Research Article ISSN: 3143-4371

Innovative Insights in Digital Health

Variational Integration of Epistemic, Utilitarian, and Normative Interactions for Mental Wellbeing

Edouard Siregar[1,2] *

[1]*Seji Labs, Estepona, 29680, Spain.*

[2]*New York Academy of Sciences; AI focus group, New York City, NY, USA.*

Received Date: 20 Jul 2026;
Accepted Date: 27 Aug 2026;
Published Date: 12 Sep 2026.

***Correspondence:** *Edouard Siregar, Seji Labs, Estepona 29680, Spain, and New York Academy of Sciences, AI Focus Group, New York City, NY, USA.*

Citation: Edouard Siregar. Variational Integration of Epistemic, Utilitarian, and Normative Interactions for Mental Wellbeing. Innov Insights Digit Health. 2026; 2(3): 1-15.

ABSTRACT

Mental wellbeing emerges from the continuous integration of perception, decision-making, and social regulation. We propose a unified mathematical framework based on the Field Theory of Representation (FTR), in which these interacting processes are modeled as a variational field defined over a learned semantic manifold.

Within FTR, an agent's cognitive state evolves according to a Lagrangian density whose potential term V (E, U, N) captures the coupled influence of epistemic evidence E, utilitarian objectives U, and normative constraints N. These three components represent complementary forces that jointly shape adaptive cognition and behavior.

Inference is formulated as trajectories on a latent manifold $\mathcal{M} \subset \mathbb{R}^d$, where the geometry encodes the semantic organization of experience. Mental wellbeing corresponds to stationary solutions of the associated action functional, balancing representational efficiency, accurate environmental modeling, goal-directed behavior, and normative coherence.

FTR therefore provides a geometric framework for studying how cognitive systems maintain stable representations while continuously adapting to changing sensory, behavioral, and social demands. The theory generates quantitative hypotheses for large-scale neurocognitive dynamics, cognitive resilience, stress adaptation, and dysregulation associated with neurological and psychiatric disorders.

By expressing these phenomena within a common variational formalism, FTR offers a principled bridge between mathematical models of cognition, computational neuroscience, and future artificial intelligence systems whose adaptive behavior remains aligned with human wellbeing.

Keywords: *Field Theory of Representation (FTR); Mental Wellbeing; Variational Framework; Cognitive Systems; Computational Neuroscience.*

1. Introduction

Artificial intelligence increasingly combines learned representations, reasoning, planning, decision-making, and adaptation. In contemporary systems, however, these processes are usually implemented through different mechanisms. Representation learning constructs an internal state space; inference operates on that representation; optimization changes model parameters; planning constructs action sequences; and safety or alignment constraints are often introduced as additional objectives or external control mechanisms.

This separation is computationally useful, but it leaves open a fundamental theoretical question:

What mathematical object simultaneously describes the representation on which reasoning occurs and the evolution of that representation through learning?

We investigate this question through a variational formulation of adaptive cognition. The central proposal is that a cognitive system should be described by two coupled mathematical objects:

$$\gamma(t), \qquad \mathfrak{G}(t).$$

The first is a cognitive trajectory representing reasoning or decision evolution. The second is a Cognitive Geometry describing the latent representation space in which that trajectory evolves.

The distinction can be summarized as

$$\boxed{\begin{gathered}\text{reasoning} : \mathcal{M} \to \mathcal{M},\\ \text{learning} : \mathcal{G} \to \mathcal{G},\\ \text{adaptive cognition} : \mathcal{M} \times \mathcal{G} \to \mathcal{M} \times \mathcal{G}.\end{gathered}}$$

Here $\mathcal{G}$ denotes a space of admissible Cognitive Geometries. The essential consequence is that learning does not merely move the cognitive state within a fixed representation space. It can also modify the geometry of that space.

The framework is motivated by a broader Field Theory of Representations (FTR), in which latent representations are treated geometrically and reasoning is described by constrained variational dynamics [13, 14]. The present paper is independent of a sequence of earlier formulations. It develops a separate theoretical contribution centered specifically on the coupling between trajectory dynamics and geometry dynamics. Earlier FTR work provides conceptual and mathematical background, but the present paper focuses on the consequences of allowing the representation geometry itself to evolve.

A second component of the framework concerns the interaction between three classes of constraints:

$$\mathrm{E}, \quad \mathrm{U}, \quad \mathrm{N}.$$

The field E represents epistemic considerations such as evidence and information quality; U represents utilitarian or task-directed objectives; and N represents normative constraints. These quantities are not assumed to be interchangeable. Their interaction determines which cognitive trajectories are admissible and how competing trajectories are evaluated.

The resulting framework is therefore:

$$\boxed{\begin{gathered}\text{representation geometry} + \text{trajectory dynamics}\\ + \text{constraint fields} + \text{geometry evolution.}\end{gathered}}$$

The contribution of this paper is theoretical. We do not claim that current AI systems have already been shown empirically to instantiate the proposed geometry. Instead, we formulate mathematical objects, derive their coupled variational equations, establish analytical results under explicit assumptions, and identify empirical consequences that can be tested in future work.

The paper has four principal objectives.

First, we formalize Cognitive Geometry as a dynamical object rather than a fixed representation space.

Second, we formulate reasoning and learning as coupled variational processes.

Third, we show how several higher-order cognitive phenomena can be interpreted as different regimes of the same geometric dynamics.

Fourth, we examine how this formulation may contribute to trustworthy AI, including interpretability, safe adaptation, alignment, and the mathematical study of mental wellbeing.

1.1 Relation to Existing Work

The present framework lies at the intersection of variational approaches to cognition, information geometry, dynamical systems, computational neuroscience, and representation learning. Its central objective, however, differs from that of any one of these traditions: to formulate epistemic, utilitarian, and normative influences as interacting geometric fields within a common variational description of adaptive cognition.

A particularly important point of comparison is Active Inference and the Free Energy Principle [7, 8]. These frameworks describe perception, learning, and action through variational principles and emphasize the continuous coupling between an agent and its environment. In Active Inference, the primary variational object is generally expressed in terms of probabilistic beliefs, generative models, and variational free energy. The present work instead formulates cognition in terms of trajectories on a latent representation manifold together with the geometric structures that organize those trajectories. The two approaches are therefore potentially complementary: probabilistic belief dynamics may be represented within a latent geometric framework, while the present theory focuses specifically on the geometry of representation and the interaction of multiple cognitive constraint fields.

The framework is also related to information geometry, which equips statistical models with Riemannian metrics, affine connections, and curvature [1, 2]. Information geometry provides important mathematical tools for understanding the intrinsic structure of statistical model spaces. The present approach transfers this geometric perspective to latent representation spaces. Here, the metric is interpreted as a structure of cognitive or semantic similarity, while affine transport and curvature provide mathematical structures through which reasoning, contextual dependence, and geometric adaptation may be studied.

More broadly, the theory belongs to the tradition of dynamical and em bodied approaches to cognition, which emphasize that cognition is a continuously evolving process rather than a sequence of isolated symbolic operations [4, 16]. Related perspectives in cognitive science have also modeled behavior in terms of interacting control systems, predictive processes, and dynamical stability. The present framework develops this perspective in a variational-geometric direction by explicitly distinguishing between the dynamics of cognitive trajectories and the slower evolution of the geometric structures that constrain those trajectories.

The proposed treatment of mental wellbeing is also informed by computational and dynamical approaches to mental health. Recent work has increasingly emphasized that psychiatric and psychological phenomena may be studied through computational models of inference, learning, uncertainty, decision-making, and adaptive control [9, 10, 11]. Dynamical systems approaches likewise motivate the study of resilience, perturbation, recovery, and transitions between stable and unstable regimes [5, 12]. The present work does not propose a clinical model or diagnostic biomarker. Instead, it provides a mathematical framework within which questions concerning stability, adaptation, resilience, stress, and dysregulation can be formulated in terms of coupled geometric and variational dynamics.

Finally, the framework connects with modern representation learning and geometric approaches to machine learning. Contemporary neural systems learn high-dimensional latent representations whose structure is increasingly recognized as central to generalization, transfer, and adaptation. Geometric methods provide mathematical tools for analyzing such structures and their evolution [6]. However, representation learning typically emphasizes the optimization of parameters or embeddings for a particular objective. The present theory instead treats the representation geometry itself as an explicit dynamical object and introduces epistemic, utilitarian, and normative fields as distinct components of the variational landscape governing cognitive evolution.

The contribution of this work is therefore not to replace these existing approaches, but to provide a common mathematical language capable of relating them. The proposed Field Theory of Representations treats cognition as constrained variational evolution on an adaptive latent geometry, while mental wellbeing is considered as a potential application domain in which epistemic, goal-directed, and normative constraints must remain dynamically integrated.

The novelty of the present work lies in the explicit variational integration of epistemic, utilitarian, and normative fields on an evolving latent geometry, together with a two-level distinction between reasoning trajectories and geometric adaptation. This formulation permits mental wellbeing to be studied as a problem of dynamical stability, resilience, and constraint coherence without identifying wellbeing with any single scalar objective or claiming a direct clinical biomarker.

2 Cognitive Geometry

The basic representational object of the framework is a latent manifold

$$\mathcal{M}.$$

A point

$$x \in \mathcal{M}$$

represents a latent cognitive state. The manifold structure is intended to capture relationships among such states that are not adequately described by coordinates alone.

Definition 2.1 (Cognitive Geometry) *A Cognitive Geometry is a tuple*

$$\boxed{\mathfrak{G} = (\mathcal{M}, g, \nabla, E, U, N),}$$

where $\mathcal{M}$ is a latent manifold, g is a metric, ∇ is an affine connection, and E, U, and N are epistemic, utilitarian, and normative constraint fields, respectively.

The metric determines local similarity:

$$ds^2 = g_{ij}(x)\, dx^i\, dx^j.$$

Thus nearby points in the metric correspond to nearby latent states according to the representational geometry induced by the model.

The affine connection specifies how tangent information is transported along the manifold. For a trajectory

$$\gamma(t),$$

the covariant derivative

$$\nabla_{\dot\gamma}\dot\gamma$$

provides a coordinate-independent description of acceleration.

The curvature tensor associated with ∇ describes the failure of parallel transport to be path independent. In a cognitive interpretation, curvature may therefore encode structural dependence on the path by which a representation has been reached.

The fields

$$E, \quad U, \quad N$$

introduce constraint structure beyond purely geometric similarity. A generic constraint potential can be written

$$V(x) = V(E(x), U(x), N(x)).$$

More generally, the three fields may enter independently:

$$V = V_E(E) + V_U(U) + V_N(N) + V_{EUN}(E, U, N).$$

The final term permits explicit interaction among epistemic, utilitarian, and normative considerations.

This distinction is important. A system may encounter a state that is epistemically informative but normatively undesirable, or useful for a task but epistemically unsupported. The framework does not assume that these criteria coincide.

2.1 Representation Geometry and Cognitive Dynamics

Given a fixed geometry, a cognitive trajectory is a smooth map

$$\gamma : [0, T] \to \mathcal{M}.$$

The tangent vector

$$\dot{\gamma}(t)$$

describes the instantaneous direction of cognitive evolution.
A generic cognitive Lagrangian may be written

$$L_C = \frac{1}{2} g_{\gamma(t)}(\dot{\gamma}, \dot{\gamma}) - V(E, U, N; \gamma).$$

The corresponding action is

$$S[\gamma; \mathfrak{G}] = \int_0^T L_C(\gamma, \dot{\gamma}; \mathfrak{G})\, dt.$$

Stationarity,

$$\delta S = 0,$$

produces the cognitive Euler–Lagrange equations. In local coordinates,

$$\ddot{x}^k + \Gamma^k_{ij} \dot{x}^i \dot{x}^j = -g^{k\ell} \partial_\ell V,$$

for the simple kinetic-plus-potential form above.

This equation illustrates the basic separation between geometry and constraint. The Christoffel symbols describe geometric transport, while the gradient of V describes constraint-driven acceleration.

2.2 Why the Geometry Must Be Allowed to Evolve

A fixed representation geometry is sufficient for a system whose environment, task distribution, and internal organization remain stationary. Adaptive systems, however, encounter new observations and new tasks.

Consequently, the geometry may be represented as

$$\mathfrak{G}(t) = (\mathcal{M}_t, g_t, \nabla_t, E_t, U_t, N_t).$$

Learning then becomes evolution in the space

$$\mathcal{G} = \{\mathfrak{G} : \mathfrak{G} \text{ satisfies the admissibility conditions}\}.$$

The central conceptual distinction of the present paper is therefore

$$\boxed{\begin{array}{c}\text{reasoning} = \text{dynamics within } \mathcal{M},\\ \text{learning} = \text{dynamics within } \mathcal{G}.\end{array}}$$

This distinction provides the foundation for the coupled theory developed below.

Box Summary. Cognitive Geometry specifies the latent manifold, semantic metric, representation transport, and epistemic, utilitarian, and normative constraints. Reasoning evolves on this geometry. Learning changes the geometry itself.

3 Two-Level Variational Dynamics

The preceding section established Cognitive Geometry as the mathematical object that organizes latent representation. We now introduce two distinct levels of variational evolution.

At the first level, reasoning determines trajectories on the current geometry. At the second level, learning modifies the geometry itself.

3.1 Trajectory Dynamics

For fixed

$$\mathfrak{G} = (\mathcal{M}, g, \nabla, E, U, N),$$

reasoning is represented by

$$\gamma : [0, T] \to \mathcal{M}.$$

and the Cognitive Action

$$S[\gamma; \mathfrak{G}] = \int_0^T L_C(\gamma, \dot{\gamma}; \mathfrak{G})\, dt.$$

Stationary trajectories satisfy

$$\delta_\gamma S = 0,$$

The resulting Euler–Lagrange equations describe cognitive evolution under the geometric and constraint structure of $\mathfrak{G}$. During this variational problem,

$$\delta\mathfrak{G} = 0.$$

The geometry is therefore treated as an environmental structure for reasoning.

3.2 Geometry Dynamics

Learning is represented by a trajectory

$$t \mapsto \mathfrak{G}(t)$$

in the space $\mathcal{G}$ of admissible Cognitive Geometries. Define the geometric learning action

$$\mathcal{L}[\mathfrak{G}] = \int_0^T L_G(\mathfrak{G}, \dot{\mathfrak{G}})\, dt.$$

Here $\dot{\mathfrak{G}}$ denotes the appropriate tangent to the space of geometric structures. Depending on the model, this may include

$$\dot{g},\ \dot{\nabla},\ \dot{E},\ \dot{U},\ \dot{N}.$$

The learning stationarity condition is

$$\delta_{\mathfrak{G}} \mathcal{L} = 0.$$

Learning is therefore not identified with optimization over a vector of parameters. Parameterized models may provide computational realizations of $\mathfrak{G}$, but the mathematical object being evolved is the geometric structure.

3.3 The Coupled Action

Reasoning and learning cannot generally be separated because the reasoning action depends on $\mathfrak{G}$. We therefore define

$$\boxed{A[\gamma, \mathfrak{G}] = S[\gamma; \mathfrak{G}] + \lambda \mathcal{L}[\mathfrak{G}], \quad \lambda > 0.}$$

The first term governs trajectory dynamics and the second governs geometric adaptation. The forward relation is

$$\mathfrak{G} \to \gamma,$$

because the geometry determines the reasoning dynamics. The backward relation is

$$\gamma \to \mathfrak{G},$$

because cognitive experience can contribute to geometric adaptation. Thus

$$\boxed{\mathfrak{G} \to \gamma \to \mathfrak{G}.}$$

3.4 Coupled Euler–Lagrange Equations

The first variation is

$$\delta A = \delta_\gamma S + \delta_{\mathfrak{G}} S + \lambda \delta_{\mathfrak{G}} \mathcal{L}.$$

Variation with respect to γ gives

$$\boxed{\delta_\gamma S = 0.}$$

Variation with respect to $\mathfrak{G}$ gives

$$\boxed{\delta_{\mathfrak{G}} S + \lambda \delta_{\mathfrak{G}} \mathcal{L} = 0.}$$

The second equation contains the feedback from reasoning into learning. Thus, even if $\mathcal{L}$ describes an intrinsic geometric learning cost, the actual geometric evolution is influenced by the dependence of reasoning on the geometry.

Box Summary. The theory has two levels:

$$\gamma(t) \subset \mathcal{M} \ \text{ and } \ \mathfrak{G}(t) \in \mathcal{G}.$$

Reasoning evolves trajectories; learning evolves the geometry. Their interaction is governed by a coupled action.

4 Coupling Between Reasoning and Learning

The two variational levels are not independent. A reasoning trajectory is conditioned by the geometry at which it is evaluated, while experience generated by reasoning contributes to future geometric adaptation.

4.1 Reasoning on Fixed Geometry

For a fixed geometry,

$$\mathfrak{G} = \text{constant},$$

the reasoning problem is

$$\delta S[\gamma; \mathfrak{G}] = 0.$$

The metric determines semantic distance, the connection determines transport, curvature determines path dependence, and the fields

$$E, \ U, \ N$$

determine the constraint landscape. This corresponds to conventional variational reasoning on a fixed representation space.

4.2 Learning as Geometry Evolution

Learning instead solves

$$\delta \, \mathcal{L} \, [\mathfrak{G}] = 0$$

in $\mathcal{G}$. The object of adaptation may include

$$g_{t,} \ \nabla_t, \ \ E_t, \ U_t, \ N_t$$

Thus learning can alter semantic organization, transport, memory structure, and rational constraints.

4.3 Feedback

Suppose the geometry at time t is

$$\mathfrak{G}_t.$$

Reasoning produces a trajectory

$$\gamma_t^*$$

satisfying

$$\delta S[\gamma; \mathfrak{G}_t] = 0.$$

The resulting experience contributes to geometric adaptation:

$$\mathfrak{G}_t \to \mathfrak{G}_{t+1}.$$

Future reasoning then solves

$$\delta S[\gamma; \mathfrak{G}_{t+1}] = 0.$$

Hence the system follows

$$\boxed{\mathfrak{G}_t \to \gamma_t \to \mathfrak{G}_{t+1} \to \gamma_{t+1} \to \cdots}$$

There is therefore no permanently fixed optimization landscape.

4.4 Coupled Cognitive Evolution

Definition 4.1 (Coupled Cognitive Evolution). *A coupled cognitive evolution is a pair*

$$(\gamma(t), \mathfrak{G}(t))$$

satisfying the trajectory and geometry variational equations simultaneously.

The complete state belongs to

$$\boxed{(\gamma(t), \mathfrak{G}(t)) \in \mathcal{M} \times \mathcal{G}.}$$

The product-space formulation is important because it distinguishes the present framework from an algorithm that merely alternates two optimizers. The mathematical object is a joint variational system.

Box Summary. Learning changes the optimization landscape used by future reasoning, while reasoning generates the experience that changes the landscape. Adaptive cognition is therefore a closed feedback system rather than a sequence of independent optimization problems.

5 Hierarchical Cognitive Functionals

The coupled framework can accommodate several cognitive processes through functionals defined at different levels.

5.1 Cognitive Action

Definition 5.1 (Cognitive Action). *For an admissible trajectory γ,*

$$\boxed{S[\gamma] = \int_0^T L_C(x, \dot{x}; \mathfrak{G}) \, dt.}$$

Stationarity,

$$\delta S = 0,$$

determines reasoning dynamics.

5.2 Attentional Effort

Attention is modeled as computational effort associated with realizing a
trajectory:

$$\boxed{A[\gamma] = \int_0^T L_A(x, \dot{x}; \mathfrak{G})\, dt.}$$

5.3 Memory

Memory is represented through history dependence:

$$\boxed{\mathcal{M}[\gamma] = \int_0^T L_M(x, \dot{x}; \mathfrak{G}, \mathcal{H}_t)\, dt,}$$

where

$$\mathcal{H}_t = \{\gamma(s) : 0 \leq s \leq t\}.$$

Curvature, parallel transport, and accumulated geometric deformation can provide mechanisms for representing such history dependence.

5.4 Decision Value

Decision-making evaluates trajectories using epistemic, utilitarian, and normative criteria:

$$\boxed{\mathcal{D}[\gamma] = \alpha_E J_E[\gamma] + \alpha_U J_U[\gamma] - \alpha_N J_N[\gamma].}$$

The coefficients

$$\alpha_E, \alpha_U, \alpha_N \geq 0$$

specify the relative importance of the three criteria.

A simple combined trajectory problem is

$$\delta(S + \mathcal{D}) = 0.$$

5.5 Learning

Learning acts on geometry:

$$\boxed{\mathcal{L}[\mathfrak{G}] = \int_0^T L_G(\mathfrak{G}, \dot{\mathfrak{G}})\, dt.}$$

The stationary condition is

$$\delta\mathcal{L} = 0.$$

5.6 Functional Hierarchy

The resulting hierarchy is

$$\boxed{\begin{aligned} &S \rightarrow \text{reasoning dynamics,} \\ &A \rightarrow \text{attentional effort,} \\ &\mathcal{M} \rightarrow \text{historical dependence,} \\ &\mathcal{D} \rightarrow \text{trajectory evaluation,} \\ &\mathcal{L} \rightarrow \text{geometry adaptation.} \end{aligned}}$$

These functionals need not be interpreted as separate software modules. They are mathematical descriptions of distinct aspects of a common cognitive process.

6 The Hierarchy of Cognitive Dynamics

The functionals define a hierarchy according to the mathematical object being evolved.

6.1 Geometry

At the highest structural level,

$$\boxed{\mathfrak{G} = (\mathcal{M}, g, \nabla, E, U, N).}$$

It defines the representational environment in which all lower-level dynamics occur.

6.2 Learning

Learning modifies $\mathfrak{G}$:

$$\delta\,\mathcal{L}\,[\mathfrak{G}] = 0.$$

It therefore changes the environment available to future reasoning.

6.3 Reasoning

Reasoning evolves

$$\gamma(t) \subset \mathcal{M}$$

according to

$$\delta S[\gamma; \mathfrak{G}] = 0.$$

6.4 Attention

Attention evaluates the computational effort required to realize a trajectory:

$$A[\gamma].$$

6.5 Memory

Memory introduces history dependence through

$$\mathcal{M}[\gamma].$$

6.6 Decision

Decision evaluates competing trajectories through

$$\mathcal{D}[\gamma].$$

6.7 Unified Hierarchy

The complete hierarchy can be summarized as

$$\boxed{\begin{aligned} &\textbf{Geometry} \rightarrow \text{representation space,} \\ &\textbf{Learning} \rightarrow \text{geometry adaptation,} \\ &\textbf{Reasoning} \rightarrow \text{trajectory evolution,} \\ &\textbf{Attention} \rightarrow \text{resource allocation,} \\ &\textbf{Memory} \rightarrow \text{history dependence,} \\ &\textbf{Decision} \rightarrow \text{trajectory selection.} \end{aligned}}$$

The hierarchy is predominantly top-down but dynamically closed. Decisions generate experience, experience affects learning, learning modifies geometry, and the modified geometry changes future reasoning.

Box Summary. The framework does not require independent computational modules for every cognitive function. Instead, cognition is organized according to the mathematical objects that evolve and the functionals used to evaluate them.

7 Mathematical Theory of Coupled Cognitive Systems

We now formulate the mathematical object on which the coupled theory is based.

7.1 Definitions

Definition 7.1 (Coupled Cognitive State). *A coupled cognitive state is*

$$\boxed{C(t) = (\gamma(t), \mathfrak{G}(t)).}$$

Definition 7.2 (Total Cognitive Functional). *The total functional is*

$$\boxed{\mathcal{J} = S + A + \boldsymbol{\mathcal{M}} + \mathcal{D} + \lambda \boldsymbol{\mathcal{L}}}$$

The individual terms may depend on different components of the coupled state.

Definition 7.3 (Stationary Coupled Solution). *A coupled state*

$$(\gamma^*, \mathfrak{G}^*)$$

is stationary if

$$\boxed{\delta\mathcal{J}\,[\gamma^*, \mathfrak{G}^*] = 0.}$$

7.2 Analytical Assumptions

The following assumptions are sufficient for a basic existence result.

Assumption 7.4 (Admissible Space). *The admissible coupled state space*

$$\mathfrak{X} = \mathfrak{X}_\gamma \times \mathfrak{X}_\mathfrak{G}$$

is a nonempty reflexive Banach space or an appropriate weakly compact subspace.

Assumption 7.5 (Coercivity). The functional $\mathcal{J}$ is coercive on $\mathfrak{X}$.

Assumption 7.6 (Weak Lower Semicontinuity). *The functional $\mathcal{J}$ is weakly lower semicontinuous*.

Assumption 7.7 (Differentiability). *The first variation*

$$D\mathcal{J}$$

exists on the admissible space and is continuous.

7.3 Existence

Theorem 7.8 (Existence of a Coupled Minimizer). *Under the preceding assumptions, the coupled functional*

$$\mathcal{J} : \mathfrak{X} \to \mathbb{R} \cup \{+\infty\}$$

admits at least one minimizer

$$(\gamma^*, \mathfrak{G}^*) \in \mathfrak{X}\,.$$

If the minimizer lies in the interior of the admissible set and $\mathcal{J}$ is Frechet differentiable there, then

$$D\mathcal{J}\,(\gamma^*, \mathfrak{G}^*) = 0.$$

Proof. By coercivity, a minimizing sequence cannot escape to infinity. Reflexivity provides a weakly convergent subsequence. Weak lower semicontinuity then implies that the weak limit attains the infimum. If the minimizer is an interior point and $\mathcal{J}$ is differentiable, the first-order necessary condition gives

$$D\mathcal{J}\,(\gamma^*, \mathfrak{G}^*) = 0.$$

7.4 Uniqueness

Existence does not by itself imply uniqueness.

Theorem 7.9 (Uniqueness Under Strict Convexity). *Suppose, in addition, that $\mathcal{J}$ is strictly convex on the admissible convex set $\mathfrak{X}$. Then the coupled minimizer is unique.*

Proof. If two distinct minimizers existed, strict convexity would imply that their midpoint has strictly smaller functional value, contradicting minimality.

This result is deliberately stronger than a local positive-definiteness claim. Local positivity of the Hessian or second variation can establish local minimality, but global uniqueness requires an appropriate global convexity condition.

7.5 Local Stability

Suppose

$$(\gamma^*, \mathfrak{G}^*)$$

is a stationary solution and that the second variation satisfies

$$\delta^2\mathcal{J}\,[(\delta\gamma, \delta\mathfrak{G}), (\delta\gamma, \delta\mathfrak{G})] \geq c\,(\,||\delta\gamma||^2 + ||\delta\mathfrak{G}||^2)$$

for some $c > 0$ in a neighborhood of the stationary solution. Then the stationary pair is a strict local minimizer of $\mathcal{J}$.

Under the additional assumptions required to identify the variational problem with an appropriate dynamical system, this local coercivity provides the standard starting point for a Lyapunov stability analysis.

7.6 Representation Result

Corollary 7.10 (Coupled Representation). *Every admissible minimizer of $\mathcal{J}$ admits a representation as a coupled pair*

$$(\gamma^*, \mathfrak{G}^*)$$

satisfying the corresponding stationary equations.

The result generalizes fixed-geometry variational reasoning by allowing the representation geometry to participate in the optimization.

7.7 Consistency

If the admissible geometry remains within the constraint-compatible class during evolution, then stationary trajectories remain evaluated relative to the same admissibility structure.

Thus consistency is represented mathematically by preservation of the admissible set rather than by an independent post hoc constraint.

Box Summary. The coupled theory can provide existence under coercivity and lower semicontinuity, uniqueness under strict convexity, and local stability under appropriate second-variation conditions. These conclusions require explicit analytical assumptions and should not be interpreted as unconditional properties of arbitrary AI systems.

8 Emergent Cognitive Phenomena

The coupled formulation provides mathematical interpretations of several phenomena usually treated as separate learning mechanisms.

8.1 Adaptation

Definition 8.1 (Adaptive Cognitive System). *A system is adaptive if*

$$\dot{\mathfrak{G}}(t) \neq 0$$

and the geometric evolution depends on realized cognitive experience.

The feedback structure

$$\gamma \to \mathfrak{G} \to \gamma$$

makes adaptation intrinsic to the theory.

Proposition 8.2. *Whenever the trajectory-to-geometry coupling is nonzero, the realized reasoning trajectory can influence subsequent geometric evolution.*

This is a structural statement about the model, not an empirical claim that all real cognitive systems satisfy the assumed coupling.

8.2 Transfer Learning

Let

$$\mathfrak{G}_1,\ \mathfrak{G}_2$$

be two admissible Cognitive Geometries. A geometric transfer map is a smooth map

$$\Phi : \mathfrak{G}_1 \to \mathfrak{G}_2$$

that preserves the structures relevant to the transferred reasoning problem. In the ideal case,

$$S_2[\Phi \circ \gamma] = S_1[\gamma]$$

up to an admissible transformation. Transfer learning can then be interpreted as transferring geometric structure rather than merely numerical parameter values.

