Towards a Langlands Program for Collective Intelligence: A Categorical Foundation for Variational Dynamics on Constrained Spaces

Jonas Hallgren
Equilibria Network
jonas@eq-network.org

February 27, 2026

Abstract

This paper proposes steps towards a unified mathematical framework for collective intelligence that bridges game theory, social physics, active inference, and computational social science through categorical foundations. We propose that diverse collective intelligence mechanisms can be understood as constrained variational dynamics operating on graph-theoretic interaction spaces. By conceptualizing different collective systems as specialized constraint categories acting on the same underlying mathematical structure, we point at correspondences between equilibrium concepts, phase transitions, and institutional emergence. We hope that this categorical perspective might enable compositional modeling of hybrid systems, multi-scale analysis connecting micro-behaviors to macro-patterns, and principled design of robust coordination mechanisms. We point at how symmetries might constrain the space of possible transformations. How conservation laws might emerge through from applying Noether's theorem, and how renormalization could be used as a foundation for coarse-graining and finding collective agents.

1 Introduction

As artificial intelligence (AI) systems increasingly interact, they form complex networks of relationships giving rise to emergent behaviors across multiple scales [30, 22]. These collective phenomena—ranging from the coordination of autonomous vehicles to the dynamics of algorithmic trading—reveal patterns that transcend the properties of individual agents. Recent work highlights the potential for such multi-agent systems to exhibit destabilizing dynamics, emergent forms of agency, and cascading effects through interconnected networks [20].

Understanding and steering these collective phenomena is crucial, presenting a fascinating opportunity for collaboration across disciplines. Game theory and cooperative AI offer precise tools for modeling strategic interactions [12]. Social physics provides methods for analyzing emergent behaviors and phase transitions [5]. Active inference brings sophisticated techniques for understanding adaptive behavior and belief updating [17]. While each tradition captures essential aspects, a unified mathematical language that preserves these diverse insights while enabling systematic integration remains an open challenge [12].

This challenge parallels historical developments in other fields. Just as the Langlands program revealed deep connections between number theory and harmonic analysis, we propose that profound structural correspondences exist between seemingly disparate approaches to collective intelligence. Uncovering these connections promises not only to enhance theoretical understanding but also to provide practical tools for designing and governing the increasingly complex ecosystems of interacting AI agents [17].

We propose that **categorical foundations** interpreted through **variational dynamics on constrained spaces** offer a powerful unifying framework. Graphs naturally express both discrete agent interactions and continuous field-like properties, bridging complementary perspectives [36, 8]. By conceptualizing multi-agent systems as transformations of graph states encoding relational information, rather than mere collections of independent agents, we can capture local strategic interactions and global emergent patterns within a common mathematical framework [15].

This categorical, graph-centric approach enables three fundamental capabilities crucial for a unified theory:

1. **Compositional Representation**: Diverse collective intelligence mechanisms—networks, markets, hierarchies—can be expressed as distinct constraint categories acting on the same underlying graph structure, revealing their essential similarities and differences [11].

2. **Multi-scale Analysis**: System properties can be examined across scales, from local agent behaviors encoded in node/edge attributes to emergent network structures, providing insights into how agency and intelligence manifest at different levels of organization [30, 17].

3. **Dynamic Evolution**: The temporal evolution of complex systems can be modeled as sequences of graph transformations that preserve mathematical structure while capturing essential changes, often driven by variational principles like free energy minimization [26, 22].

This perspective draws inspiration from field theories in physics, focusing on transformations of relational structures rather than individual agent trajectories. Its strength lies in its universality; the same formalisms can describe information flow in neural networks, preference aggregation in markets, or consensus formation in democratic systems, reflecting deep mathematical isomorphisms [17].

This paper serves as an invitation to researchers across disciplines to collaboratively develop this unified mathematical language for collective intelligence. By establishing clear mathematical connections between traditionally separate domains, we aim to reveal hidden structures and foster new avenues for insight and application. Our goal is not to replace existing approaches but to enrich them mutually through a common foundation grounded in graph transformations and variational information dynamics.

2 Categorical Perspectives on Collective Intelligence

When multiple agents share a common resource, a fundamental tension emerges between individual benefit and collective sustainability. This scenario—exemplified by fisheries, grazing lands, and atmospheric commons—provides an ideal test case for comparing different theoretical approaches to collective intelligence [12].

We ask the question: what mathematical structures best capture the dynamics that lead to either sustainable cooperation or collective failure?

In order to find the answer to this question, we will first take a look at the existing research. First, we construct a general category that we can start off our exploration from. We then point out four research traditions that have developed powerful but distinct approaches to this question.

2.1 The Categorical Foundation

We begin by formalizing interaction spaces as categorical objects:

Definition 1 (Interaction Space)**.** *An interaction space* $S = (\mathcal{G}, \mathcal{F}, \mathcal{C})$ *consists of:*

- $\mathcal{G}$*: The interaction graph encoding agent connections and resource flows*
- $\mathcal{F}$*: The information flow structure determining observability and influence patterns*
- $\mathcal{C}$*: The constraint algebra specifying rules governing belief updates and action selection*

Different mathematical traditions can now be understood as imposing different constraint categories on these interaction spaces, leading to characteristic dynamics and equilibrium behaviors.

2.2 Four Categorical Perspectives on Resource Management

To understand how these approaches complement each other, we can compare how each tradition models the same resource management challenge. These are simplified versions of the actual descriptions in these fields and there are a lot of caveats to be made yet we're looking to point out the overview, rather than the detailed view:

Tradition	**Mathematical Formulation**	**Key Variables**	**Constraint Category**
Game Theory	$U_i(a_i, a_{-i}) = r_i(a_i, a_{-i}) + \gamma V_i(s')$	Agent utilities, reward functions, action strategies	$\mathcal{C}_{GT}$: Factorization constraints
Social Physics	$H = -\sum_{ij} J_{ij}\sigma_i\sigma_j - \sum_i h_i\sigma_i$	Spin configurations, coupling strengths, external fields	$\mathcal{C}_{SP}$: Energy conservation
Active Inference	$F = D_{KL}[q(\phi)\|\|p(\phi\|m)] - \mathbb{E}_q[\log p(o\|\phi)]$	Beliefs, generative models, free energy	$\mathcal{C}_{ActInf}$: Markov blanket preservation
Computational Social Science	$s_{t+1} = M(s_t, \{a_i^t\}_{i=1}^N, \xi_t)$	Agent actions, system states, stochastic factors	$\mathcal{C}_{CSS}$: Institutional constraints

Table 1: Categorical Perspectives on Resource Management

Each mathematical framework reveals different aspects of the same underlying categorical structure:

From the Game Theory lens, resource management appears as a sequential social dilemma [29]—a repeated game where agent i's utility (U_i) depends on both immediate rewards (r_i) and expected future value (V_i). The constraint category $\mathcal{C}_{GT}$ enforces factorization of agent beliefs and utility maximization. This perspective highlights the misalignment between individual incentives and collective welfare, revealing temporal trade-offs between immediate gains and future sustainability. Researchers have developed sophisticated mechanisms to bridge this gap, including commitment strategies [21], intrinsic social motivation [24], and institutional design [27, 12]. However, this approach faces computational complexity when the constraint category becomes too rigid for large-scale systems [19].

Through the Social Physics lens, the same problem appears as a dynamical system with multiple stable states [5]. The constraint category $\mathcal{C}_{SP}$ imposes energy conservation and nearest-neighbor interactions. The Graph Laplacian $L = D - W$ emerges as a fundamental operator, with eigenvalue spectra controlling diffusion dynamics and phase transition boundaries [39]. The spectral gap λ_2 determines mixing times and convergence rates of collective belief updating. This view reveals how seemingly stable systems can experience sudden, irreversible collapse when critical thresholds are crossed. Researchers have identified early warning signals for impending collapse [37], tipping points that trigger regime shifts [38], and resilience measures that quantify a system's ability to withstand perturbations [23, 1].

The Active Inference lens operates under constraint category $\mathcal{C}_{ActInf}$ which preserves Markov blanket structure while minimizing variational free energy [17, 28]. This bridges individual cognition and collective behavior through unified information-theoretic principles, with agents maintaining beliefs about hidden states while acting to minimize surprise [16, 26, 22]. This perspective connects directly to the Free Energy Principle in computational neuroscience and provides a framework for understanding how different mechanisms (consensus, optimization) emerge as strategies for minimizing variational objectives within the constraints of the information topology [7, 9, 10].

The Computational Social Science lens employs constraint category $\mathcal{C}_{CSS}$ which allows institutional emergence while maintaining network structure [6]. The system state (s_t) evolves based on agent actions ($\{a_i^t\}$) and stochastic factors (ξ_t). This approach illuminates how norms and institutions emerge to govern resource use and emphasizes the importance of communication and monitoring mechanisms. Researchers have identified design principles for successful governance [32], developed field methods for measuring cooperation [25], and studied how social network structure affects outcomes [2].

2.3 The Universal Induction Functor

We now turn to the generation of a potential categorical basis that may describe all of these systems. This is a model and all models are wrong, some are useful, we hope this one is useful.

Principle 1 (A Relational Functor for Collective Dynamics)**.** *There exists a natural functor $\mathbf{F} : \mathcal{CI} \to \mathcal{DYN}$ that maps the category of collective intelligence structures to the category of their dynamical systems. This mapping is defined as follows:*

- ***Objects****: The functor maps a specific relational structure of a collective, S, to the dynamical system, $\mathcal{D}_S$, that governs its evolution.*

  An object $S \in \mathcal{CI}$ is a tuple $S = (\mathcal{A}, \mathcal{G}, \mathcal{R})$ representing the instantaneous, static architecture of the collective, where:

  - *$\mathcal{A}$ is the set of agents, defined by their internal states and generative models (i.e., their policies and beliefs).*
  - *$\mathcal{G}$ is the interaction graph, defined by its topology of nodes and edges, representing the pathways of influence and observation.*
  - *$\mathcal{R}$ is the embodied credit assignment mechanism (e.g., market prices, voting rules, hierarchical reviews), which specifies the prior preferences of the collective agent.*

  The corresponding object $\mathbf{F}(S) = \mathcal{D}_S$ in $\mathcal{DYN}$ is the dynamical system whose temporal evolution is constrained by a gradient flow on a free energy landscape:

$$\dot{\mathbf{s}} = -\nabla F_S(\mathbf{s}) \tag{1}$$

where the variational free energy functional F_S is canonically determined by the collective's relational structure $S = (\mathcal{A}, \mathcal{G}, \mathcal{R})$.

- ***Morphisms**: The functor maps a structural transformation in the collective, ϕ, to the corresponding transformation of its dynamics, T_ϕ.*

  A morphism $\phi : S \to S'$ in $\mathcal{CI}$ represents a structural change to the collective (e.g., an agent learning and updating its model in $\mathcal{A}$, a network link being added to $\mathcal{G}$, or an institutional rule changing in $\mathcal{R}$).

  The corresponding morphism $\mathbf{F}(\phi) = T_\phi : \mathcal{D}_S \to \mathcal{D}_{S'}$ in $\mathcal{DYN}$ represents the induced transformation on the space of dynamics itself. This reflects how the evolution of the collective's structure reshapes its own future evolutionary landscape.

This principle asserts that the dynamics of a collective intelligence are not arbitrary but are a direct, predictable consequence of its underlying relational architecture, governed by the universal imperative to minimize variational free energy.

2.4 Complementary Strengths and Limitations

While each perspective offers valuable insights, they also have characteristic limitations that reflect what they were designed to illuminate:

Tradition	**Key Strengths**	**Characteristic Limitations**
Game Theory	• Precise modeling of strategic interactions • Clear identification of incentive misalignments • Principled mechanism design approaches	• Typically limited to small agent populations • Often assumes complete information • Struggles with heterogeneous agent capabilities
Social Physics	• Identification of system-level phase transitions • Powerful prediction of critical thresholds • Modeling of self-organization from local interactions	• Often abstracts away individual strategic behavior • Struggles with heterogeneous agent incentives • Typically descriptive rather than prescriptive
Active Inference	• Unifying theoretical framework across scales • Principled uncertainty handling and belief updating • Natural multi-agent coordination mechanisms	• Computational complexity of belief propagation • Limited incorporation of learning and adaptation • Requires simplification of complex scenarios
Computational Social Science	• Rich modeling of emergent norms and institutions • Incorporation of empirical human behavior • Attention to communication and monitoring	• Often lacks formal mathematical guarantees • Results may depend on implementation details • Challenges in generalizing across contexts

Table 2: Complementary Strengths and Limitations

These complementary patterns of strengths and limitations aren't merely coincidental—they reflect fundamental differences in what each approach is designed to illuminate. Game theory excels at modeling strategic interactions but struggles with scale [12]. Social physics effectively identifies emergent patterns but often abstracts away agent heterogeneity [5]. Active inference provides principled accounts of adaptive behavior but faces computational challenges [17]. Computational social science offers rich empirical insights but sometimes lacks formal rigor [6].

2.5 Toward a Unified Framework

These complementary strengths suggest the value of an integrated approach. A unified mathematical framework would ideally:

- Preserve the strategic reasoning capabilities of Game Theory
- Incorporate the phase transition insights of Social Physics
- Maintain the principled adaptation mechanisms of Active Inference
- Include the institutional emergence aspects of Computational Social Science

We propose that graph theory provides a natural integration framework [8]. By representing resource management systems as graphs where:

- Nodes represent agents and resources
- Edges represent interactions (harvesting, monitoring, communication)
- Node attributes capture state variables (resource levels, agent preferences)
- Edge attributes represent interaction parameters (harvesting rates, communication channels)

The system dynamics can then be expressed as transformations of this graph structure:

$$G_{t+1} = \mathcal{T}_{\mathcal{C}}(G_t) \tag{2}$$

Where G_t represents the graph at time t and $\mathcal{T}_{\mathcal{C}}$ is a transformation operator that updates the graph based on agent actions and system dynamics while respecting the constraints imposed by category $\mathcal{C}$ [11].

This representation allows us to define operators that correspond to key concepts across disciplines:

Concept	Discipline	Graph Representation
Nash Equilibrium	Game Theory	Stable node strategy profiles where no agent has incentive to change [19]
Phase Transition	Social Physics	Critical changes in graph connectivity patterns as control parameters vary [5]
Free Energy Minimum	Active Inference	Attractors where Markov blanket structure is preserved [17]
Institutional Rules	Computational Social Science	Constraints on permissible graph transformations that maintain system integrity [32]

Table 3: Mapping Disciplinary Concepts to Graph Representations

This unified representation enables new forms of analysis that were previously difficult to achieve, including:

1. Understanding how strategic interactions scale with increasing numbers of agents by examining how graph properties change with size [8].
2. Identifying critical nodes and edges whose removal would trigger phase transitions in system behavior [5].
3. Characterizing stability conditions in terms of graph spectral properties [39].
4. Studying how different institutional rules constrain the evolution of graph structure [32, 11].

These capabilities point toward a deeper understanding of collective intelligence that draws on the complementary strengths of different research traditions [30]. By providing a common mathematical language based on graph theory, we create new possibilities for both theoretical integration and practical application [36].

3 Variational Dynamics and Categorical Graph Transformations

In Section 2, we examined how different disciplines represent collective intelligence systems through graph-theoretic structures. While these representations capture the instantaneous state of collective systems, they leave open the crucial question of how these structures evolve over time. This section addresses that question by proposing a unifying framework centered on **graph transformations** as the fundamental objects of study [11]. Building upon this, we integrate concepts from information theory, graph signal processing, and variational principles to formalize how these transformations process information constrained by the underlying graph topology [41, 39, 13, 17?].

3.1 Graph Dynamics as Information Processing

The examples from Section 2 illustrate that collective intelligence dynamics, across diverse modeling traditions, can be expressed as transformations acting on attributed graphs $G = (V, E, \alpha, \beta)$.

- Game Theory: $G_{t+1} = T_{game}(G_t)$ (Strategy/attribute updates)
- Social Physics: $G_{t+1} = T_{physics}(G_t, \theta, \xi_t)$ (Regime shifts)
- Active Inference: $G_{t+1} = T_{inference}(G_t)$ (Belief updates)
- Computational Social Science: $G_{t+1} = T_{social}(G_t, \xi_t)$ (Institutional evolution)

These transformations, regardless of their specific form, represent processes that update the system's state based on information. We formally define the transformation:

Definition 2 (Graph Transformation Function)**.** *A graph transformation function $T : \mathcal{G} \to \mathcal{G}$ maps from the space of attributed graphs $\mathcal{G}$ to itself, potentially modifying vertex attributes α, edge attributes β, and graph topology (V, E). T encapsulates the information processing rules of the collective intelligence system [11].*

The core idea is that the graph G itself defines an **information topology** [41, 39], constraining how information flows and how updates (T) are performed. This topology can be captured by the graph's adjacency structure W, its Laplacian $L = D - W$, or the precision matrix Q of an equivalent probabilistic model like a Gaussian Markov Random Field (GMRF), where non-zero entries signify conditional dependencies [41].

3.2 Variational Dynamics and Information-Theoretic Measures

A powerful unifying principle, supported by findings across multiple relevant fields [17, 40, 22], suggests that collective intelligence dynamics T often drive the system towards states that minimize a variational objective, typically related to **prediction error** or **surprise**. This can be formalized using Kullback-Leibler (KL) divergence or its bound, the **variational free energy**.

If $q(G_{t+1})$ represents the distribution of states the system actually transitions to under its intrinsic dynamics and stochastic influences, and $p(G_{t+1}|G_t, T)$ represents a generative model's prediction (implicitly encoded by the agent or system structure and its transformation rules T), the dynamics can often be understood as implicitly minimizing:

$$\mathcal{L}(G_t) = D_{KL}(q(G_{t+1}) \parallel p(G_{t+1}|G_t, T)) \tag{3}$$

Minimizing this functional aligns the system's state with its implicit model (or predictions), reducing surprise and enabling adaptive behavior, coordination, and learning [16, 26, 22]. This perspective connects directly to the Free Energy Principle in computational neuroscience [17, 28] and provides a framework for understanding how different mechanisms (consensus, optimization) emerge as strategies for minimizing variational objectives within the constraints of the information topology [7, 9, 10].

Key aspects of these variational dynamics can be quantified using information-theoretic tools:

- **Information Flow:** Mutual information ($I(\alpha_i(t+1); \alpha_j(t))$) or related measures quantify the influence transmitted along the topological pathways defined by G_t.
- **Information Change/Surprise:** The KL-divergence itself (Eq. 3), or changes in state probability distributions $P(G)$, measure the information cost or surprise associated with the transformation T.
- **Information Integration & Complexity:** Graph entropy ($I_f(G)$ [13]), where the functional f probes local topological properties (e.g., derived from L or Q), quantifies the information content or structural complexity of the resulting state G_{t+1}. Other complexity measures (Φ) capture emergent organization. The variational dynamics potentially involve a trade-off between minimizing surprise and maintaining an appropriate level of structural complexity/entropy.

The state of the system G_t can also be viewed as a point on a statistical manifold endowed with an information metric (e.g., Fisher Information Metric derived from the graph structure L or Q), with transformations T tracing trajectories on this manifold [3, 41, 4].

3.3 The Unifying Lens: Variational Information Dynamics on Graphs

We propose that collective intelligence emerges from **variational dynamics operating on information topologies**. The graph structure G (represented topologically by (V, E) and geometrically/probabilistically by W, L, or Q) defines the constraints and pathways for information processing. Graph transformations T implement the system's dynamics, driven by the fundamental principle of minimizing surprise or prediction error (quantified information-theoretically, often via KL-divergence or variational free energy).

This perspective offers a simplified yet potent unifying lens:

- It grounds dynamics in the relational **graph structure** (information topology).
- It employs universal principles from **information theory and variational inference** (surprise minimization).
- It naturally integrates concepts from:
  - **Graph Signal Processing**: Utilizing L for spectral analysis, defining topology, and implementing filtering operations [39].
  - **Probabilistic Graphical Models**: Using Q (from GMRFs) for conditional independence structure and belief propagation [41, 33].
  - **Active Inference/Free Energy Principle**: Providing the variational optimization framework [17, 35, 22, 26].
  - **Network Science**: Employing graph entropy and complexity measures to characterize states [13**?**].

This framework, focused on graph transformations interpreted through variational information dynamics, satisfies the requirements established at the end of Section 2 for a unified mathematical language. It provides a foundation for describing, analyzing, and comparing diverse collective intelligence systems, and sets the stage for exploring the structural correspondences proposed in the next section.

4 Conjectured Mathematical Bridges

The framework developed in the previous sections, centered on graph transformations and variational information dynamics, suggests potential deep connections between existing, seemingly disparate mathematical approaches to collective intelligence. We propose the following conjectures as starting points for a research program aimed at uncovering these unifying structures.

Conjecture 1 (The Equilibria-Attractor Correspondence)**.** *Strategic equilibria in game-theoretic frameworks [35] and attractors in dynamical systems might represent manifestations of the same underlying mathematical structures, specifically as configurations that minimize variational free energy or related information functionals [18, 40], when viewed through appropriate transformations of graph representations.*

This tentative correspondence suggests a potential bridge between domains that have developed largely independently. While the complete relationship remains elusive, suggestive connections appear in specific cases: In potential games, Nash equilibria correspond to local minima of potential functions, which also represent fixed points (attractors) in gradient dynamical systems. In evolutionary settings with replicator dynamics, rest points often coincide with Nash equilibria of the underlying game. Network games exhibit equilibrium properties influenced by graph topology [31] in ways reminiscent of how connectivity shapes attractor basins in networked dynamical systems. Active Inference models explicitly connect policy selection (related to equilibria) with belief updating dynamics converging to attractors [15, 28, 22].

What makes this correspondence particularly intriguing is that both equilibria and attractors seem to represent configurations where certain transformations preserve structure—points where the system finds a form of stability or minimal surprise under the relevant dynamics [18]. This pattern hints at deeper mathematical connections waiting to be uncovered, potentially framed as fixed points or minima of the variational dynamics described in Section 3.

If this correspondence could be formalized more precisely, it might enable translation of analytical techniques (e.g., stability analysis from dynamics to game theory, equilibrium refinement concepts from game theory to dynamics) and suggest novel hybrid methodologies.

Conjecture 2 (The Scale Bridging Principle)**.** *There may exist systematic mathematical relationships connecting microscale agent interactions (local information processing) with macroscale system behaviors (emergent patterns), with particular classes of graph transformations (e.g., spectral filtering, diffusion, renormalization-like operations [14]) potentially serving as formal bridges between these levels of description.*

This conjecture addresses one of the central mysteries in collective intelligence: how do complex global patterns reliably emerge from simple local interactions? While complete mathematical formulations remain distant, we observe tantalizing patterns across domains: Statistical physics has developed renormalization group methods that systematically connect microscopic interactions with macroscopic phase transitions [14]. Opinion dynamics research reveals how network topology shapes global convergence properties from local updating rules [31]. Evolutionary systems demonstrate how spatial interaction structures influence population-level strategy diffusion. Active Inference provides a framework where local agent interactions based on minimizing free energy lead to emergent collective behaviors and shared models [18, 16, 22]. Graph Signal Processing explicitly develops tools for analyzing signals and processes across different scales on graphs [39].

Our graph-theoretic and variational framework suggests a tentative unifying perspective. Local transformations T_i, perhaps driven by local free energy minimization or belief propagation [15, 33], aggregate through the network structure (Q, L). The resulting global transformation T and emergent collective state might be understood through analyzing how these local operations compose and how information propagates across graph spectral [41, 39] or spatial scales.

The boundary between discrete graph-based representations and continuous field theories represents a particularly challenging frontier of this conjecture. In the limit of infinitely many nodes with infinitesimal influences, graph Laplacians converge to continuum operators [39] and graph transformations might approximate field transformations, yet the precise mathematical bridge remains incompletely understood. This boundary case appears across domains—from statistical physics to neural field theories to continuous opinion dynamics—suggesting its fundamental importance.

Conjecture 3 (The Computational Correspondence Principle)**.** *Computational problems arising in different mathematical approaches to collective intelligence (e.g., finding equilibria, predicting attractors, optimizing control, performing inference) may share fundamental structural similarities when they describe the same underlying phenomena, potentially enabling systematic translation of algorithms and complexity results across domains.*

This speculative correspondence suggests that the computational challenges encountered in different approaches might be more than superficially similar. Finding Nash equilibria in general games is PPAD-complete [19], while predicting long-term behavior in certain dynamical systems is computationally intractable—might these challenges represent manifestations of the same underlying computational structure related to the complexity of managing information flow and belief coordination on the graph?

Similarly, approximation algorithms for graph optimization problems often share structural characteristics with methods for approximating game-theoretic equilibria. Inference in graphical models (like Bayesian networks or factor graphs used in Active Inference and related fields [33, 28, 15]) and computation in graphical games exhibit related computational patterns and challenges [19]. These parallels hint at deeper connections waiting to be formalized.

Our framework suggests that if different mathematical approaches ultimately describe the same underlying variational dynamics on an information topology, then the computational challenges they present might admit formal relationships, potentially allowing algorithmic insights (e.g., efficient message passing schedules [28], complexity bounds) to transfer across domains.

4.1 From Conjectures to Research Program

These conjectured correspondences, while intriguing, remain primarily speculative. Transforming them into a rigorous mathematical program requires a methodological approach that acknowledges both their promise and the substantial work needed to establish their validity:

4.1.1 Research Methodology

Rather than attempting to establish general principles immediately, progress likely requires studying specific collective intelligence problems (e.g., resource management [34], consensus [7], collective motion [22], strategic interaction [35]) through multiple mathematical lenses. By examining concrete cases, we might discern structural patterns that generalize across approaches.

Where partial correspondences emerge, we can work to extend them through careful analysis of the underlying mathematical structures. This process might require developing new mathematical tools specifically designed to bridge existing frameworks—potentially drawing from:

- **Category Theory** to formalize mappings between different graph representations and transformation types.

- **Information Geometry** [3] to understand the structure of belief spaces on graphs and the dynamics of variational inference.
- **Graph Theory and Spectral Analysis** [41, 39] to characterize information topology and its role in dynamics and scale bridging.
- **Computational Complexity Theory** [19] to establish formal relationships between algorithmic challenges across different frameworks.
- **Variational Methods and Statistical Physics** [18] to provide the core optimization principles.

Computational experiments provide another essential avenue, allowing us to test conjectured correspondences against concrete scenarios (e.g., simulating factorised Active Inference agents [35] or GMRF dynamics [41]) and identify their boundaries and limitations. These experiments might reveal where correspondences break down as well as where they hold, both outcomes providing valuable insight.

4.1.2 Practical Horizons

Despite their speculative nature, these conjectured correspondences suggest promising directions for practical applications. If the Equilibria-Attractor Correspondence holds, spectral methods from graph theory [41, 39] might provide unified approaches for early warning signals of transitions in diverse collective systems [37].

The Scale Bridging Principle, if developed further, could inform multi-scale design approaches for engineering robust collective AI systems or intervening in social systems, perhaps leveraging insights from hierarchical active inference [18]. The Computational Correspondence Principle might enable systematic translation of efficient algorithms (e.g., message passing [28], variational inference techniques) between domains.

More broadly, these correspondences might facilitate knowledge transfer across the diverse landscape of collective intelligence research—from biological systems [22] to social networks [5] to artificial multi-agent ensembles [35, 33]—revealing common variational principles operating on information topologies.

The program outlined here is necessarily incomplete and preliminary. We offer these conjectures not as definitive claims but as an invitation to collaborative exploration of the mathematical foundations of collective intelligence [30, 12]. By working across traditional disciplinary boundaries, we may discover unexpected connections that deepen our understanding of how intelligent collectives function across domains—from neural systems to social groups to multi-agent AI.

In the spirit of the Langlands program, the greatest value may lie not in whether these specific conjectures prove exactly correct, but in the novel mathematical perspectives and methods that emerge from their investigation. We invite researchers from all relevant disciplines [8, 36] to join in refining, extending, or even refuting these proposed correspondences as we work toward a more unified understanding of collective intelligence.

5 Conclusion

We began with the challenge of fragmented approaches to collective intelligence across disciplines employing disparate mathematical formalisms. Through categorical analysis, we have shown this fragmentation to reflect different constraint categories operating on the same underlying mathematical structure rather than fundamental theoretical differences.

This categorical perspective reveals collective intelligence as emerging from constrained variational dynamics on graph-structured interaction spaces. Different mathematical traditions—game theory, social physics, active inference, and computational social science—impose different constraint categories that lead to characteristic forms of equilibrium and dynamics, but all operate through the same universal induction functor that maps categorical structure to collective behavior.

The synthesis provides several key insights:

- **Equilibrium types** emerge from constraint categories rather than being fundamental primitives
- **Strategic reasoning** becomes categorical belief propagation with constraints
- **Physical dynamics** and game-theoretic learning are instances of the same universal categorical operations
- **System design** becomes constraint category engineering

This categorical framework offers the mathematical infrastructure for a synthetic science of collective intelligence—moving beyond description to principled construction of coordination mechanisms. It provides a path beyond the computational limitations of traditional game theory and the reductionist abstractions of social physics toward a design theory for complex, multi-scale systems.

Most importantly, this approach transforms our relationship to collective intelligence from passive observation to active engineering. By understanding how categorical constraints shape collective behavior, we can design interaction spaces that naturally induce beneficial coordination patterns—systems that are not only powerful and efficient but also robust, fair, and aligned with human flourishing.

The future of collective intelligence lies not in predicting what emerges naturally, but in consciously designing the categorical constraints that govern emergence itself. This is a future where mathematics becomes not just a tool for analysis but a language for synthesis, enabling us to architect coordination patterns that serve human values and societal goals.

References

[1] Mahault Albarracin, Maxwell Ramstead, Riddhi J. Pitliya, Ines Hipolito, Lancelot Da Costa, Maria Raffa, Axel Constant, and Sarah Grace Manski. Sustainability under Active Inference. *Systems*, 12(5):163, May 2024.

[2] Michael J Alexander, Michael S Warren, Alex Chang, Emmi Ekaette, Raymon Fankun, Madian Nielson, Niki Pour, Philip Todd, Nicolas Turgeon, and Chao Yan. Network structure and community evolution on reddit. *arXiv preprint arXiv:1502.06222*, 2015.

[3] Shun-ichi Amari. *Information Geometry and Its Applications*, volume 194 of *Applied Mathematical Sciences*. Springer, 2016.

[4] Xerxes D. