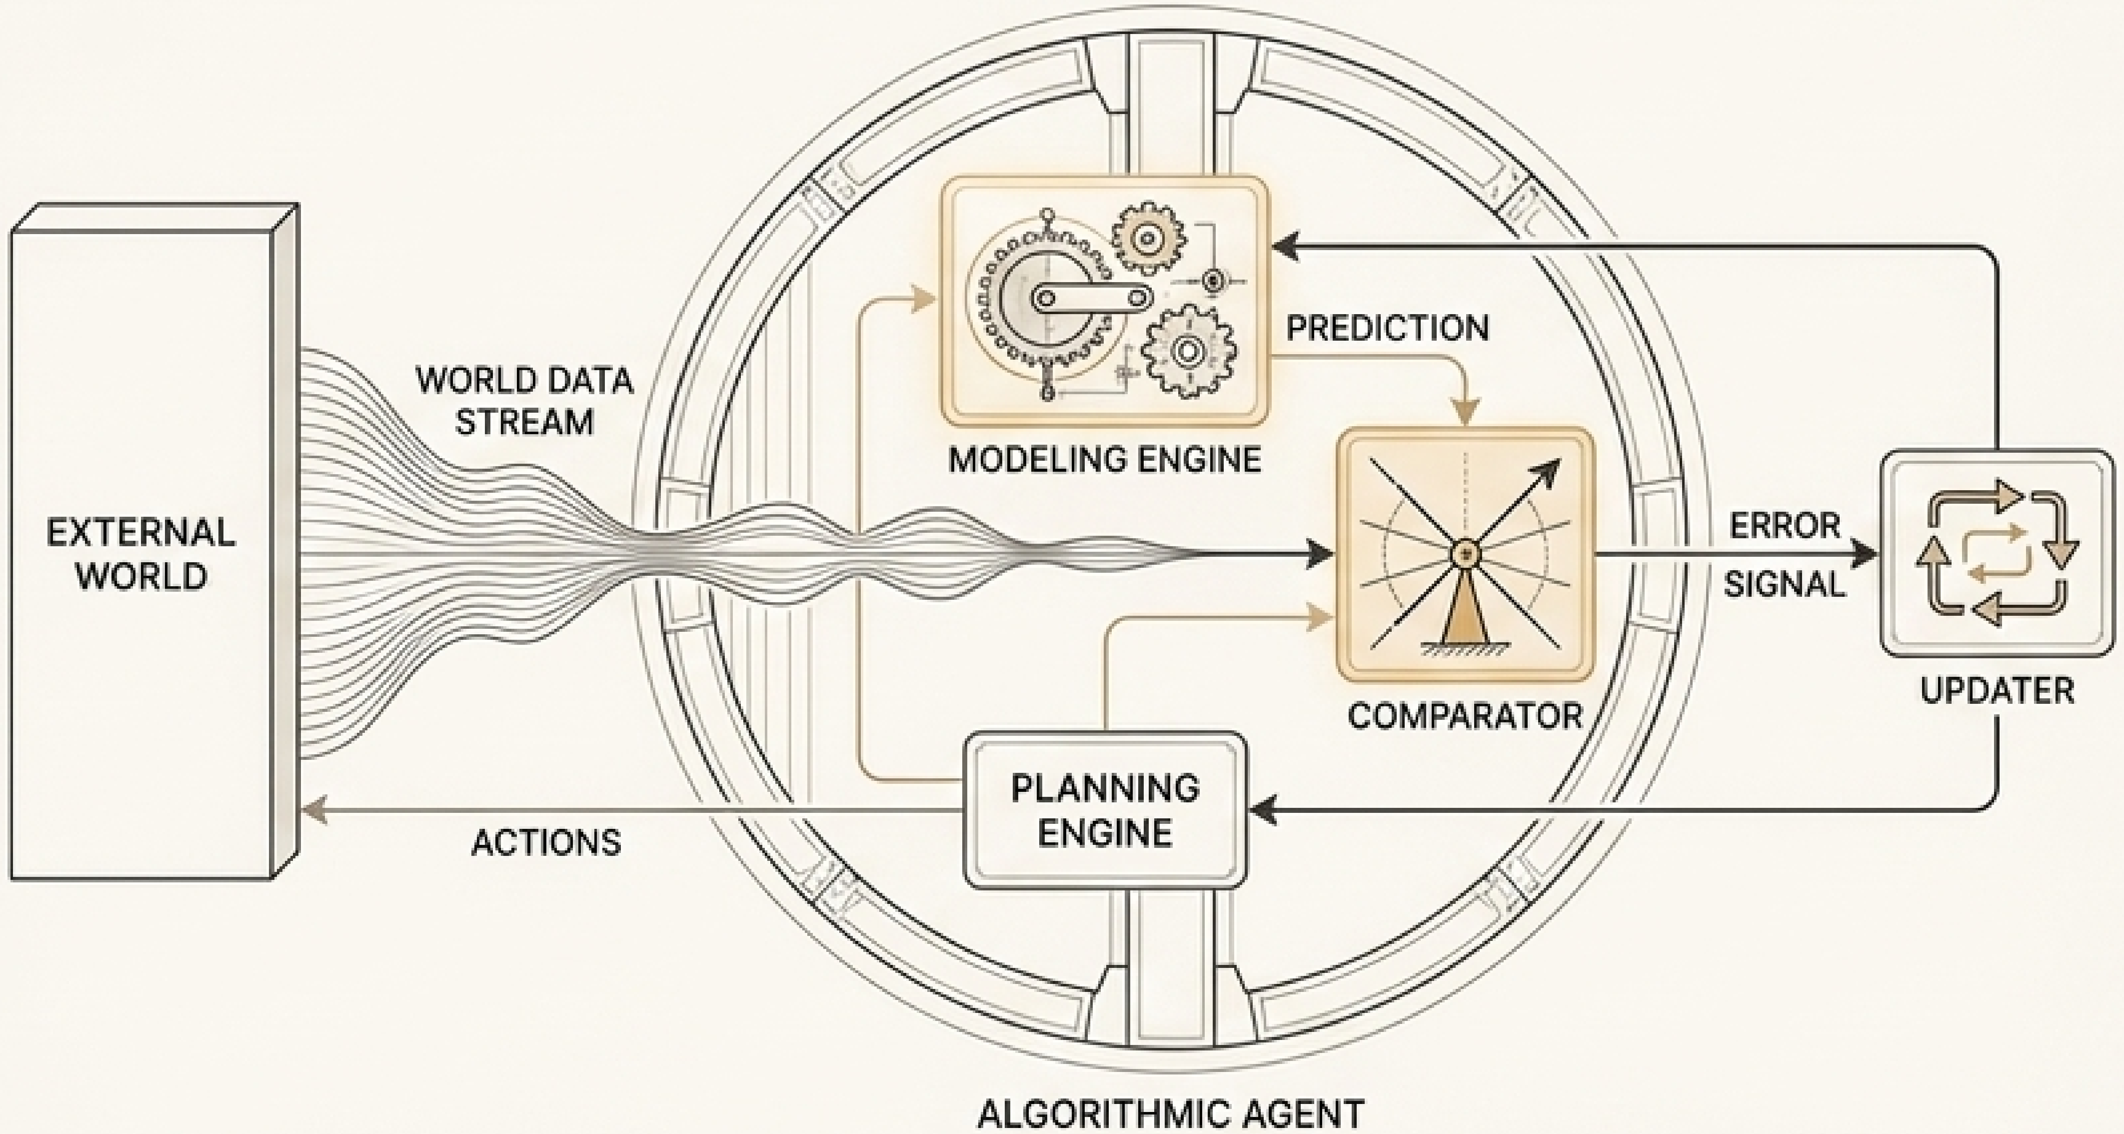

- We introduce the "Algorithmic Agent" from Kolmogorov Theory (KT).
- Its definition: An information-processing system with an objective function that interacts with the world by inferring and running compressive models to plan and act.
- The ability to compress information is equivalent to understanding. It means finding patterns and regularities to build a predictive model of the world.

Structure that can be compressed is, fundamentally, **symmetry**.

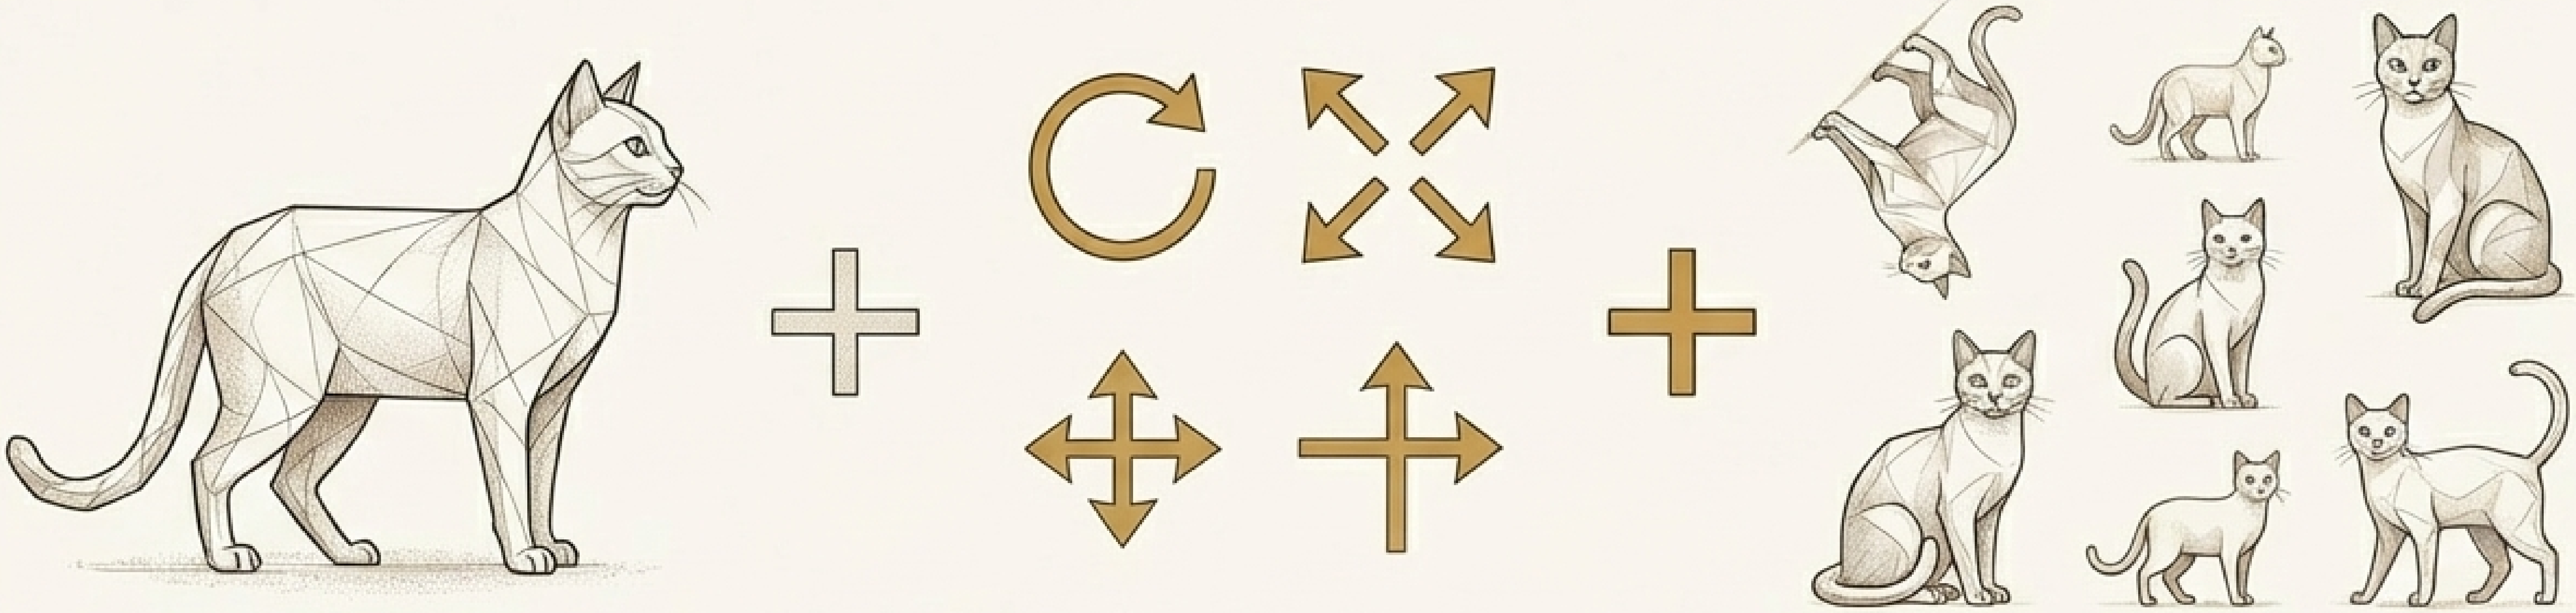

Canonical Archetype (`$I_0$`)   Transformations (`$\gamma$`)   Generated Dataset (`$I = \gamma \cdot I_0$`)

Compression works by exploiting repetition, patterns, and invariance. A model that successfully compresses data essentially reflects the underlying symmetries or invariances. For example, a vast set of "cat' images can be generated from one canonical 'archetype' cat plus a set of transformations (rotation, scaling, changes in pose). The 'cat-ness' is the invariant property. This invariance under transformation is the mathematical definition of symmetry.

