
A MULTI-LAYER DYNAMICAL ARCHITECTURE FOR HUMAN COGNITION: FROM CONTINUOUS THOUGHT TO STRUCTURAL SELF-CONSISTENCY

A PREPRINT

© **Gou Zhang Cai**

Independent Researcher

(Cognitive Systems, AI)

Email: gouzhangcai0@gmail.com

April 14, 2026

ABSTRACT

This paper proposes a human cognitive architecture based on nonlinear dynamical systems, using a spatial dynamical representation as its core framework. The model introduces a three-layer structure—comprising a Cognitive Field, Particle Flow, and Pattern Primitives—to characterize the continuous evolution of thought processes.

Within this framework, a Gravitational Firmware Unit (GCU) is defined as a slow-varying latent structure that captures long-term stability and self-consistency in cognition. In addition, a Benefit-Seeking Principle is introduced as an intrinsic mechanism for directional decision-making, governing the evolution of thought trajectories through structural constraints rather than explicit optimization objectives. Furthermore, a Field-Landscape Model is constructed to establish a dynamical mapping between personality traits and behavioral patterns, enabling the representation of cognitive tendencies as evolving structural configurations.

In contrast to conventional models that emphasize convergence toward predefined targets, the proposed architecture treats cognition as a continuously evolving dynamical system. Through structural bias and multi-timescale coupling, it enables a closed-loop interaction between behavior and learning, allowing the system to maintain long-term behavioral consistency while preserving local adaptability.

Although the present work remains at a theoretical stage, it provides a unified dynamical perspective for modeling human-like cognition and offers a novel conceptual framework for the development of artificial general intelligence.

Keywords Cognitive Architecture · Dynamical Systems · Human Cognition · Multiscale Dynamics · Self-Consistency · Decision-Making Mechanisms · Emergent Behavior

1. Introduction

One of the central goals of cognitive science is to develop computational models that can explain and simulate human thought processes. Over decades of research, a variety of paradigms have been proposed, including symbolic approaches, connectionist models, and hybrid frameworks [1]. These approaches have demonstrated distinct advantages across different tasks and have been widely adopted in both cognitive science and artificial intelligence.

From a modeling perspective, these paradigms differ substantially in their representational assumptions. Symbolic approaches typically rely on explicit structures and rule-based reasoning, offering strengths in formal systems and logical consistency [2]. In contrast, connectionist approaches emphasize distributed representations and statistical learning, exhibiting strong performance in pattern recognition and generalization [3]. While each paradigm captures important aspects of cognition, they also exhibit limitations in representing continuous cognitive processes, long-term structural evolution, and multi-timescale behavioral dynamics. These differences reflect underlying modeling trade-offs rather than inherent superiority of one approach over another.

In recent years, dynamical systems theory has increasingly been introduced into cognitive modeling, providing a formal framework for describing cognition as a continuous-time evolving process [4]. This perspective treats cognitive states as trajectories in a state space, offering a natural way to capture temporal dependencies and context-sensitive behavior. In neuroscience, attractor-based dynamical systems have been used to explain the stability of working memory and the persistence of neural states [5], while neural population dynamics have been employed to characterize the continuous evolution of cortical activity.

This line of research has also begun to intersect with modern machine learning. For instance, continuous-time neural networks and neural ordinary differential equation models [6], along with their extensions to irregular time series [7], provide computational tools for modeling temporal dynamics. In addition, energy-based models [8] offer a unified perspective by linking state evolution and learning through optimization principles. Despite these advances, several open challenges remain. In particular, it is still unclear how to introduce stable structural constraints into continuous dynamical systems, and how to coordinate short-term dynamics with long-term structural evolution across multiple timescales. These challenges have motivated the development of multi-scale dynamical models and hierarchical temporal frameworks.

Overall, the dynamical systems perspective should not be viewed as a replacement for existing cognitive models, but rather as a complementary framework. It emphasizes cognition as a continuous process unfolding in a structured state space, providing a unified language for describing dynamics, structure, and long-term evolution, and serving as an important bridge between cognitive science and computational modeling.

In this work, we propose a dynamical cognitive architecture that explicitly integrates multi-layer structural representation, multi-timescale dynamics, and intrinsic decision mechanisms into a unified framework. Rather than focusing on direct engineering implementation, this study aims to provide a theoretical perspective on human-like cognition by modeling it as a structured dynamical system. Although the framework remains at an early theoretical stage, it offers a conceptual foundation for exploring alternative pathways toward artificial general intelligence, particularly in contrast to purely data-driven approaches.

2 Core Differences from Mainstream AI Models

The proposed architecture draws inspiration from several key characteristics of human cognition and differs from mainstream artificial intelligence models in the following aspects.