8.3 Meta-Learning

Learning evolves

$$\mathfrak{G}.$$

Meta-learning can be interpreted as evolving the law that determines the evolution of $\mathfrak{G}$. Schematically,

$$\boxed{\begin{aligned} &\text{Reasoning} \to \gamma, \\ &\text{Learning} \to \mathfrak{G}, \\ &\text{Meta-learning} \to \text{learning law.} \end{aligned}}$$

This creates a hierarchy of variational spaces.

8.4 Continual Learning

Continual learning corresponds to a continuous path

$$\mathfrak{G}(t) \in \mathcal{G}$$

that remains compatible with previously established representational structures.

In this interpretation, catastrophic forgetting can be studied as a failure of geometric continuity, excessive deformation, or instability of the constraint structure.

This does not imply that all forgetting in practical neural networks is geometric. Rather, it provides a testable mathematical hypothesis.

8.5 Collective Cognition

For N interacting agents, define

$$\mathfrak{G}_1, \ldots, \mathfrak{G}_N.$$

The joint state belongs to

$$\boxed{\mathcal{G}_{\text{collective}} = \prod_{i=1}^{N} \mathcal{G}_i}$$

A collective action may contain interaction terms

$$\mathcal{J}_{\text{int}} = \sum_{i<j} K_{ij}(\mathfrak{G}_i, \mathfrak{G}_j).$$

Collective reasoning then becomes coupled evolution on the product space. Communication, coordination, shared representations, and distributed learning can consequently be represented by geometric interaction terms.

Box Summary. Adaptation, transfer, meta-learning, continual learning, and collective cognition can be represented as increasingly rich forms of geometry evolution. The framework turns these phenomena into mathematical hypotheses rather than assuming separate cognitive modules for each.

9 Implications for Trustworthy Artificial Intelligence

The explicit representation of cognitive geometry provides a potential mathematical basis for several properties associated with trustworthy AI.

9.1 Interpretable Cognitive Geometry
The tuple

$$\mathfrak{G} = (\mathcal{M}, g, \nabla, E, U, N).$$

provides explicit objects that can, in principle, be inspected.

$$\mathcal{M} \to \text{latent organization},$$
$$g \to \text{semantic similarity},$$
$$\nabla \to \text{transport},$$
$$E \to \text{epistemic constraints},$$
$$U \to \text{task objectives},$$
$$N \to \text{normative constraints}.$$

Interpretation therefore becomes a structural problem rather than solely a post hoc feature-attribution problem.

9.2 Safe Geometry Adaptation
Learning should not be unrestricted deformation of $\mathfrak{G}$. Instead, define an admissible set

$$\mathcal{G}_{\text{safe}} \subseteq \mathcal{G}.$$

Learning becomes

$$\min_{\mathfrak{G}(t) \in \mathcal{G}_{\text{safe}}} \mathcal{L}[\mathfrak{G}].$$

Safety is therefore represented as a property of the admissible geometry evolution.

9.3 Alignment as Constraint Evolution
Alignment is represented through the fields

$$E,\ U,\ N.$$

The normative component N can encode restrictions that should remain valid throughout adaptation.

More generally, alignment can be formulated as maintaining a trajectory within an admissible constraint region:

$$\gamma(t) \in C_{\text{aligned}}(\mathfrak{G}(t)).$$

This formulation makes alignment dynamic. If $\mathfrak{G}$ evolves, the constraint structure evolves as well.

9.4 Robustness
Suppose

$$\gamma^*, \mathfrak{G}^*$$

is a stationary solution. A desirable robustness property is continuous dependence:

$$\|\delta\mathfrak{G}_0\| \text{ small} \longrightarrow \|\delta\gamma(t)\| \text{ controlled}.$$

Such a statement requires additional assumptions on the coupled dynamical operator. The framework therefore provides a mathematical language in which robustness can be defined and tested rather than assuming that geometric structure automatically guarantees it.

9.5 Trustworthy Cognitive Hierarchy
The resulting architecture can be summarized as

Geometry → representation,
Learning → geometry evolution,
Reasoning → trajectory optimization,
Decision → constraint evaluation,
Action → environmental interaction.

Trustworthiness is therefore not introduced as an isolated final module. Instead, it can be studied as a collection of mathematical properties of the coupled system.

Box Summary. Interpretability comes from explicit geometry, safety from constrained geometry evolution, alignment from explicit constraint fields, and robustness from stability and sensitivity properties of the coupled dynamics.

10 Mental Wellbeing as an Application Domain
The original motivation for the present framework arose from the problem of understanding how epistemic, utilitarian, and normative influences interact in cognition relevant to mental wellbeing. The broader theory developed in this paper does not identify wellbeing with cognition itself. Rather, mental wellbeing provides an important application domain in which the coupled variational structure of cognition can be given a concrete interpretation.

The central hypothesis is that wellbeing-relevant cognitive states arise when epistemic modeling, goal-directed behavior, and normative regulation remain sufficiently coherent while the cognitive system continues to adapt to changing environmental conditions. In this sense, wellbeing is treated as a dynamical property of an adaptive cognitive system rather than as a single scalar quantity or a fixed point in representation space.

This interpretation is deliberately mathematical and does not constitute a clinical model or diagnostic framework. Its purpose is to formulate hypotheses that can subsequently be tested using behavioral, neural, and computational data.

10.1 Epistemic, Utilitarian, and Normative Interaction
Let

$$V(E, U, N)$$

denote a potential describing the interaction among epistemic evidence, goaldirected utility, and normative constraints.

The three components represent distinct but interacting influences on cognitive evolution. The epistemic field describes the degree to which a trajectory is supported by available information or evidence. The utilitarian field represents task objectives, preferences, or goal-directed consequences. The normative field represents constraints associated with logical, social, ethical, or operational admissibility.

The important feature of the formulation is that these components are not required to be independently optimized. Their interaction is represented by the joint potential

$$V(E, U, N),$$

which may contain both additive and non-additive terms. For example, one may write

$$V(E, U, N) = V_E(E) + V_U(U) + V_N(N) + V_{EU}(E, U) + V_{EN}(E, N) + V_{UN}(U, N) + V_{EUN}(E, U, N).$$

The interaction terms allow the model to represent situations in which the effect of one cognitive constraint depends upon the state of another. For a cognitive trajectory

$$\gamma : [0, T] \rightarrow \mathcal{M},$$

a minimal cognitive Lagrangian is

$$L_C = \tfrac{1}{2} g(\dot{\gamma}, \dot{\gamma}) - V(E, U, N).$$

The corresponding action is

$$S[\gamma] = \int_0^T \left[\tfrac{1}{2} g(\dot{\gamma}, \dot{\gamma}) - V(E, U, N) \right] dt.$$

Stationary cognitive trajectories satisfy

$$\delta S[\gamma] = 0.$$

The resulting dynamics describe cognition as a compromise among geometric transport, epistemic evidence, goal-directed utility, and normative constraints. Importantly, the framework does not require a universal ordering among these influences. Their relative importance may depend on the cognitive state, environmental context, and evolving Cognitive Geometry.

This provides a mathematical interpretation of a central feature of wellbeing: cognitive behavior is generally not governed by a single objective, but by the interaction of multiple partially compatible constraints.

10.2 Wellbeing as a Dynamical Property

Within the present framework, mental wellbeing should not be identified with a particular value of the potential

$$V(E, U, N).$$

Nor should it be identified with stationarity alone. A stationary cognitive state may be stable but poorly adapted to its environment, while a continuously changing state may remain healthy if its evolution is coherent and recoverable.

We therefore introduce wellbeing as a property of the dynamics. Let

$$\mathcal{W}[\gamma, \mathfrak{G}]$$

be a wellbeing functional that evaluates properties such as dynamical stability, adaptability, constraint coherence, and recovery following perturbation. A general form may be written as

$$\mathcal{W} = w_S W_S + w_A W_A + w_C W_C + w_R W_R,$$

where

$$W_S, \quad W_A, \quad W_C, \quad W_R$$

measure stability, adaptability, constraint coherence, and resilience, respectively. The coefficients

$$w_S, \; w_A, \; w_C, \; w_R \geq 0$$

determine the relative importance assigned to these properties in a particular application.

This formulation separates two concepts that are often conflated. The variational principle determines the admissible cognitive dynamics, whereas the wellbeing functional evaluates whether those dynamics possess properties associated with healthy adaptation.

A wellbeing-relevant regime may therefore be characterized by a cognitive trajectory and geometry satisfying

$$\delta S = 0$$

together with bounded perturbation sensitivity and sustained constraint coherence.

This distinction is important because it prevents the theory from making the strong claim that every stationary cognitive configuration is necessarily desirable.

10.3 Stability and Cognitive Resilience

A central property of wellbeing-relevant cognition is the ability to absorb perturbations without producing disproportionate changes in cognitive behavior.
Let

$$(\gamma^*, \mathfrak{G}^*)$$

denote a stationary coupled solution and consider perturbations

$$\gamma = \gamma^* + \delta\gamma, \qquad \mathfrak{G} = \mathfrak{G}^* + \delta\mathfrak{G}.$$

The linearized coupled variational dynamics may be represented schematically as

$$\frac{d}{dt}\begin{pmatrix} \delta\gamma \\ \delta\mathfrak{G} \end{pmatrix} = \mathbf{J}(t)\begin{pmatrix} \delta\gamma \\ \delta\mathfrak{G} \end{pmatrix}$$

where

$$\mathbf{J}(t)$$

is the linearized evolution operator around the stationary solution. A local resilience measure can then be defined by

$$R_T = \sup_{t \in [0,T]} \frac{\|\delta\gamma(t)\| + \eta\,\|\delta\mathfrak{G}(t)\|}{\|\delta\gamma(0)\| + \|\delta G(0)\|},$$

where

$$\eta > 0$$

sets the relative scale between trajectory and geometric perturbations. Small values of

$$R_T$$

indicate that perturbations remain contained over the observation interval, whereas large values indicate increased sensitivity of the coupled cognitive system.

More generally, resilience may be characterized through the spectrum of the linearized operator

$$\mathbf{J}.$$

Configurations for which perturbations decay correspond to locally stable regimes, while configurations exhibiting growing modes may indicate susceptibility to amplification.

This formulation provides a bridge between the abstract stability results of the preceding sections and measurable properties of cognitive systems.

10.4 Stress as Geometric Perturbation

Environmental stress can be represented as a perturbation of the fields that constrain cognitive evolution. In particular,

$$\mathrm{E} \to \mathrm{E} + \delta\mathrm{E}, \quad \mathrm{U} \to \mathrm{U} + \delta\mathrm{U}, \quad \mathrm{N} \to \mathrm{N} + \delta\mathrm{N}.$$

These perturbations alter the cognitive potential,

$$\mathrm{V}\,(\mathrm{E}, \mathrm{U}, \mathrm{N}) \to \mathrm{V}\,(\mathrm{E} + \delta\mathrm{E}, \mathrm{U} + \delta\mathrm{U}, \mathrm{N} + \delta\mathrm{N}),$$

and consequently modify both the stationary trajectories and, in the fully coupled theory, the Cognitive Geometry itself:

$$\mathfrak{G} \to \mathfrak{G} + \delta\,\mathfrak{G}.$$

The response of the system can therefore be studied through the sensitivity mapping

$$(\delta\mathrm{E}, \delta\mathrm{U}, \delta\mathrm{N}) \to (\delta\gamma, \delta\mathrm{G}).$$

A well-adapted cognitive system would be expected to exhibit bounded responses and recovery toward an admissible region of the Cognitive Geometry. Conversely, persistent amplification, loss of constraint compatibility, or irreversible geometric deformation may characterize regimes of maladaptive dynamics.

This interpretation does not assume that any particular geometric quantity is a clinical biomarker. Rather, it provides a mathematical language in which stress response and adaptation can be operationalized and subsequently compared with empirical measurements.

10.5 Geometric Interpretation of Dysregulation

The same formalism permits a qualitative distinction between stable adaptation and dysregulation.
Let

$$\Omega_{\mathrm{adm}} \subset \mathcal{G}$$

denote a region of Cognitive Geometry in which epistemic, utilitarian, and normative constraints remain mutually compatible.

A cognitively stable regime may then be represented by

$$\mathfrak{G}\,(t) \in \Omega_{\mathrm{adm}}$$

with bounded trajectory perturbations and controlled geometric evolution.

Dysregulation, in contrast, may correspond to one or more of the following geometric conditions:

1. increased sensitivity of cognitive trajectories to small perturbations;
2. loss of compatibility among epistemic, utilitarian, and normative constraint fields;
3. rapid or poorly controlled deformation of the Cognitive Geometry;
4. emergence of unstable modes in the linearized coupled dynamics;
5. failure to return toward an admissible region following environmental perturbation.

These conditions are mathematical hypotheses rather than diagnostic criteria. Their value lies in providing measurable quantities that may be compared with behavioral or neural indicators of cognitive stability and adaptation.

10.6 Recovery and Return Dynamics

An additional property relevant to wellbeing is recovery following perturbation.

Suppose that at time

$$t_0$$

the cognitive system is perturbed away from a stable configuration

$$(\,\gamma^*, \mathfrak{G}^*\,).$$

Define the deviation

$$\mathrm{D}(t) = \mathrm{d}\boldsymbol{\mathcal{M}}(\gamma(t), \gamma^*(t)) + \eta \mathrm{d}\mathcal{G}\,(\mathfrak{G}\,(t), \mathfrak{G}^*\,(t)),$$

where

$$\mathrm{d}\boldsymbol{\mathcal{M}}$$

and

$$\mathrm{d}\mathcal{G}$$

denote appropriate distances on trajectory and geometry spaces. A recovery condition may then be expressed as

$$\mathrm{D}(t) \to 0 \quad \text{as} \quad t \to \infty,$$

or, for finite observation intervals, as a monotonic or bounded reduction in

$$\mathrm{D}(t)$$

The associated recovery time,

$$\boldsymbol{\tau}_{\mathrm{R}} = \inf\,\{t > t_0 : \mathrm{D}(t) \leq \epsilon\,\mathrm{D}\,(t_0)\}\,,$$

provides a possible quantitative measure of cognitive recovery.

This formulation suggests that resilience should not be measured only by the magnitude of the immediate response to stress. The temporal dynamics of return toward a stable cognitive regime may provide an additional and potentially more informative measure.

10.7 Potential Experimental Program

The mathematical framework generates several empirical hypotheses that can be tested using behavioral, neural, or artificial cognitive systems.

1. Latent representations inferred from neural or behavioral data may exhibit measurable geometric structure, including locally stable metrics, transport structure, and curvature.
2. Changes in task demands, environmental context, or social conditions may produce measurable changes in the estimated Cognitive Geometry.
3. Stable cognitive regimes may correspond to regions of reduced sensitivity of the coupled trajectory–geometry dynamics.
4. Stressful perturbations may produce characteristic changes in the estimated geometry and constraint fields.
5. Recovery from perturbation may be associated with a return toward previously stable geometric regimes.
6. Individuals or artificial systems exhibiting greater behavioral resilience may display lower perturbation sensitivity or shorter geometric recovery times.
7. Persistent maladaptive behavior may correspond to unstable or highly sensitive regions of the coupled variational dynamics.

These hypotheses suggest a concrete empirical workflow. First, latent representations may be estimated from behavioral or neural observations. Second, local geometric quantities may be reconstructed from these representations. Third, perturbations in task, sensory, or social conditions may be introduced experimentally. Finally, the resulting trajectory and geometric responses may be compared with the predictions of the coupled variational model.

Such experiments would allow the theory to be evaluated independently of its interpretive claims: the central question would be whether the proposed geometric quantities improve the prediction of cognitive stability, adaptation, and recovery relative to appropriate non-geometric baselines.

10.8 Scope and Interpretation

The present formulation should be interpreted as a mathematical framework for generating and testing hypotheses about cognition relevant to mental wellbeing. It does not claim that

$$\mathcal{W}, \quad R_T, \quad \tau_{\mathrm{R}},$$

or any other geometric quantity introduced here is itself a validated clinical measure.

In particular, the theory does not provide a diagnostic procedure for neurological or psychiatric disorders. Establishing clinical validity would require independent empirical studies, appropriate population samples, validated psychological and neurological measures, and careful comparison with existing clinical models.

The contribution of the present framework is instead to provide a common mathematical language in which stability, adaptation, perturbation response, and recovery can be studied as properties of coupled cognitive dynamics.

Box Summary. Mental wellbeing is treated as an application domain of the Field Theory of Representations rather than as its definition. Epistemic evidence, utilitarian objectives, and normative constraints interact through the potential V (E, U, N), generating constrained cognitive trajectories on an evolving Cognitive Geometry. Wellbeing-relevant regimes are hypothesized to combine stability, adaptability, constraint coherence, and recovery from perturbation. Stress can be modeled as a perturbation of the constraint fields, while resilience can be quantified through the sensitivity and recovery of the coupled trajectory–geometry system. These constructions generate testable hypotheses for behavioral, neural, and artificial cognitive systems, but should not be interpreted as validated clinical measures without empirical validation.

11 Discussion

The contribution of this paper is mathematical rather than architectural. It proposes a common formalism in which representation, reasoning, learning, and constraint evolution can be described as coupled geometric dynamics.

11.1 Relation to Cognitive Architectures

Classical cognitive architectures such as ACT-R, Soar, Sigma, and CLARION decompose cognition into interacting functional mechanisms [3,15].

The present framework addresses a different question. Rather than proposing another computational architecture, it asks whether these processes can be represented within a common mathematical language.

The correspondence is

$$\begin{aligned}
\text{representation} &\rightarrow \text{geometry},\\
\text{reasoning} &\rightarrow \text{variational trajectory},\\
\text{learning} &\rightarrow \text{geometry evolution},\\
\text{decision} &\rightarrow \text{constraint evaluation}.
\end{aligned}$$

The framework is therefore architecture-independent. Symbolic, neural, hybrid, and embodied systems could potentially instantiate different realizations of the same geometric formalism.

11.2 Relation to Active Inference

The framework is conceptually related to Active Inference and the freeenergy principle [7, 8].

Both approaches use variational principles and emphasize adaptive interaction between internal states and the environment.

The principal distinction is the mathematical object being optimized. Active Inference generally formulates variational free energy over probabilistic beliefs and generative models. The present framework formulates action functionals over latent representation geometries.

The two perspectives may therefore be complementary. Probabilistic representation spaces can be treated as particular examples of geometric structures, while variational free-energy objectives can provide particular choices of cognitive functionals.

11.3 Relation to Information Geometry

Information geometry provides a mathematical precedent for assigning metrics, connections, and curvature to spaces of statistical models [1, 2].

The present framework differs primarily in its object of interpretation. The primary geometry is intended to represent latent semantic organization rather than merely the parameter manifold of a statistical model.

Nevertheless, information-geometric tools may provide practical methods for constructing or estimating components of $\mathfrak{G}$.

11.4 Relation to Representation Learning

Modern representation learning produces high-dimensional latent spaces. The present theory proposes treating such spaces as empirical candidates for $\mathcal{M}$.

This creates an experimentally testable program:

trained representation → estimated geometry → trajectory analysis → geometry evolution.

The theory therefore does not require a replacement of neural networks. Instead, neural networks may serve as computational realizations from which geometric quantities can be estimated.

11.5 Theoretical Contribution

The main theoretical contribution is the explicit separation and coupling of two variational levels:

$$\gamma(t) \in \mathcal{M} \text{ and } \mathfrak{G}\,(t) \in \mathcal{G}\,.$$

This allows representation learning to be formulated as geometry evolution rather than simply parameter optimization.

The resulting product-space formulation,

$$(\,\gamma, \mathfrak{G}\,) \in \mathcal{M} \times \mathcal{G}\,,$$

provides a common domain for reasoning and learning.

11.6 Practical Significance

The framework suggests practical directions in:

- geometric representation learning;
- continual learning;
- transfer learning;
- adaptive planning;
- interpretable latent-state models;
- alignment and constraint-aware learning;
- multi-agent representation sharing;
- computational models of cognitive resilience.

The practical significance is therefore prospective. The present paper does not claim that these systems have already been implemented or validated.

Box Summary. The framework is complementary to existing cognitive architectures, Active Inference, information geometry, and representation learning. Its distinctive proposal is to treat the latent representation geometry itself as a dynamical object coupled to reasoning trajectories.

12 Limitations

The framework is intentionally theoretical, and several limitations must be made explicit.

First, the paper does not provide a general algorithm for estimating

$$\mathfrak{G} = (\mathcal{M}, g, \nabla, E, U, N).$$

from arbitrary empirical data.

Second, the relationship between neural representation spaces and the proposed Cognitive Geometry remains an empirical question.

Third, the fields

$$E, U, N$$

require operational definitions in concrete applications.

Fourth, the analytical results depend on assumptions concerning admissibility, coercivity, compactness or reflexivity, differentiability, and convexity. Real neural systems need not satisfy these assumptions globally.

Fifth, stability of a variational minimizer is not automatically equivalent to behavioral or clinical stability. Additional dynamical assumptions are required to establish such a correspondence.

Sixth, the mental wellbeing interpretation is a theoretical application program and should not be interpreted as a validated medical or psychiatric model.

Finally, the theory does not yet establish that the proposed geometric quantities improve empirical AI performance relative to existing methods.

These limitations define rather than invalidate the research program. They identify the measurements, algorithms, and experiments required to test the framework.

13 Future Directions

Several research directions follow directly from the theory.

13.1 Differentiable Geometric Learning

A first direction is the construction of differentiable models in which

$$g, \quad \nabla, \quad E, \quad U, \quad N$$

are learned jointly with task representations.

The optimization problem would take the form

$$\min_{\gamma,\, \mathfrak{G}} \mathcal{J}\,[\gamma, \mathfrak{G}].$$

This would provide a computational realization of the coupled variational framework.

13.2 Empirical Geometry Estimation

A second direction is the estimation of geometric quantities directly from latent representations.

Given observations

$$\{x_i\}_{i=1}^{n},$$

one could estimate local metrics, connections, and curvature and test whether these quantities predict reasoning trajectories or behavioral transitions.

13.3 Continual Learning

The geometry-evolution interpretation suggests algorithms that constrain updates to preserve selected geometric invariants:

$$I_k(\mathfrak{G}_{t+1}) \approx I_k(\mathfrak{G}_t).$$

Such invariants could potentially reduce catastrophic representational drift.

13.4 Transfer and Meta-Learning

Geometric transfer maps could be learned between related tasks:

$$\Phi_{ij} : \mathfrak{G}_i \to \mathfrak{G}_j.$$

Meta-learning could optimize the class of permissible maps and geometric adaptation laws rather than only task-specific parameters.

13.5 Embodied Agents

An embodied system naturally closes the loop:

$$\text{environment} \to \text{perception} \to \mathcal{M} \to \gamma \to \text{action} \to \text{environment}.$$

The geometry may then evolve through experience, creating a complete agent–environment variational loop.

13.6 Trustworthy AI

Future work should operationalize trustworthiness as measurable properties of the coupled system.

Possible quantities include:

$$\begin{aligned} &\text{robustness} \to \text{sensitivity of } \gamma \text{ to perturbations},\\ &\text{alignment} \to \text{constraint preservation},\\ &\text{interpretability} \to \text{recoverability of geometric structure},\\ &\text{safety} \to \text{invariance of admissible regions}. \end{aligned}$$

13.7 Mental Wellbeing

For mental wellbeing, future studies could combine behavioral, physiological, and neural data to estimate latent cognitive trajectories and investigate whether changes in geometric stability correlate with resilience, stress adaptation, or dysregulation.

Such studies would be necessary before making any clinical interpretation.

Box Summary. The principal next step is empirical realization: estimate Cognitive Geometries from data, implement coupled geometry–trajectory optimization, and test whether geometric quantities predict adaptation, robustness, decision-making, and cognitive resilience.

14 Conclusion

We have developed a variational framework for adaptive cognition in which reasoning trajectories and the representation geometry supporting those trajectories evolve jointly.

The central mathematical distinction is

$$\boxed{\begin{aligned} \gamma(t) &\subset \mathcal{M},\\ \mathfrak{G}(t) &\in \mathcal{G}. \end{aligned}}$$

Reasoning is represented by variational evolution of γ on a current Cognitive Geometry, while learning is represented by variational evolution of $\mathfrak{G}$ in the space of admissible geometries.

The complete cognitive state is therefore

$$\boxed{C(t) = (\gamma(t), \mathfrak{G}(t)).}$$

The central feedback loop is

$$\boxed{\mathfrak{G}_t \to \gamma_t \to \mathfrak{G}_{t+1}.}$$

This formulation provides a unified mathematical language for representation, reasoning, learning, attention, memory, decision-making, and adaptation.

The epistemic, utilitarian, and normative fields

$$E, U, N$$

provide a mechanism through which different requirements can interact without being reduced a priori to a single cognitive objective.

Several phenomena commonly treated as separate mechanisms can then be interpreted as different regimes of the same coupled system:

$$\begin{aligned} &\text{adaptation} \to \text{geometry evolution},\\ &\text{transfer} \to \text{geometric transport},\\ &\text{meta-learning} \to \text{evolution of adaptation laws},\\ &\text{continual learning} \to \text{continuous geometric evolution},\\ &\text{collective cognition} \to \text{coupled product-space dynamics}. \end{aligned}$$

The framework also provides a possible mathematical foundation for trustworthy AI. Interpretability can be studied through explicit geometric structure, alignment through constraint fields, safety through admissible geometry evolution, and robustness through stability and sensitivity analysis.

Mental wellbeing represents one important application domain because it requires simultaneous consideration of epistemic adequacy, goal-directed behavior, and normative or social constraints. The present framework does not claim to provide a clinical theory; rather, it identifies mathematical quantities whose empirical relationship to cognitive resilience and dysregulation can be investigated.

The principal claim of the paper is therefore not that cognition is already known to obey the proposed equations. It is that adaptive cognition can be formulated as a mathematically coherent coupled variational problem in which representation geometry and reasoning dynamics are treated as simultaneous objects of study.

In this sense, cognition is not merely dynamics within a representation space. It is the coupled evolution of trajectories and the geometric structure from which those trajectories emerge.

Final Box Summary.

Adaptive cognition = reasoning dynamics + evolving representation geometry.

The framework provides a common variational language connecting latent representation, reasoning, learning, epistemic evidence, utility, normative constraints, trustworthy AI, and potential models of mental wellbeing. Its central empirical challenge is now to determine whether the proposed geometric quantities can be estimated from real cognitive and artificial systems and whether they predict behavior beyond existing representation and learning models.

Conflict of Interest and Funding

The author declares no conflict of interest and received no specific funding for this work.

References

1. Amari S. Information Geometry and Its Applications. Tokyo: Springer; 2016.
2. Amari S, Nagaoka H. Methods of Information Geometry. Providence, RI: American Mathematical Society; 2000.
3. Anderson JR, Bothell D, Byrne MD, Douglass S, Lebiere C, et al. An integrated theory of the mind. Psychol Rev. 2004; 111: 1036-1060.
4. Beer RD. Dynamical approaches to cognitive science. Trends Cogn Sci. 2000; 4: 91-99.
5. Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU. Complex networks: Structure and dynamics. Phys Rep. 2006; 424: 175-308.
6. Bronstein MM, Bruna J, LeCun Y, Szlam A, Vandergheynst P. Geometric deep learning: Going beyond Euclidean data. IEEE Signal Process Mag. 2017; 34: 18-42.
7. Friston K. The free-energy principle: A unified brain theory? Nat Rev Neurosci. 2010; 11: 127-138.
8. Friston KJ, Parr T, Pezzulo G. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. Cambridge, MA: MIT Press; 2022.
9. Friston KJ, Stephan KE, Montague PR, Dolan RJ. Computational psychiatry: The brain as a phantastic organ. Lancet Psychiatry. 2014; 1: 148-158.
10. Huys QJM, Maia TV, Frank MJ. Computational psychiatry as a bridge from neuroscience to clinical applications. Nat Neurosci. 2016; 19: 404-413.
11. Montague PR, Dolan RJ, Friston KJ, Dayan P. Computational psychiatry. Trends Cogn Sci. 2012; 16: 72-80.
12. Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, et al. Early-warning signals for critical transitions. Nature. 2009; 461: 53-59.
13. Siregar E. AI with symbolic empathy: Shannon-Neumann insight guided logic. Cogn Comput. 2026; 18: 7.
14. Siregar E. Why safe AI starts with how we design it, not how we control it. The Science Matters. 2026. https://thescience-matters.org/why-safe-ai-starts-with-how-we-design-it-not-how-we-control-it/
15. Sun R. The CLARION cognitive architecture: Extending cognitive modeling to social simulation. In: Sun R, ed. Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation. New York, NY: Cambridge University Press; 2005: 79-100.
16. Thelen E, Smith LB. A Dynamic Systems Approach to the Development of Cognition and Action. Cambridge, MA: MIT Press; 1994.