We can use the language of Lie Groups—the mathematics of continuous symmetry—to formally define a "generative model."

- A **generative model** is a set of rules for creating data (e.g., all possible hand images).
- A **Lie group** is a mathematical object that describes continuous transformations (like rotations, scaling, or articulation of joints).
- **The Proposal (Definition 2.2):** A simple generative model *is* a Lie group. Any specific data instance (any cat image) can be generated by applying a group transformation γ to an arbitrary reference instance I_0.

$$I = \gamma \cdot I_0$$

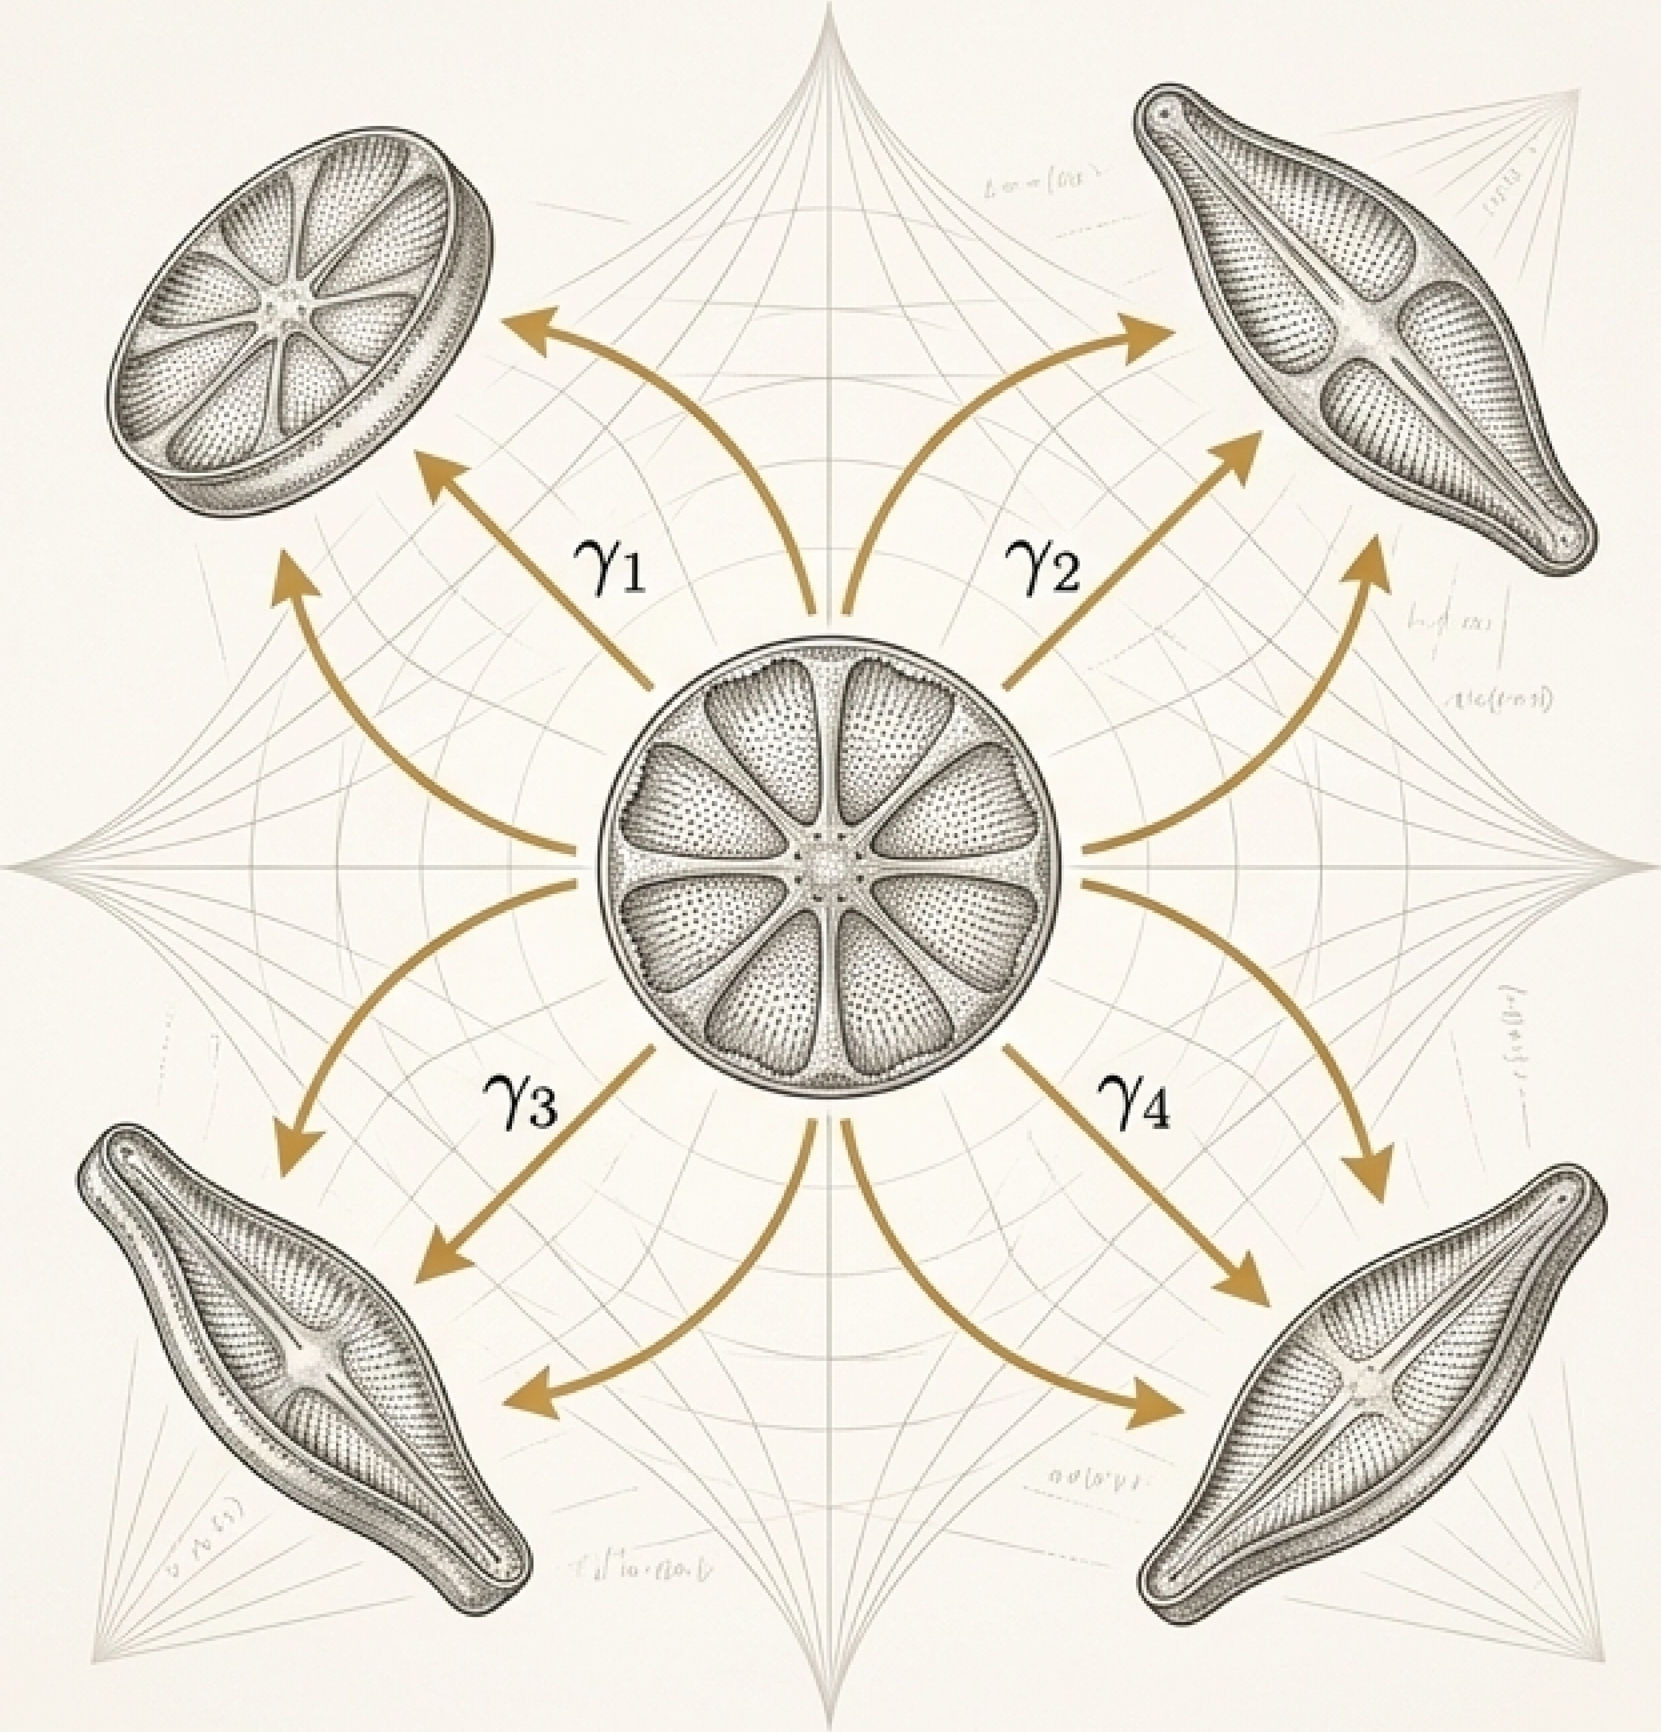

The complexity of a vast Lie group is captured by a small set of “generators”—the rules for infinitesimal steps.

Any complex transformation within a Lie group can be constructed by repeatedly applying a few simple, infinitesimal transformations (the group’s “generators”). This is deeply connected to algorithmic simplicity: a short program (the generators) can produce a vast, complex dataset. This recursive, compressive nature makes Lie groups the ideal language for describing simple, powerful world models.

$$g = \exp\left[\sum \theta_k T_k\right]$$

A full transformation g is generated by exponentiating a combination of infinitesimal steps T_k.

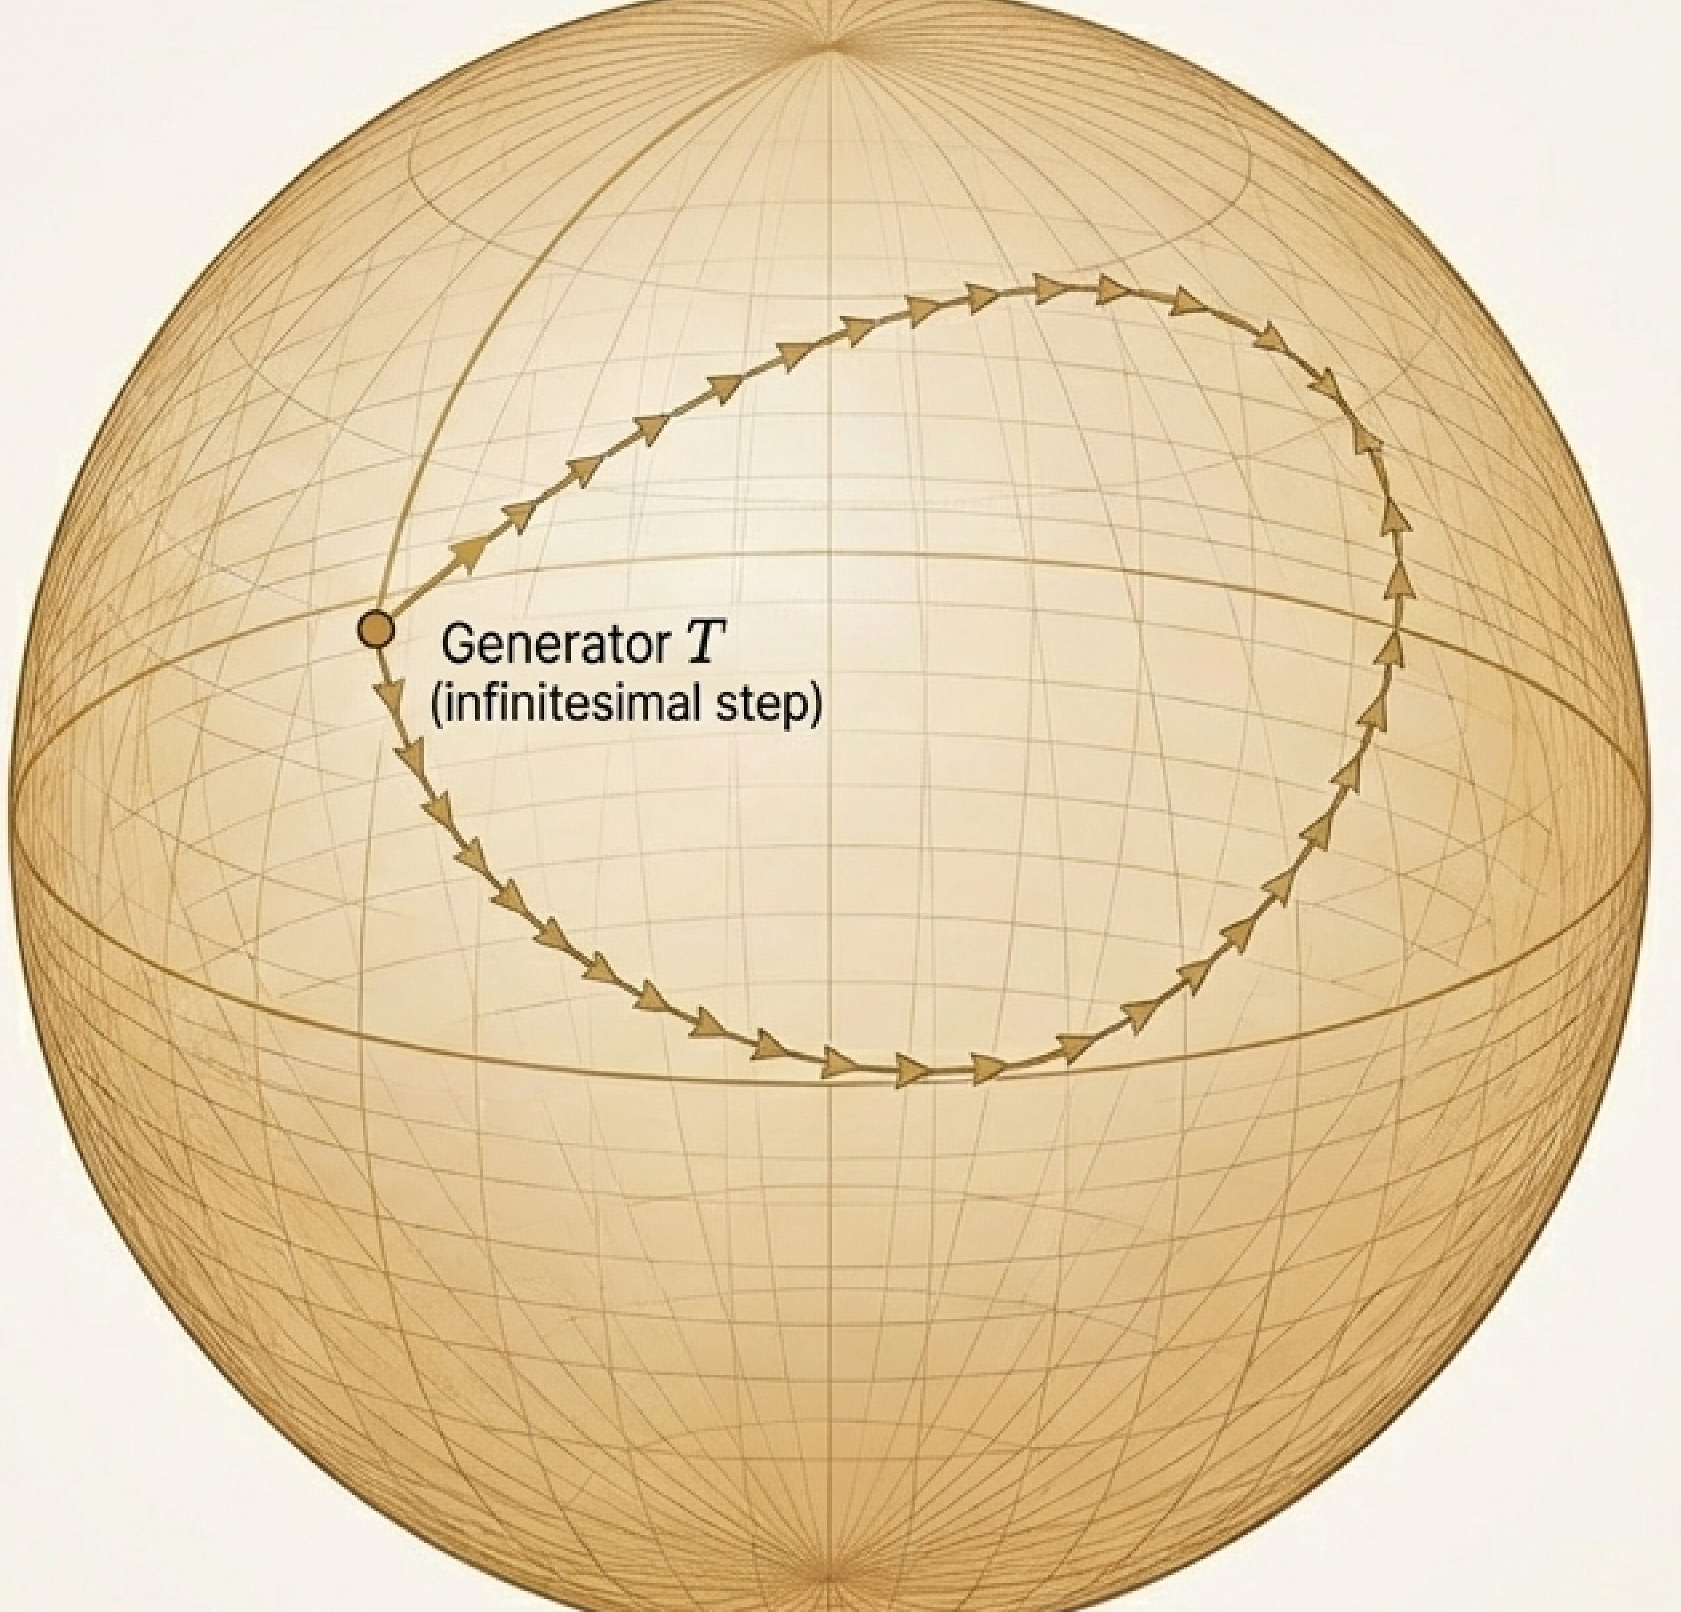

A successful agent must constrain its own dynamics to continuously match the world's data stream.

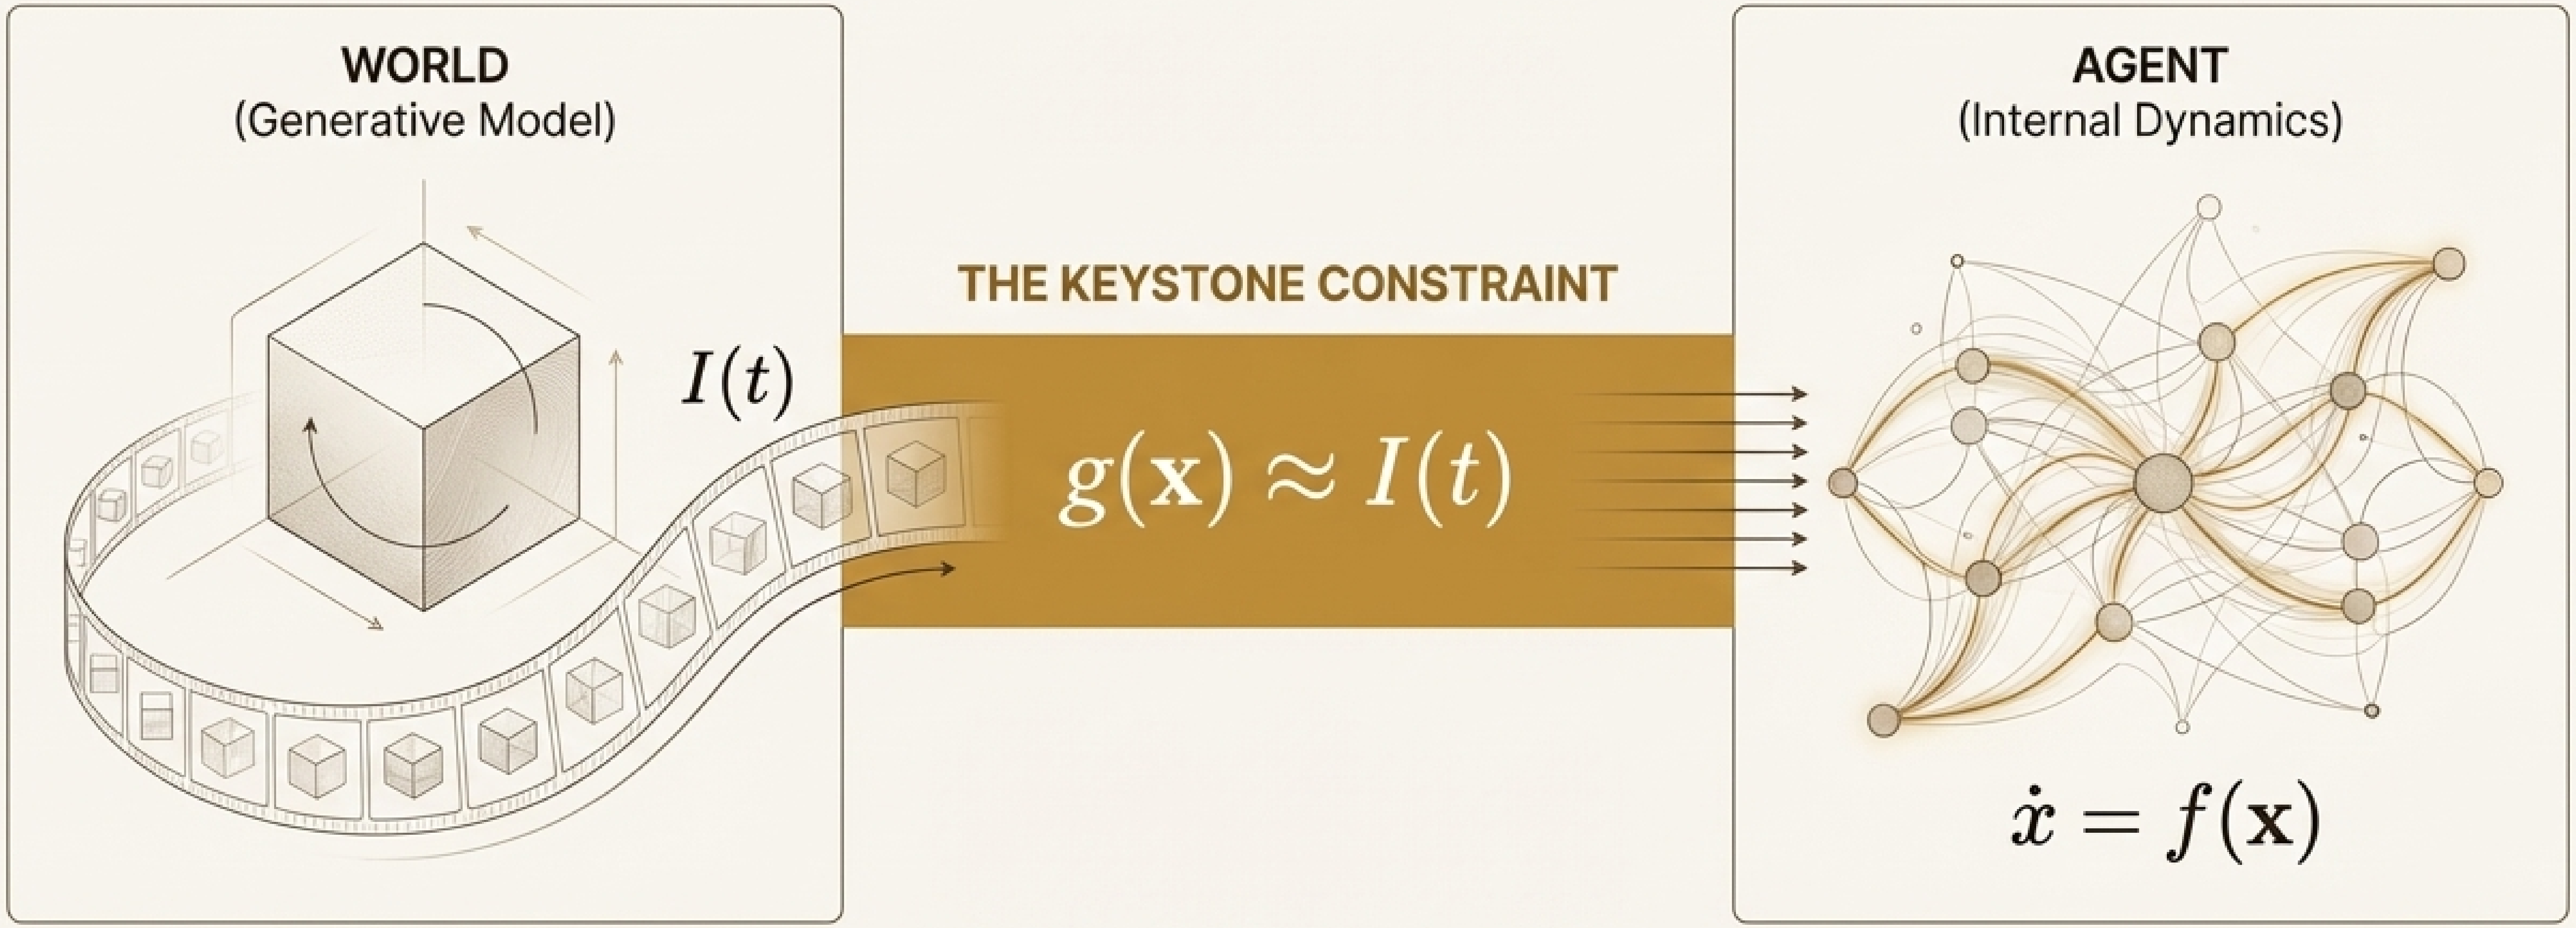

An agent isn't a passive observer; it has its own internal dynamics (e.g., neural activity described by $\dot{x} = f(\mathbf{x})$). The agent's 'Comparator' imposes a crucial constraint: a part of its internal state must mirror the sensory input. This creates a **Differential Algebraic Equation**: the agent's internal dynamics ($\dot{x} = f(\dots)$) are not free but are yoked to an algebraic conconstraint ($g(\mathbf{x}) \approx I(t)$).

The world-tracking constraint forces the agent's internal system to inherit the symmetries of the world model.

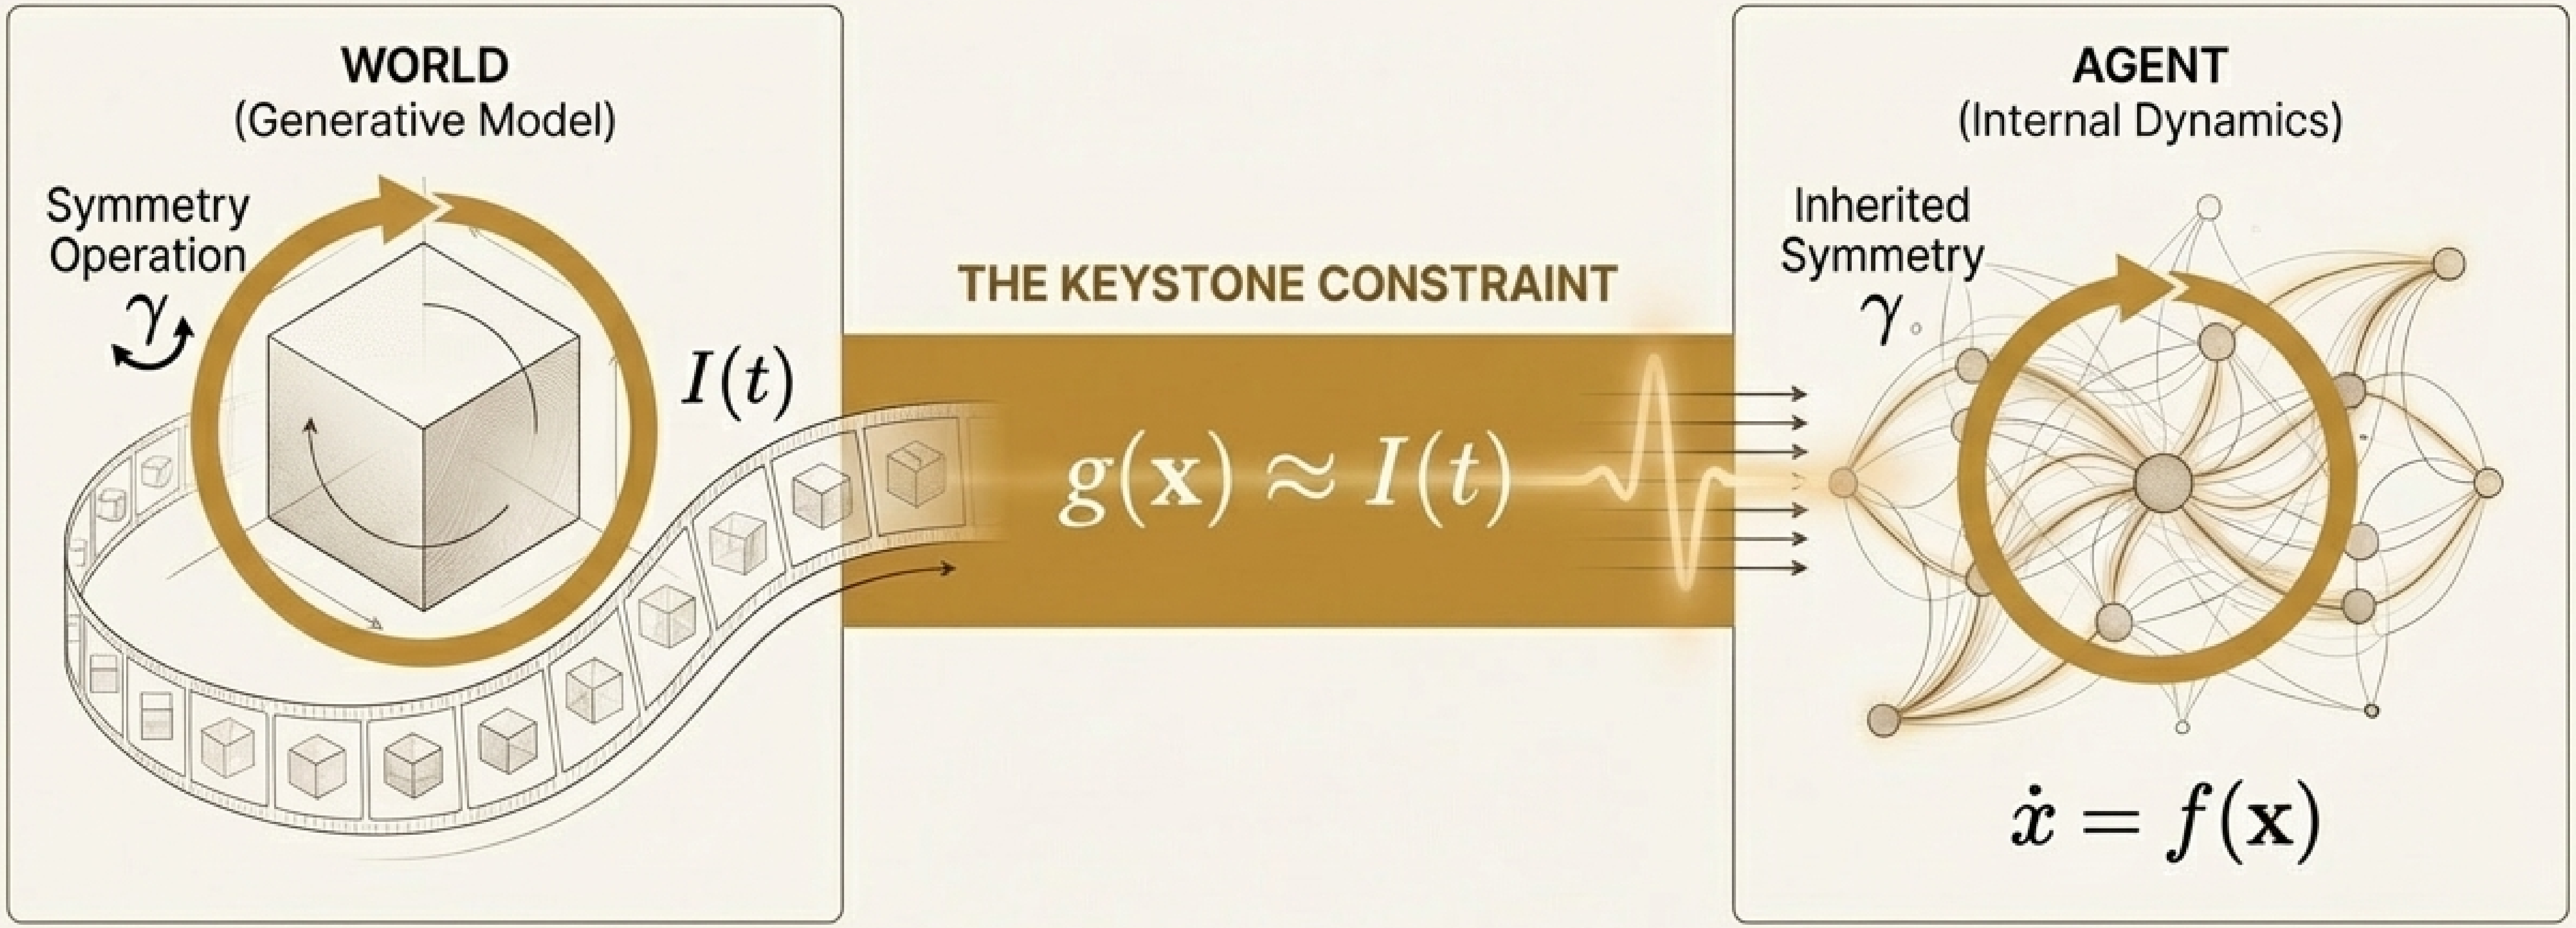

If the world's data is invariant under a transformation (e.g., rotating a hand results in another valid hand image), then a successful tracking agent must also be invariant to that change.

This is a mathematical necessity. For the constraint $g(\mathbf{x}) \approx I(t)$ to hold true across all possible hand images, the agent's internal rules ($f(\mathbf{x})$) and parameters ($\mathbf{w}$) must be structured to respect the hand's symmetries.

In short: the agent must carry the symmetries of the world's generative model.

In any dynamical system, symmetries give rise to conserved quantities—aspects of the system that remain constant.

- This is a fundamental principle in physics, articulated by Noether's Theorem.
- When an agent's dynamics inherit the world's symmetries, it effectively creates conservation laws within its neural state space.
- Example: For a static input $I(t) = C$, the world-tracking constraint $g(\mathbf{x}) \approx C$ becomes a set of conserved quantities for the agent's dynamics.
- The agent's internal trajectory is now confined to a subspace where these quantities $g(\mathbf{x})$ are constant.

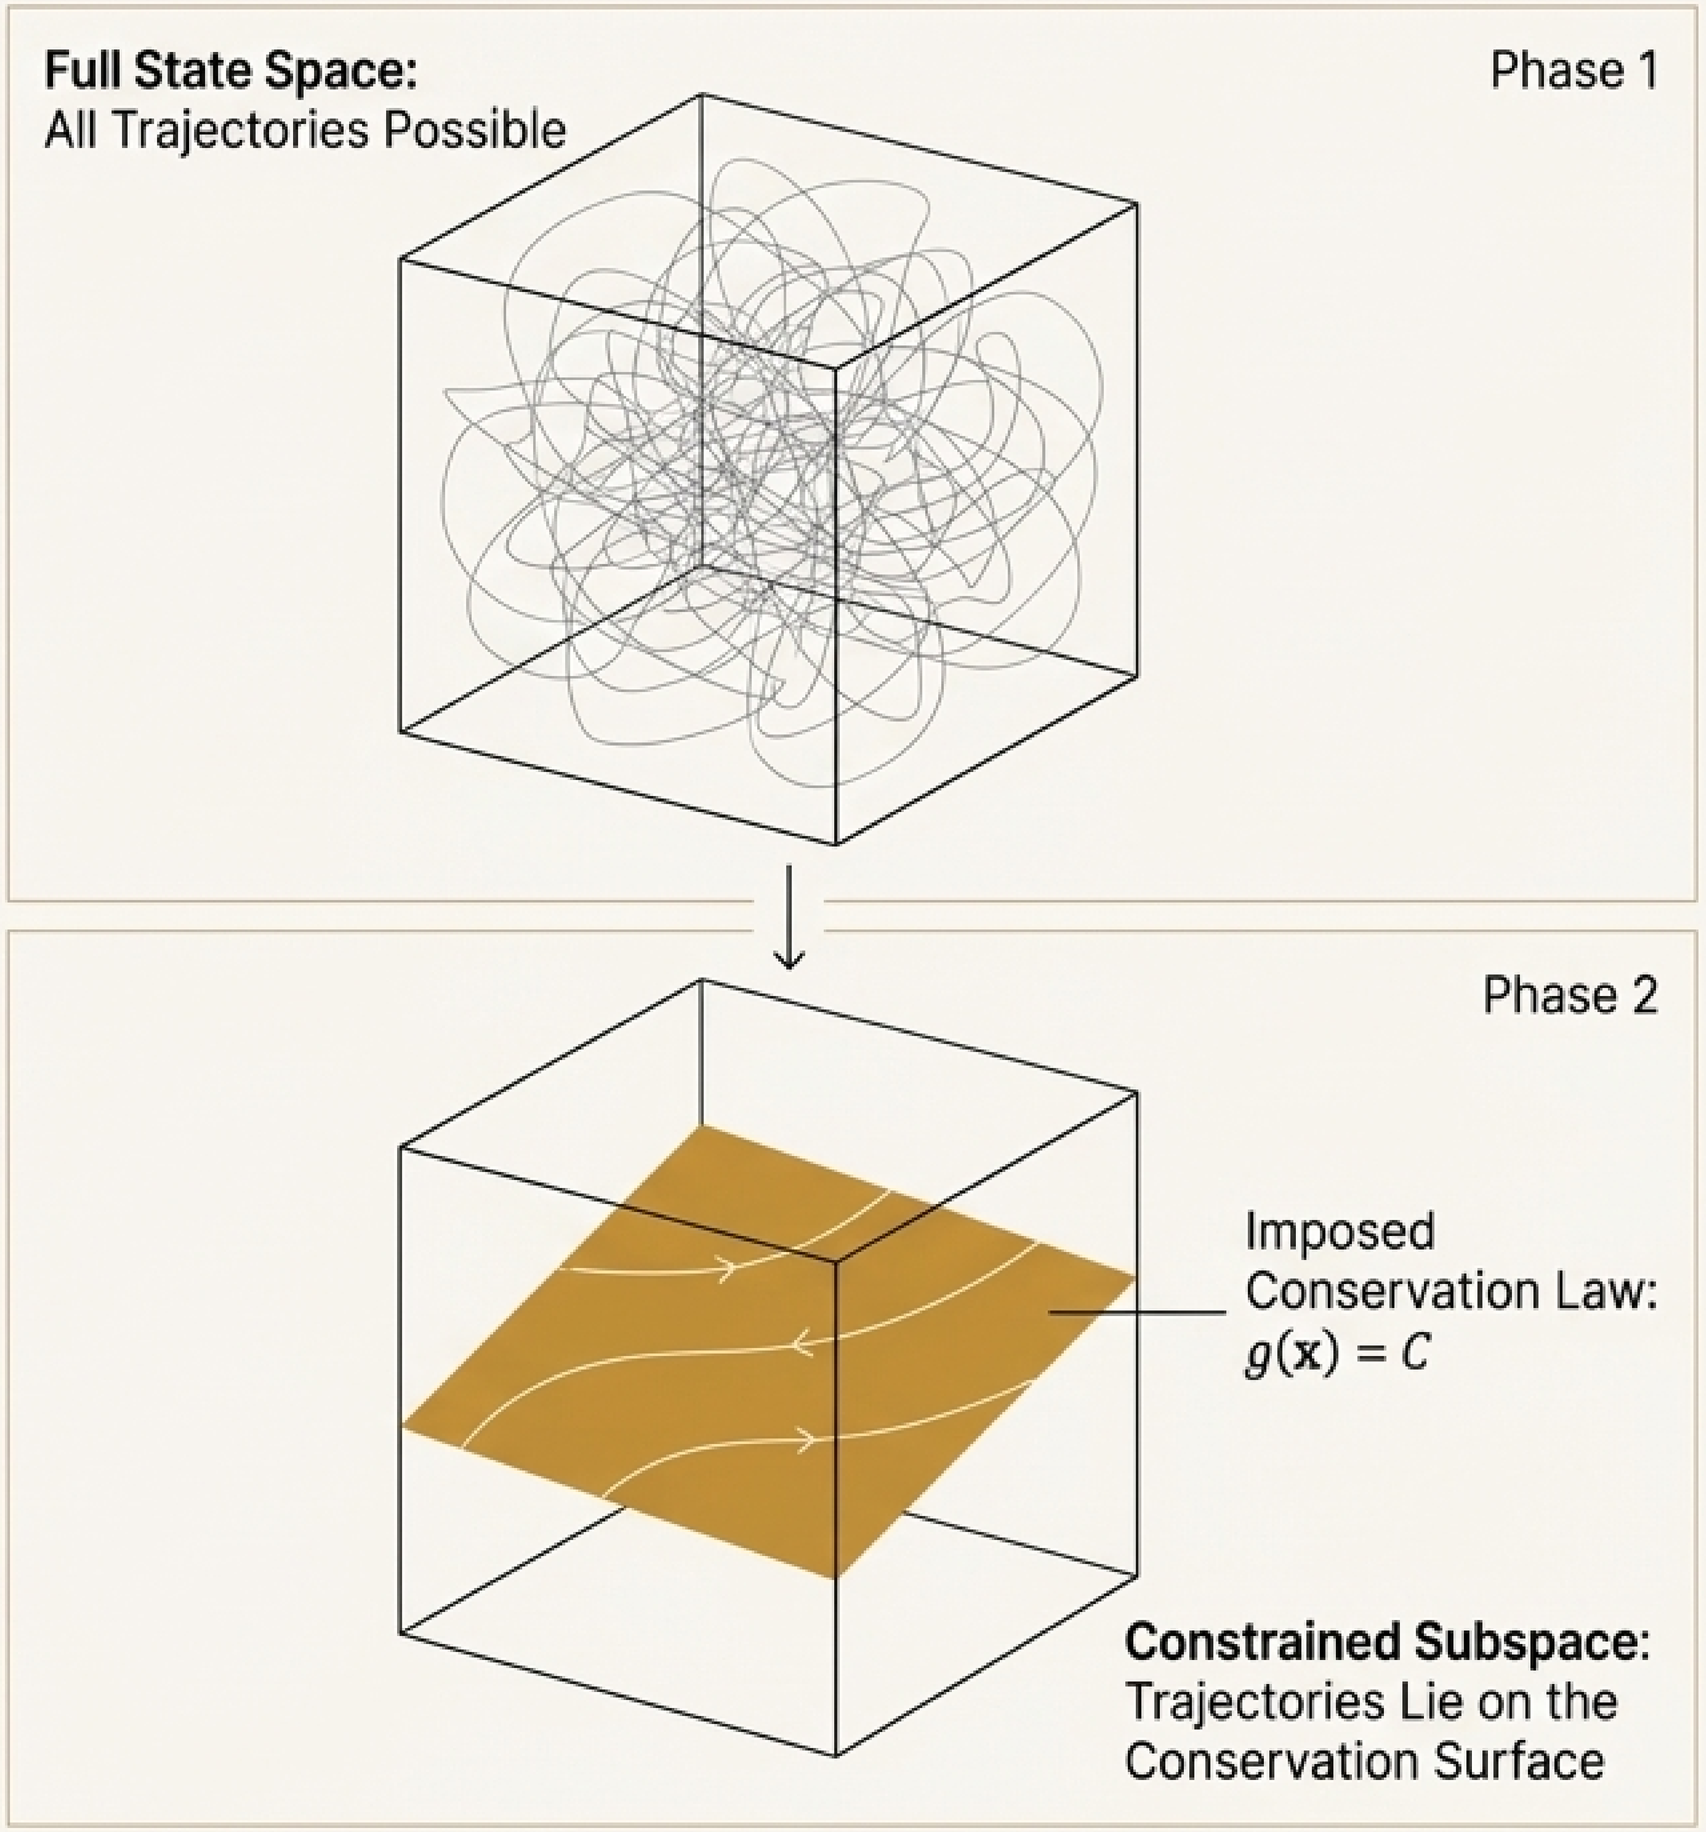

The low-dimensional manifold is the geometric expression of the conservation laws imposed by world-tracking.

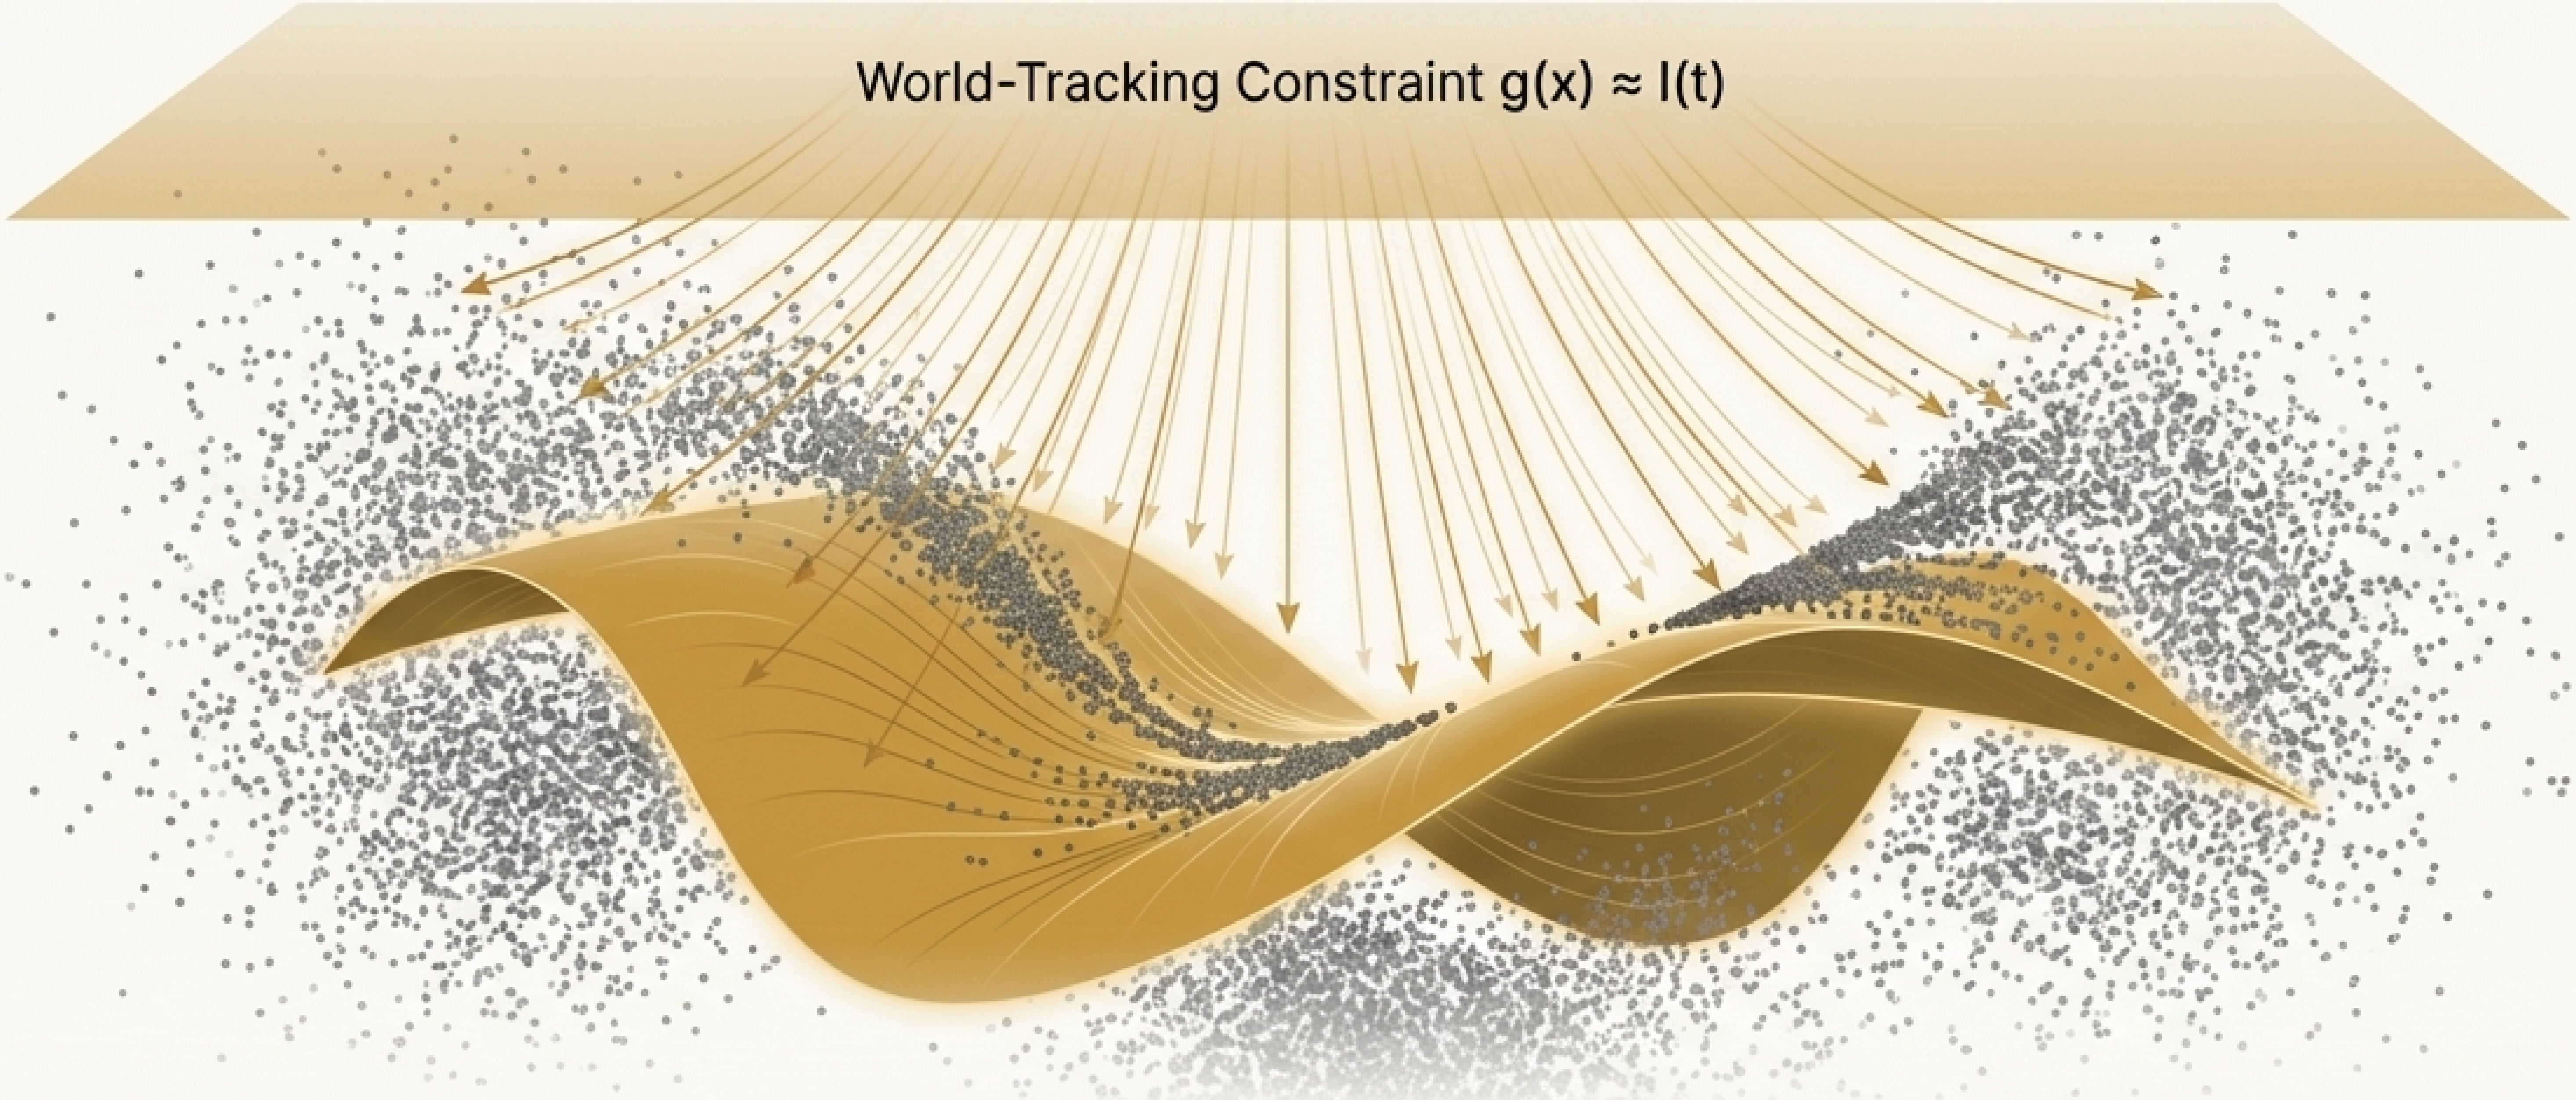

The thousands of constraints from tracking an image ($g(x) \approx I(t)$) act as a massive set of conservation laws for the billions of neurons in the brain. These laws force the system's trajectory out of the full, high-dimensional state space and onto a much smaller, highly constrained subspace. **This subspace is the invariant manifold.**

A simple chain of logic explains the emergence of structured dynamics.

1. The Mystery

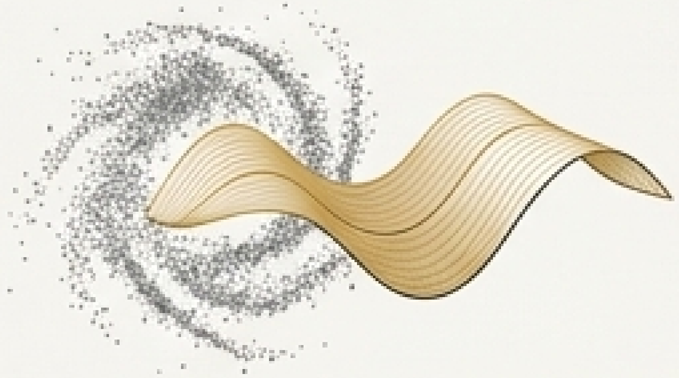

Natural high-dimensional data lives on simple, low-dimensional manifolds. *Why?*

2. The Clues

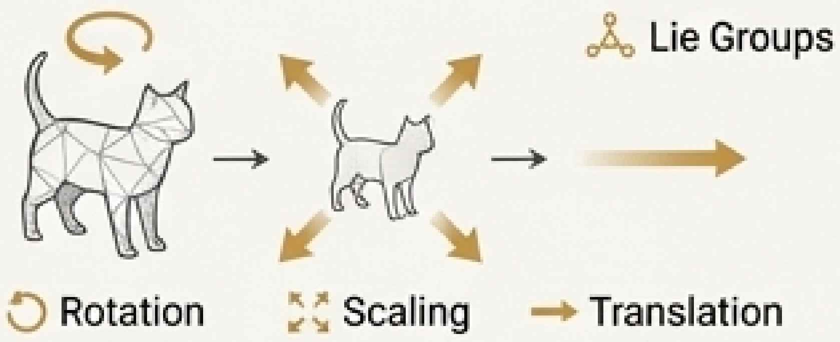

Agents are compressors. Compression is fundamentally about identifying symmetry, which is mathematically described by Lie groups.

3. The Mechanism

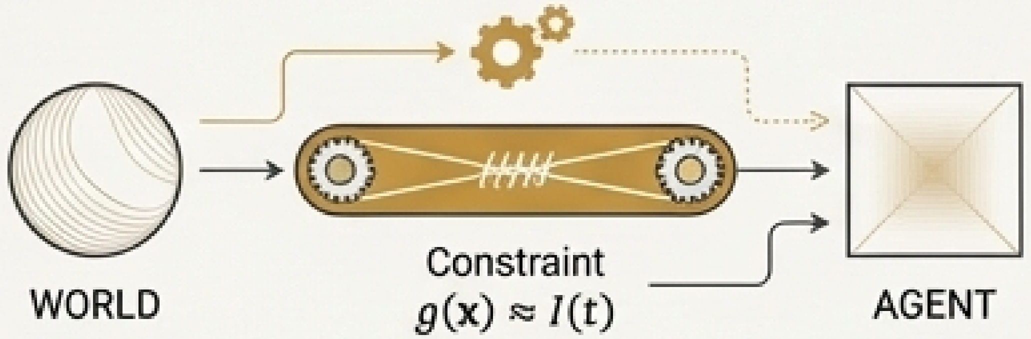

To track the world, an agent's dynamics must be constrained ($g(\mathbf{x}) \approx I(t)$). This forces the agent to inherit the world's symmetries.

4. The Resolution

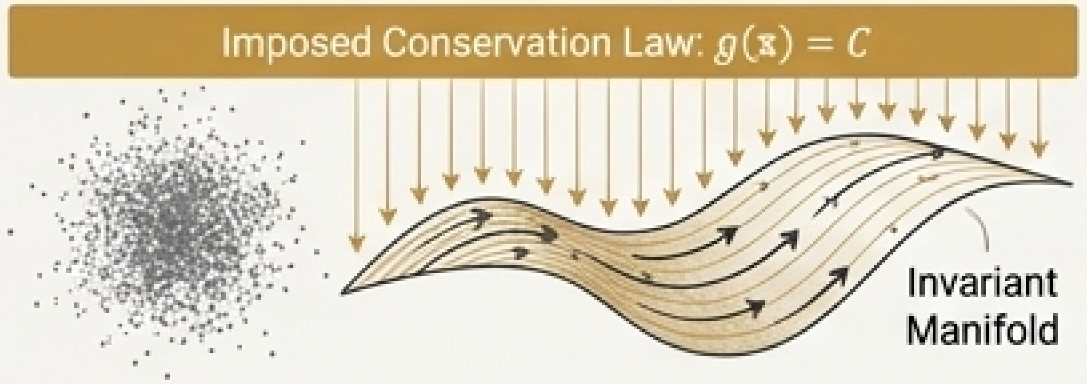

Inherited symmetries create internal conservation laws. These laws restrict the agent's dynamics to a subspace, which *is* the low-dimensional manifold.

This framework offers a new, principled way to interpret brain data and states of consciousness.

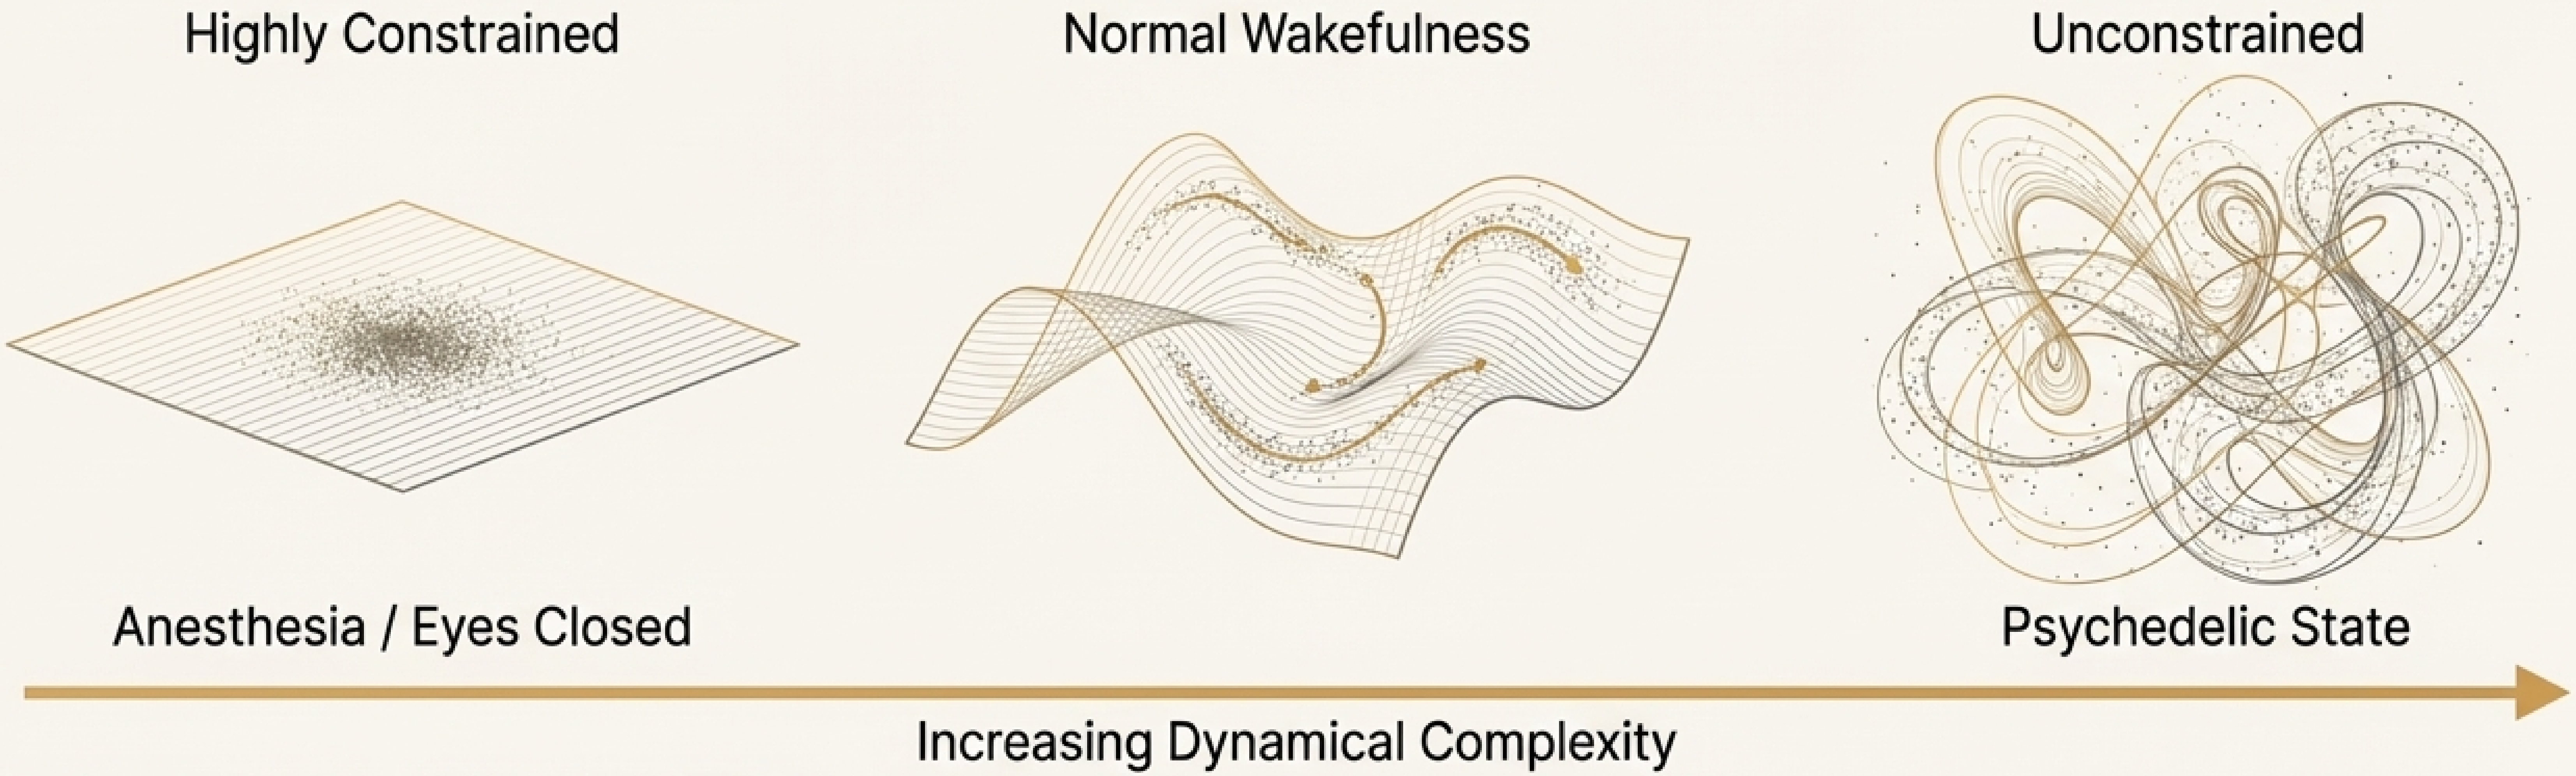

Explaining Brain State Complexity:

Highly Constrained States (Eyes Closed, Anesthesia): Sensory input is clamped. The world-tracking constraint becomes fixed, leading to highly constrained, **lower-dimensional** (less complex) dynamics, like **alpha rhythms.**

Unconstrained States (Psychedelics): World models are disrupted. Constraints are lifted, allowing dynamics to explore a much **larger, higher-dimensional** (more complex) space.

This provides a theoretical basis for using manifold analysis on neuroimaging data (EEG, fMRI).

We can design more efficient and robust AI by building symmetry directly into its architecture.

Brute Force Learning

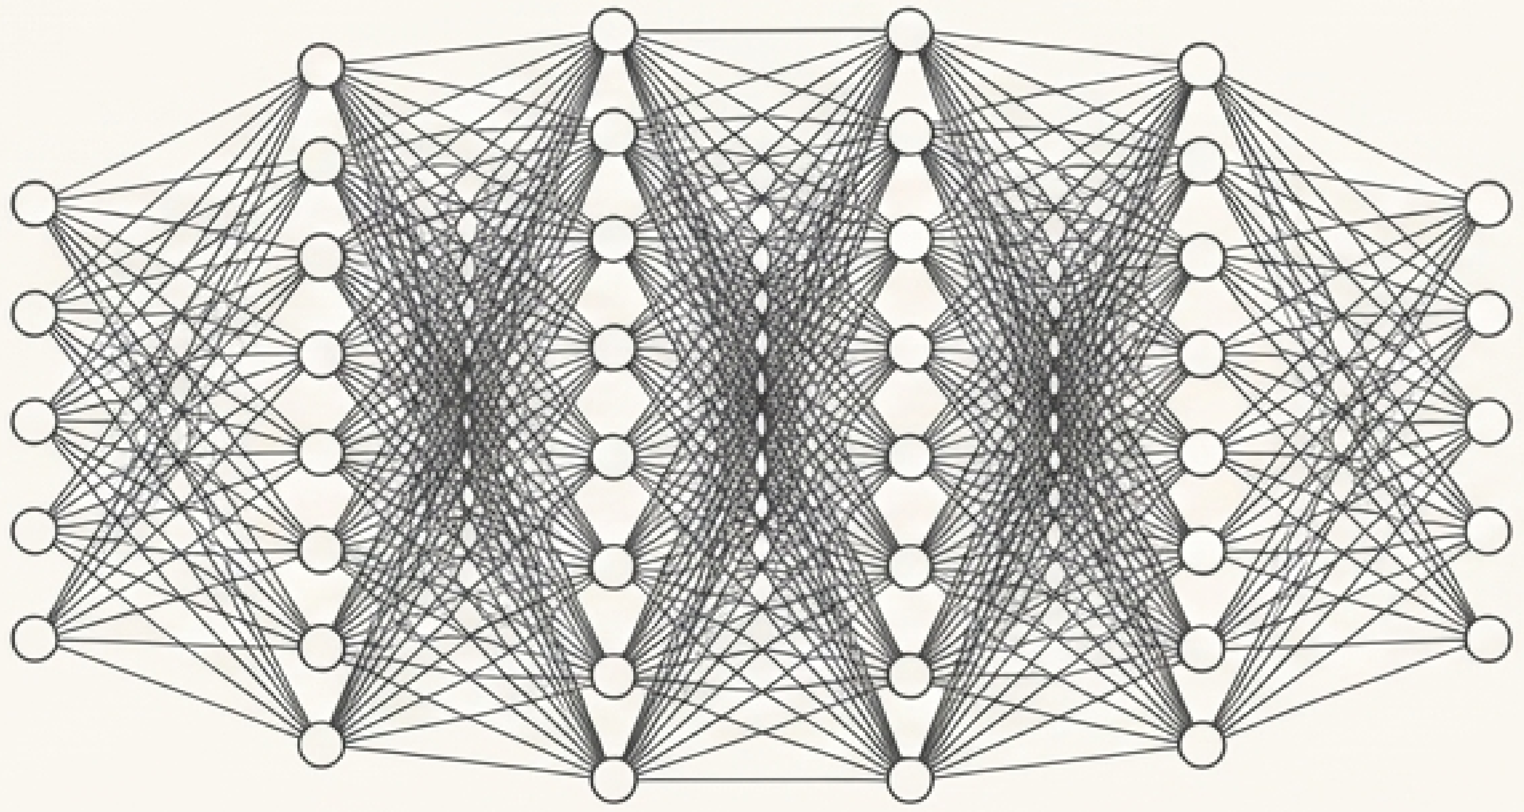

Symmetry-Aware Design

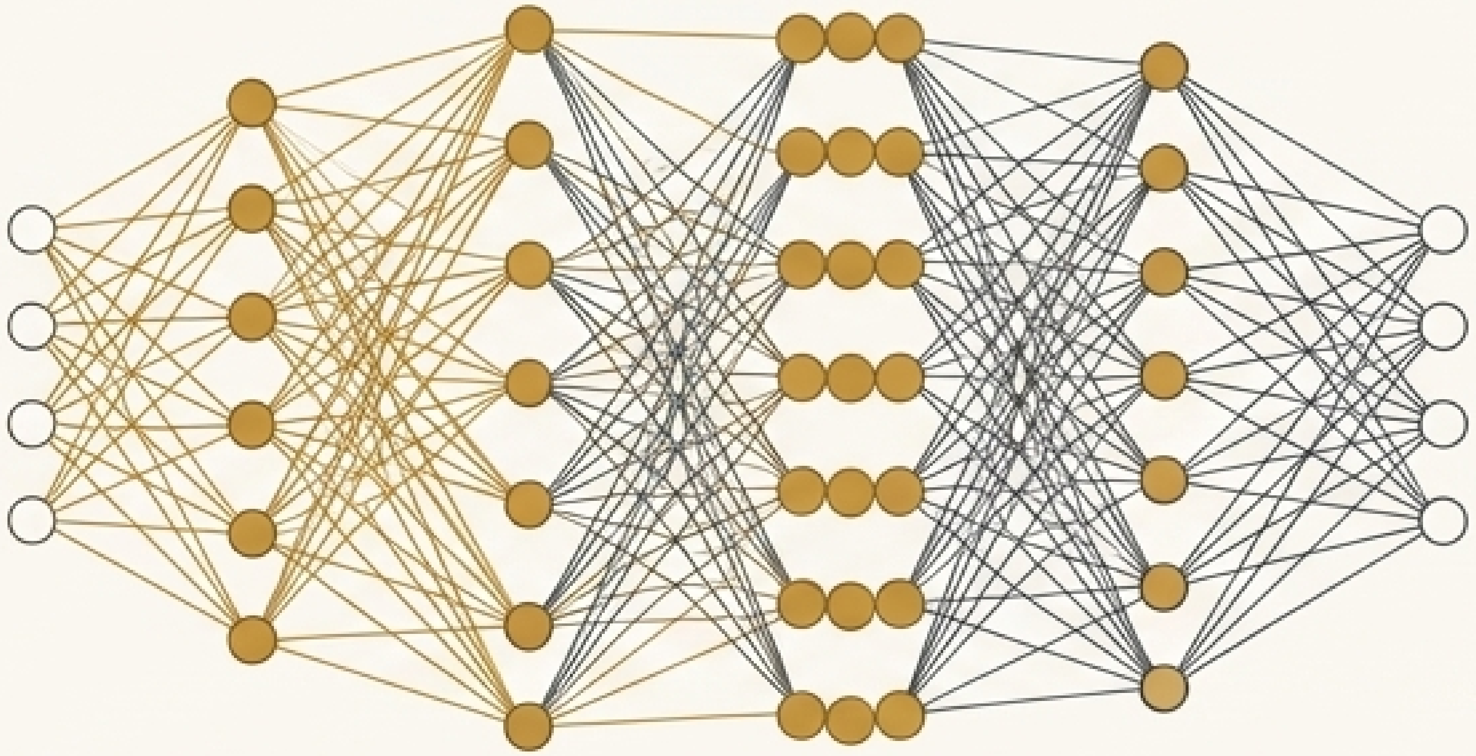

This theory explains *why* architectures like Convolutional Neural Networks (which have built-in translational symmetry) are so effective for images. They don't have to learn this symmetry from scratch.

A New Design Principle: Instead of forcing networks to learn symmetries from massive datasets, we can impose them as structural constraints from the start. This leads to superior data efficiency and better generalization.

Future AI could be designed to discover and exploit the specific Lie group symmetries of any given problem domain.

The convergent principles of algorithmic simplicity and mathematical symmetry provide a deep insight into intelligence itself.

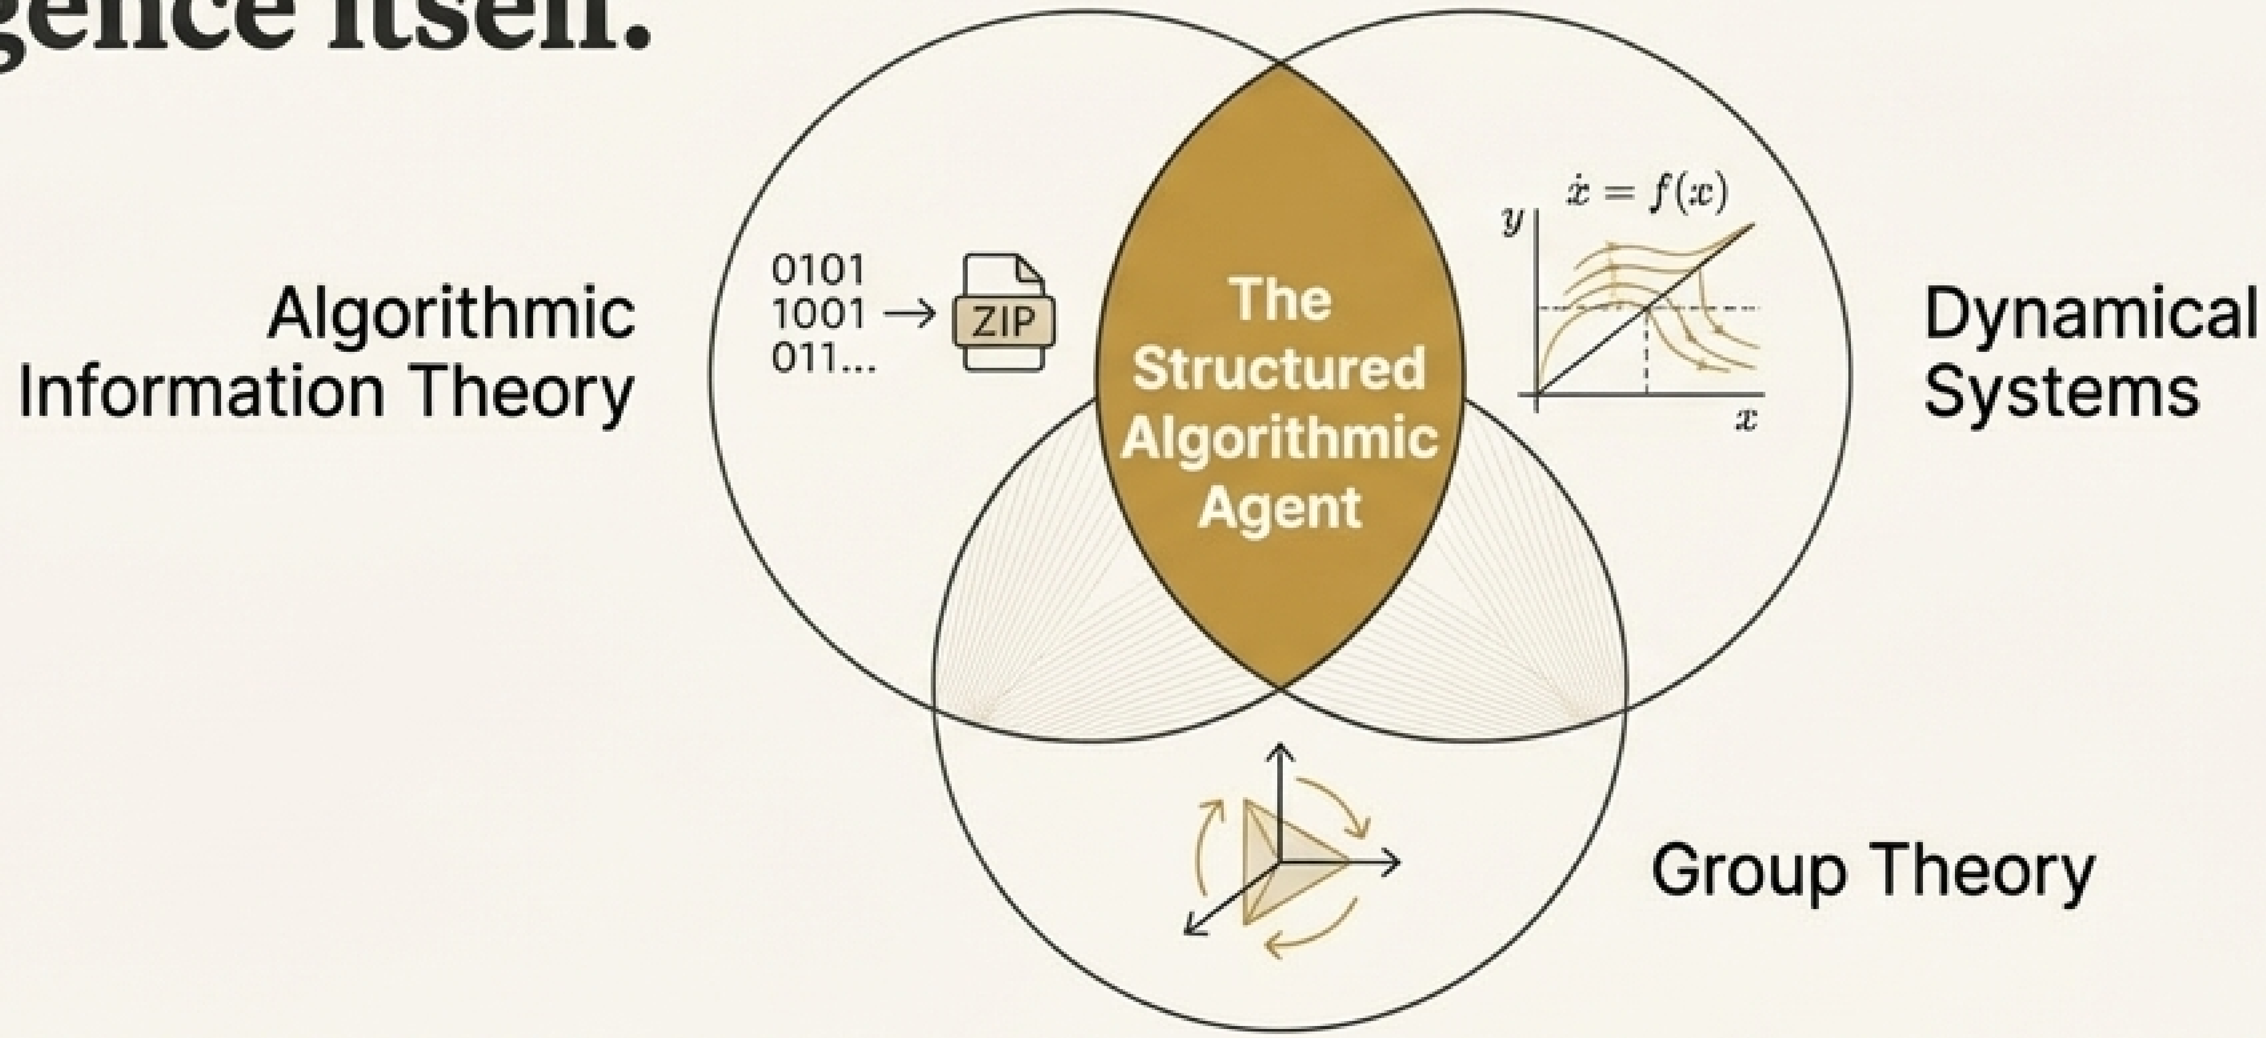

This framework connects the high-level principles of information theory (compression) with the low-level mechanics of dynamical systems (neural equations) via the language of group theory.

Symmetry is not merely an aesthetic feature of the world; it is a necessary consequence of any agent successfully modeling a compressible reality.

This suggests a path forward: the search for intelligence is, in large part, a search for the right symmetries.

A Legacy of Symmetry

The historical trajectory of group theory—the mathematical theory of symmetry—is a testament to its power. From Galois and Ruffini revealing the symmetries of equations, to Sophus Lie applying them to continuous dynamics, to Emmy Noether's profound theorem connecting symmetry to the conservation laws that govern our universe. This work continues that quest, applying the lens of symmetry to understand the sicture of intelligence itself.