2.1 Objective: From Convergence Optimization to Continuous Evolution

Mainstream deep neural networks (DNNs) and large language models (LLMs) are typically trained by minimizing a loss function, aiming to converge to a parameter configuration that approximates the data distribution. Once training is completed, model parameters remain largely fixed, with subsequent updates usually performed through fine-tuning.

In contrast, the proposed architecture does not target a single convergence point. Instead, it models cognition as a continuously evolving dynamical process. The system continuously updates its internal structure and state distribution based on incoming inputs and historical states, without a clearly defined “training completion” phase, resembling an online adaptive system.

2.2 Parameter Representation: From Distributed Weights to Structured Variables

Mainstream models encode knowledge through high-dimensional parameters learned via statistical fitting, where individual parameters typically lack explicit semantic meaning.

In the proposed framework, the system state is represented by structured variables with explicit functional roles (e.g., dynamical intensity, connectivity distribution, structural position). These variables are designed to maintain interpretable relationships with system behavior, improving overall transparency and interpretability.

2.3 Attention Mechanism: From Explicit Module to Emergent State Distribution

In mainstream models, attention mechanisms are explicitly designed computational modules that selectively weight input information.

In contrast, the proposed architecture does not rely on a dedicated attention module. Instead, information selection emerges from the global state distribution of the system. The propagation paths and influence of different inputs are determined by the current structural configuration and ongoing dynamics, making attention an emergent property rather than an explicitly engineered component.

2.4 Inference Process: From Forward Computation to State Evolution

In conventional models, inference is performed as a forward computation under fixed parameters, producing outputs without modifying long-term internal states.

In the proposed framework, inference is treated as a continuous evolution of system states. Each input not only generates an output but also influences the current state of the system, with certain state changes persisting over time. This enables the accumulation and continuation of short-term cognitive states.

2.5 Learning Mechanism: From Scalar Optimization to Structural Constraints

Mainstream reinforcement learning and alignment methods typically rely on scalar rewards or loss functions to guide model behavior.

In contrast, the proposed architecture does not depend on explicit scalar rewards. Instead, behavior is guided by internal structural constraints and evolution dynamics. The system tends to evolve toward more stable or lower-cost configurations, where “cost” is implicitly defined by the current structural state rather than externally specified numerical objectives.

2.6 Self-Consistency: From External Maintenance to Structural Stability

In mainstream agent systems, consistency is often maintained through external mechanisms such as context windows, memory modules, or prompt engineering. These behaviors can be significantly altered within short timeframes.

In the proposed architecture, consistency arises primarily from slow-varying internal structures (e.g., core state distributions or long-term variables). These components change gradually over time, leading to stable behavioral biases and stylistic consistency.

2.7 Relationship Between Learning and Behavior: From Separation to Closed Loop

In mainstream models, learning (parameter updates) and inference (behavior generation) are typically treated as separate processes. The inference phase does not directly modify model parameters.

In contrast, the proposed architecture integrates learning and behavior into a unified closed-loop process. Behavioral outcomes continuously feed back into the system, influencing its internal state distribution without requiring explicit parameter updates. Short-term changes occur at the state level, while long-term effects gradually reshape core

3 Theoretical Origin and Design Rationale of the Architecture

The primary objective of the proposed architecture is to construct a unified model capable of representing the full process of human cognition and behavior. During the design process, it is necessary to balance between two extremes: on the one hand, avoiding excessive model complexity that leads to inefficient fitting of numerous microscopic variables; on the other hand, preventing oversimplification that fails to capture essential characteristics of human cognition, particularly creativity, dynamical behavior, and long-term consistency. Guided by this objective, the framework evolves progressively from initial structural hypotheses toward a dynamical modeling paradigm, ultimately forming a cognitive architecture centered on dynamical interaction processes.

3.1 Early Exploration: Limitations of the Tree-of-Thought Model

In the initial stage, the architecture design was motivated by modeling human creative cognition. A fundamental assumption is that human cognition possesses strong generative capacity: rather than being confined to a finite set of

reasoning paths, it can produce highly diverse combinations and extensions within a latent space.

Based on this assumption, early attempts adopted a Tree-of-Thought structure for modeling. In this approach, the thinking process is represented as a progressively expanding tree, where nodes correspond to intermediate cognitive states and branches represent possible reasoning directions. Combined with generative models, each node can dynamically expand new content, enabling multi-path exploration in complex tasks to a certain extent.

However, further analysis reveals several structural limitations of this approach: 1. Computational and memory complexity: The tree structure grows exponentially as branches expand. In practical applications, approximation strategies such as pruning, sampling, or beam search are required, which inevitably constrain the effective exploration space. 2. Limitations of the reasoning paradigm: Although node contents are dynamically generated, the overall process still operates as discrete path expansion and selection. It fundamentally belongs to a search-driven paradigm, where generative capacity is constrained by the current distribution and expansion strategy, making it difficult to directly represent continuous state evolution. 3. Insufficient representation of dynamical properties: The Tree-of-Thought framework organizes reasoning into discrete steps, making it suitable for staged decision-making. However, it is less capable of capturing continuous transitions, cross-level interactions, and nonlinear integration of multiple factors.

Therefore, while the Tree-of-Thought approach offers clear advantages in enhancing exploration for complex tasks, its reliance on discrete path search introduces inherent limitations in modeling continuously evolving cognitive processes. This observation motivates a transition toward modeling approaches centered on state evolution.

3.2 Core Idea: From Structural Expansion to Dynamical Modeling

Building upon the limitations identified above, this work shifts its focus toward modeling the intrinsic dynamical nature of cognition. A key observation is that human thinking does not primarily operate through explicit structural expansion, but is better characterized as a process of continuous state transitions and interactions within a bounded system.

Based on this insight, the cognitive system is formulated as a bounded yet extensible state space. Within this space, different cognitive elements are represented as dynamic states, whose evolution is jointly driven by local interactions and global constraints. Each instance of thinking can thus be interpreted as a continuous trajectory of system states, rather than a discrete selection over predefined structures.

Within this framework, dynamical interactions—including state coupling, composition, and reconfiguration—serve as the core mechanism for generating complex cognitive behaviors. Through appropriate constraints and evolution rules, the system is capable of producing highly diverse behavioral patterns within a limited representational space, while maintaining overall stability and consistency.

It is important to emphasize that this process is not unconstrained stochastic variation. Instead, system evolution is governed by a set of structural rules that regulate the direction, magnitude, and scope of state changes. These constraints ensure that, while preserving expressive capacity, the overall behavior of the system remains coherent and controllable.

4 Core Design of the Three-Layer Dynamical Cognitive Architecture


The proposed architecture models cognition as a multi-layer dynamical system evolving in a continuous space. The overall structure consists of three coupled subsystems: the outer Cognitive Field, the intermediate Particle Flow layer, and the inner Microscopic Modulation layer. These three layers operate at different scales: the Cognitive Field provides semantic structure and spatial constraints, the Particle Flow layer characterizes the primary trajectories of thought, and the Microscopic Modulation layer introduces fine-grained perturbations and individual variability. Together, they form a unified framework capable of describing continuous reasoning, associative jumps, and the emergence of stable behavioral patterns.

4.1 Outer Layer: Hollow-Sphere Cognitive Field

The Cognitive Field serves as the underlying space of the dynamical system. It is defined as a finite closed manifold embedded in a \square -dimensional Euclidean space. It can be intuitively represented as a hollow sphere, but may also be extended to higher-dimensional topological structures. Its primary role is to host cognitive and behavioral nodes,

while defining the boundary and semantic topology of the system.

Its key properties are as follows:



1. *Semantic definition of nodes*: Each point $x_i \in M$, on the manifold corresponds to an individual cognitive or behavioral node, representing a specific concept, action, perception, or abstract thought unit.

2. *Spatial encoding of relevance*: The semantic relationship between two nodes is directly represented by their spatial distance. Shorter distances indicate stronger semantic correlation, while larger distances indicate weaker association.

3. *Topological cognitive mapping*: Dense regions in the Cognitive Field correspond to highly related clusters of concepts, typically representing familiar or frequently accessed knowledge domains. Sparse regions represent abstract, peripheral, or underdeveloped concepts.

4. *Finite but extensible boundary constraint*: The Cognitive Field is a finite closed manifold, preventing unbounded cognitive expansion and uncontrolled divergence in an unstructured space. Meanwhile, as learning progresses, node density increases and the internal topology becomes more refined. The manifold itself can expand through higher-dimensional deformations projected onto its observable structure.


4.2 Intermediate Layer: Particle Flow as the Dynamical Core of Thought

The Particle Flow constitutes the dynamical core of the cognitive system and resides within the interior of the Cognitive Field. It describes the process of thinking rather than its content. In this framework, human cognition is modeled as a continuous motion of particle flow within the Cognitive Field, characterized by direction, inertia, and energy consumption.

The key idea is to unify logical reasoning and creative thinking through the dynamical properties of particle flow. Its three core attributes are:

1. *Direction and angular transitions*: The motion of particle flow is defined by a directional vector. Interactions between the flow and nodes in the Cognitive Field correspond to cognitive events: each collision represents a completed thought, and the impacted node defines its semantic outcome. The angle between consecutive trajectories determines the nature of thinking: smaller angles indicate coherent, logically consistent reasoning, while larger angles reflect discontinuity, abstraction, and creativity.
2. *Flow thickness (inertia)*: The cross-sectional magnitude of the particle flow represents cognitive inertia. Thicker flows exhibit strong directional persistence, requiring higher structural cost to change direction, corresponding to focused and stable thinking. Thinner flows are more flexible and easily redirected, but exhibit lower stability and continuity.
3. *Continuity constraint*: The particle flow follows a continuity constraint: its trajectory evolves smoothly unless perturbed by external inputs, internal mechanisms, or structural variations within the Cognitive Field. Directional changes occur in a constrained and consistent manner rather than arbitrarily.

Through these properties, the model provides a unified parameter space for both logical and creative cognition, treating them not as opposing modes, but as different dynamical states within the same system.



4.3 Inner Layer: Pattern Primitives as the Microscopic Modulation Layer

The Pattern Primitives consist of a large number of microscopic units distributed within the particle flow space. They represent the smallest composable elements of behavior and cognition, forming a transitional layer between conscious processes and underlying latent structures.

Importantly, Pattern Primitives do not directly correspond to explicit semantic content. Instead, their primary function is to modulate the trajectory of particle flow through local interactions, serving as the fundamental substrate for behavioral habits and cognitive styles.

Their key roles and properties include:

1. *Trajectory modulation:* Interactions between Pattern Primitives and particle flow induce small directional perturbations. While individual deviations are minimal, their accumulation over time can significantly alter the global trajectory of thought.
2. *Determination of cognitive characteristics:* The number and spatial distribution of Pattern Primitives determine the system's behavioral tendencies: A higher density leads to frequent interactions, resulting in cautious, refined, and self-correcting thought processes. A lower density leads to fewer interruptions, enabling faster but potentially more biased or less stable decision-making.
3. *Exploration and stochastic perturbation:* This layer introduces low-intensity stochasticity, allowing the system to deviate from existing trajectories and explore new regions of the state space. This mechanism is essential for generating novel associations and creative outcomes.
4. *Source of creativity:* The irregular, quasi-random motion of Pattern Primitives introduces micro-level perturbations into the particle flow. These perturbations may redirect trajectories toward previously inaccessible regions of the Cognitive Field, enabling breakthrough insights and creative thinking.

The presence of the Pattern Primitive layer ensures that cognitive evolution is not solely governed by macroscopic rules. Instead, it enables fine-grained variability across individuals, capturing subtle differences in thinking habits and behavioral styles, and serving as the fundamental basis for individual differentiation.

5 Slow-Variable Structure and the Dynamical Mechanism of Self-Consistency

The three-layer dynamical cognitive architecture provides a unified description of continuous reasoning and associative jumps. However, relying solely on local interactions among the Cognitive Field, the Particle Flow layer, and the Microscopic Modulation layer introduces a fundamental challenge: under persistent perturbations and exploration dynamics, the system may gradually lose a stable global direction, leading to phenomena such as trajectory dispersion and inconsistent behavioral patterns.

This issue corresponds to the risk of disordered diffusion in cognitive systems. When local dynamics lack long-term constraints, the system may exhibit strong exploratory capability but fail to maintain consistency across time scales.

Therefore, a system driven purely by fast-changing dynamics cannot form stable behavioral tendencies or long-term structural biases.

To address this limitation, the proposed framework introduces a slow-variable structure, which constrains the long-term evolution of the system without interfering with local dynamics, thereby achieving a balance between stability and flexibility.

5.1 Necessity of Introducing Slow Variables

Empirical observations of human cognition reveal a fundamental property: core behavioral tendencies, value preferences, and thinking styles do not change significantly in response to single inputs or short-term experiences, but instead exhibit strong temporal stability. In contrast, specific thoughts and behavioral responses evolve at a much faster rate.

This distinction implies that cognitive systems inherently operate across multiple time scales:

1. Fast variables: corresponding to immediate processes such as thinking, decision-making, and action
2. Slow variables: corresponding to long-term structural biases, such as behavioral tendencies and value orientations

The coupling between these two types of variables enables the system to remain adaptable in the short term while preserving consistency in the long term. Therefore, introducing slow-variable structures is essential for constructing stable cognitive systems.

5.2 Gravitational Firmware Unit (GCU) and Its Closed-Loop Evolution Mechanism

Motivated by the necessity of slow variables, the architecture introduces the Gravitational Firmware Unit (GCU) as the core long-term structural component of the system. The GCU does not directly participate in specific reasoning or decision-making processes. Instead, it acts as a slow-variable structure that continuously biases the global evolution of the system, ensuring long-term consistency without interfering with local dynamics.

5.2.1 Definition and Properties of the GCU

1. *No intrinsic driving force*: The GCU does not directly generate specific thought trajectories or behavioral outputs. Instead, it influences the distribution of system behaviors by imposing long-term structural biases. Its initial state is neutral, and its eventual configuration is shaped by the system’s historical experiences and accumulated trajectories.
2. *Slow dynamics*: The evolution of the GCU exhibits time-scale separation. Compared to the rapid dynamics of the Particle Flow and Microscopic Modulation layers, the GCU changes only gradually through long-term accumulation, forming stable structural characteristics.
3. *Indirect influence mechanism*: The GCU operates indirectly. Rather than explicitly altering trajectories, it reshapes the internal structural distribution of the system, thereby modulating the accessibility and structural cost of different evolutionary paths.

5.2.2 Mechanism of Action and Closed-Loop Evolution

Based on the above properties, the GCU and the three-layer structure together form a cross-time-scale closed-loop evolution mechanism, allowing the system to maintain global stability during continuous change. This mechanism can be summarized as follows:

1. *Structural bias modulates Pattern Primitive distribution*: The GCU does not exert a literal physical “attractive force.” Instead, it modulates the probabilistic weighting of the state space over long time scales, causing the stochastic motion of Pattern Primitives to exhibit a statistical bias toward regions associated with the GCU. As a result, regions near the GCU become more likely to be sampled and maintained, while distant regions gradually decrease in frequency over time. The distribution of Pattern Primitives thus transitions from a uniform random process to a structure-weighted stochastic process, forming a non-uniform distribution field shaped by the GCU.

2. *Distribution structure determines trajectory cost of particle flow*: The spatial distribution of Pattern Primitives defines the “collision environment” for particle flow, thereby determining the statistical feasibility of trajectories: High-density regions → frequent interactions, enabling continuous correction and extension of trajectories, resulting in low structural cost and stable paths Low-density regions → lack of corrective interactions, leading to unstable trajectories with higher structural cost Consequently, particle flow is not governed by a single driving direction, but by a set of feasible paths constrained by the underlying distribution structure.
3. *Behavioral distribution reshapes the GCU*: The long-term, high-frequency trajectories of particle flow induce statistical changes in the distribution of Pattern Primitives. Through structural feedback, these distributional changes gradually shift the effective position of the GCU. Specifically: Frequent trajectories reinforce local distributions; Reinforced distributions shift the structural weight center; The shifting weight center leads to gradual GCU drift; This establishes a closed-loop mechanism: behavioral trajectories → distribution structure → GCU bias → behavioral trajectories.

Through this process, the system forms a stable cross-scale feedback loop: short-term behavior is determined by current state dynamics, while long-term structure is shaped by behavioral distributions, which in turn constrain future behavior. This mechanism enables coordination across time scales, maintaining global consistency while preserving local flexibility.

5.3 Unified Dynamical Origin of Self-Consistency and Creativity

In this framework, self-consistency is not explicitly defined by fixed rules or parameters. Instead, it emerges from structural biases formed through long-term system evolution. This consistency is not static but arises from dynamically stable structures under multi-time-scale coupling.

Specifically, this structural consistency manifests in several ways:

1. The system tends to select similar particle flow directions under different inputs, resulting in behavior patterns that share consistent structural characteristics while retaining variability
2. Long-term structural changes occur gradually and are not significantly altered by single inputs or short-term perturbations, ensuring stable evolution
3. The coupling between fast and slow variables constrains only the long-term direction of evolution, without interfering with local dynamics

Within the same framework, creativity is not suppressed by structural constraints; instead, it is redefined and reinforced. Local dynamics remain open, allowing short-term perturbations and trajectory deviations (e.g., stochastic fluctuations of Pattern Primitives or large-angle transitions in particle flow). Thus, the system retains the ability to explore new regions of the state space.

Under this perspective, creativity is defined as: **the recombination of states and reconstruction of trajectories within an existing structural constraint, rather than unconstrained random generation.**

This formulation ensures that creativity maintains both openness and directionality, avoiding collapse into unstructured randomness.

Therefore, long-term structural bias does not act as a limiting mechanism but as a boundary-regulating structure. It provides a stable reference framework for exploration, enabling generated behaviors to retain both diversity and interpretable directionality, thereby unifying stability and creativity.

Summary: By introducing the GCU as a slow-variable core, the proposed architecture establishes a cross-time-scale regulatory mechanism within a continuous dynamical system. This mechanism allows the system to maintain short-term flexibility while forming long-term stable behavioral biases, thereby achieving self-consistency and sustainable cognitive evolution.

6 Benefit-Taxis Principle and Directional Decision-Making Mechanism

The three-layer dynamical structure, together with the GCU mechanism, establishes the state space and long-term consistency of the cognitive system. However, dynamical evolution alone does not fully address a fundamental

question: when multiple feasible trajectories coexist, what determines the system’s choice among them? In other words, the dynamical structure describes what can happen, but does not by itself explain why a particular trajectory is selected. To resolve this limitation, the proposed framework introduces the Benefit-Taxis Principle, which provides a directional selection rule for the system, transforming it from a neutral dynamical system into one with intrinsic preference structure.

6.1 Definition of the Benefit-Taxis Principle

The Benefit-Taxis Principle is not an external reward function nor an explicit utility maximization mechanism. Instead, it is a direction-selection rule determined by the system’s internal structure. Its core definition is as follows:

The cognitive system tends to evolve along directions that are most consistent with its current dynamical state and that incur the lowest structural evolution cost.

Under this definition:

1. “Benefit” is not an externally assigned scalar, but an implicit property defined by internal structural coherence
2. Decision-making is not the selection of discrete actions, but the determination of the next direction of state evolution
3. The principle operates at the level of the dynamical system, influencing trajectory evolution through differences in structural cost

6.1.1 Structural Origins of Benefit and Direction Selection

In this framework, benefit is not determined by a single variable, but emerges from the interaction of multiple structural components across different temporal and spatial scales. It can be interpreted as a tendency toward minimizing the overall structural cost under current system conditions.

Specifically, three key mechanisms jointly determine directional selection:

1. *Inertial direction of particle flow*: The current motion of particle flow represents the system’s short-term dynamical inertia. Continuing along this direction typically requires minimal structural adjustment, resulting in the lowest local cost. This forms the primary source of continuity in the system.
2. *Implicit bias from the GCU*: The long-term structural bias encoded by the GCU imposes cross-time-scale constraints on direction selection. As a slow-variable core, the GCU defines stable tendencies in long-term evolution. Trajectories aligned with this bias incur lower global restructuring costs, while deviations require higher reconfiguration effort, giving aligned paths a long-term advantage.
3. *Local perturbation cost from the Pattern Primitive layer*: The microscopic modulation layer determines the local accessibility and smoothness of trajectories. The spatial distribution of Pattern Primitives defines local structural density: Dense and continuous regions → lower perturbation cost and smoother trajectory evolution; Sparse regions → higher uncertainty and greater adjustment cost

Overall, benefit emerges as the combined effect of these structural constraints, and the system naturally prefers directions that minimize total structural cost.

From a dynamical perspective, directional selection can be interpreted as a trade-off between turning magnitude and structural cost: Small-angle transitions correspond to low-cost, locally continuous evolution and are therefore preferred; Large-angle transitions require substantial restructuring and are typically selected only when long-term benefit outweighs short-term cost

This unified mechanism provides a consistent explanation for both continuous and discontinuous thinking, as well as phenomena such as difficulty in resuming interrupted thought processes and differences in efficiency between familiar and unfamiliar tasks.

6.2 Temporal Structure of Benefit and Continuous Learning Mechanism

The Benefit-Taxis Principle exhibits a clear hierarchical structure across time scales, which can be described as two

coupled decision layers: local benefit and global benefit. These are not independent rules, but responses to structural cost at different temporal scales within the same dynamical framework.

Local benefit (short-term): Governs immediate decision-making by favoring low-cost, low-perturbation trajectories. This ensures stable responses in rapidly changing environments and dominates high-frequency decision processes, maintaining local continuity and adaptability.

Global benefit (long-term): Governs long-term structural optimization by prioritizing consistency and stability of the overall evolutionary direction. At this level, the system may accept short-term high-cost transitions in exchange for long-term structural improvement. This mechanism is primarily mediated by the GCU, which gradually adjusts long-term structural bias.

Together, these two levels form a unified decision framework, balancing short-term adaptability and long-term consistency, thereby preventing behavioral degradation under single-scale dynamics.

6.3 Continuous Decision-Making and Learning as Structural Evolution

Within this framework, decision-making is no longer an isolated instantaneous event, but a continuous process of trajectory selection within dynamical evolution. Each decision not only determines the current state but also reshapes the structural cost landscape for future evolution, introducing temporal dependency and path accumulation effects.

Specifically: Each decision alters the inertial direction of particle flow and the spatial distribution of Pattern Primitives; A single directional choice introduces short-term path dependence; Repeated similar decisions gradually shape long-term structural preferences

Over time, repeated behavioral patterns progressively reshape the system’s structural bias: frequently traversed paths decrease in structural cost, while alternative paths become relatively more costly. As a result, system behavior exhibits stable statistical patterns, including consistent style and directional tendencies.

In this process, **learning is no longer defined as error correction or loss minimization, but as: the process of modifying which directions are considered beneficial within the system.**

When a particular class of behavior or cognitive pattern is repeatedly executed, the internal structural distribution undergoes gradual adjustment. Through cross-time-scale feedback, this adjustment further influences future benefit evaluation, ultimately forming a stable closed-loop relationship:

behavior shapes structure, and structure guides behavior.

7 Dynamical Mapping Mechanism Between Personality and Behavior

Within the proposed framework, the GCU, as a slow-variable structure, determines the long-term evolutionary direction of the system, while the Particle Flow and Microscopic Modulation layers govern the generation of specific thoughts and behaviors. However, these components alone are insufficient to explain a critical question: how are abstract personality biases translated into concrete behavioral trajectories?

To address this gap, this chapter introduces three additional mechanisms—Force Field representation, Acceleration Anchor mechanism, and Benefit-Landscape structure—to establish a continuous dynamical mapping from personality to behavior, enabling abstract structural biases to be transformed into evolvable trajectories in the state space.

7.1 Force Field: Spatial Manifestation of Personality Structure

To provide an operational representation of personality bias, the framework introduces the Force Field as an intermediate structural representation. The Force Field is jointly determined by the spatial position of the GCU, the distribution of the Microscopic Modulation layer, and the dynamical properties of the Particle Flow. Its essence is a structural potential field defined over the cognitive space.

Importantly, this field is not externally imposed, but emerges naturally from the internal state of the system. Its key properties are as follows:

1. *Field shape determined by multi-layer structure:* The global shape of the Force Field arises from the interaction of multiple structural components: the GCU defines global bias, the Microscopic Modulation layer determines local variations, and the inertia and thickness of Particle Flow determine sensitivity to structural changes. Over

time, local behaviors reshape the field, while stable behavioral patterns form persistent depressions or elevations, encoding long-term preferences.

2. *Field complexity reflects personality richness:* The structural complexity of the Force Field corresponds to the expressive richness of personality. Systems with richer microstructures and higher dynamical degrees of freedom produce more complex field geometries, enabling finer-grained behavioral differentiation. Conversely, simpler structures lead to more stable and less diverse behavior patterns.
3. *Field uniqueness reflects individual differences:* Different experiences and trajectories shape distinct Force Field configurations. Even systems with identical initial conditions will diverge into different fields under different experiential inputs, resulting in individualized behavioral tendencies.

A simple spatial representation of the GCU alone is insufficient to capture the complexity of personality. In contrast, the Force Field—shaped by structural bias, slow variables, interaction dynamics, and external feedback—provides a compact yet expressive dynamical representation of personality.

7.2 Acceleration Anchors: Intrinsic Localization of Cognitive Paths

Building upon the Force Field representation, the framework introduces the Acceleration Anchor mechanism to address the problem of accessing and reactivating cognitive trajectories.

Unlike traditional approaches that rely on explicit coordinates or symbolic indexing, this mechanism does not depend on absolute spatial positions. Instead, it encodes trajectories through patterns of acceleration changes in particle flow. That is, the system records dynamic transformation patterns rather than static node locations.

The key insight is that trajectory recognition is based on similarity in dynamical patterns rather than exact matching. When the system encounters a similar acceleration pattern, the corresponding trajectory can be reactivated, enabling efficient recall.

Intuitively, this can be understood as remembering the “felt dynamics” of a thought process: when a similar dynamical pattern reoccurs, the system naturally reorients toward the corresponding trajectory.

This mechanism provides two important capabilities: Compressed experiential memory: Memory is represented implicitly through dynamical features rather than explicit storage. Trajectory recoverability: Even if a structural pathway weakens over time, it can be reactivated through matching dynamical anchors.

7.3 Benefit-Landscape: Spatial Representation of Decision-Making

By incorporating the Benefit-Taxis Principle into the Force Field structure, the system’s structural cost can be mapped into a scalar field over the cognitive space, forming a Benefit-Landscape.

Within this landscape, different directions correspond to different structural costs, and the system’s evolution can be interpreted as trajectory selection within this terrain. The structure exhibits the following key properties:

1. *Landscape height corresponds to structural cost:* Low-cost regions correspond to low “elevation,” while high-cost regions correspond to high “elevation.” Particle Flow evolves by moving toward lower-cost regions, effectively following a “downhill” trajectory.
2. *Path reinforcement effect:* Frequently traversed paths gradually decrease in structural cost, forming path-dependent structures. Conversely, rarely used paths increase in cost due to limited representational capacity, leading to natural weakening or forgetting. However, due to the presence of Acceleration Anchors, such paths can still be rapidly reactivated when needed.
3. *Intersection and associative transitions:* When multiple low-cost paths intersect, they form structural junctions. These junctions enable automatic transitions between different cognitive trajectories and serve as the structural basis for association and cross-domain thinking.
4. *Global reshaping via long-term benefit:* Long-term (global) benefit enables the system to temporarily follow high-cost paths (i.e., move “uphill”) in order to reshape the overall landscape. This process gradually forms trajectory structures that better align with long-term global benefit.

7.4 Unified Mapping from Personality to Behavior

In summary:

1. The GCU provides long-term structural bias
2. The Force Field provides spatial representation
3. The Acceleration Anchors provide trajectory memory
4. The Benefit-Landscape provides decision rules

Together, these components form a complete mapping chain:

Personality (GCU) \rightarrow Spatial Field \rightarrow Dynamic Trajectory \rightarrow Landscape-Based Decision \rightarrow Behavioral Output

Within this framework: Behavior is no longer an independent output, but the result of trajectory evolution under structural constraints. Personality is no longer an abstract descriptor, but a stable bias embedded within the dynamical system

Thus, cognition, decision-making, and behavior are unified as a continuous dynamical evolution process within a structured landscape.

8 Discussion, Limitations, and Conclusion

This study proposes a human-like cognitive modeling framework based on nonlinear dynamics, organized around a “cognitive field–particle flow–microscopic modulation structure.” It further introduces a slow-variable structure (GCU) and the Benefit-Taxis Principle to describe behavioral differences and decision mechanisms across multiple time scales. Overall, this framework attempts to provide a unified dynamical description of thinking, learning, and behavior. Its primary contribution lies in structural representation and mechanism abstraction, rather than in a directly deployable engineering system.

8.1 Key Contributions and Innovations

The main contributions of this study can be summarized as follows:

First, the proposed framework attempts to represent cognitive processes as a continuous dynamical system, rather than relying on traditional discrete symbolic processing or purely probabilistic modeling. Under this perspective, thinking, reasoning, and decision-making are described as manifestations of a unified dynamical system across different time scales, providing a structured and unified interpretation of cognition.

Second, the framework introduces a hierarchical coupling mechanism between slow variables (GCU) and fast dynamical processes, capturing the relationship between long-term stability and short-term flexibility in cognitive systems. This mechanism provides a structural explanation of how long-term preferences influence short-term behavior, without relying on external memory modules or explicit rules.

Third, at the decision-making level, the Benefit-Taxis Principle is proposed, reformulating decision-making from an explicit reward maximization problem into a structure-cost-driven path selection process. This design explains preference formation from the system’s internal structural state rather than externally defined value functions.

In addition, the force field and landscape representations serve as intermediate structures linking abstract biases to concrete behavioral trajectories. This provides a formal mechanism that maps “structural state” to “behavioral selection,” allowing personality-like biases to be expressed as dynamically evolving structures rather than static parameters.

8.2 Current Limitations and Implementation Challenges

It should be noted that this study remains at a conceptual and theoretical stage, with several limitations in formalization and practical implementation.

First, the current model remains largely descriptive rather than fully formalized mathematically. A complete

analyzable dynamical system framework has not yet been established. For example, stability conditions, convergence properties, and provable emergence mechanisms remain to be rigorously defined.

Second, numerical simulation of high-dimensional continuous dynamical systems is computationally challenging. The coupling between particle flow evolution, microscopic interactions, and slow-variable updates involves multiple time scales, placing significant demands on computational efficiency and stability. Efficient implementation methods are not yet available.

Third, an evaluation framework has not yet been established. Metrics for assessing “personality consistency,” “long-term preference stability,” or “creativity” in such systems remain undefined, limiting empirical validation and comparative analysis.

8.3 Conclusion

This paper proposes a nonlinear dynamical framework for cognitive modeling. It uses a three-layer structure—cognitive field, particle flow, and microscopic modulation units—to describe the fundamental dynamics of cognition, introduces a slow-variable structure (GCU) to capture long-term preferences and stability, and employs the Benefit-Taxis Principle to construct a structure-cost-based decision mechanism.

At the structural level, the framework provides a unified representation that integrates thinking processes, behavioral selection, and long-term preferences into a single dynamical system. Through force field and landscape representations, it establishes a mapping between abstract biases and concrete behaviors.

It should be emphasized that the current model is primarily a structured theoretical exploration. Its value lies in conceptual unification and mechanism abstraction, rather than in immediate engineering applicability.

Future work will focus on the mathematical formalization of GCU spatial weighting mechanisms, as well as exploring the integration between the cognitive field representation and large language models.

References

- [1] Newell A, Simon H A. Human problem solving[M]. Prentice-Hall, 1972.
- [2] Fodor J A, Pylyshyn Z W. Connectionism and cognitive architecture: A critical analysis[J]. *Cognition*, 1988, 28(1-2): 3–71.
- [3] LeCun Y, Bengio Y, Hinton G. Deep learning[J]. *Nature*, 2015, 521(7553): 436–444.
- [4] van Gelder T. The dynamical hypothesis in cognitive science[J]. *Behavioral and Brain Sciences*, 1995, 18(3): 615–628.
- [5] Amit D J. Modeling brain function: The world of attractor neural networks[M]. Cambridge University Press, 1992.
- [6] Chen R T Q, Rubanova Y, Bettencourt J, Duvenaud D. Neural ordinary differential equations[J]. *Advances in Neural Information Processing Systems*, 2018, 31.
- [7] Rubanova Y, Chen R T Q, Duvenaud D. Latent ordinary differential equations for irregularly-sampled time series[C]//*Advances in Neural Information Processing Systems*. 2019.
- [8] LeCun Y, Chopra S, Hadsell R, et al. A tutorial on energy-based learning[C]//*Predicting structured data*. MIT Press, 2006.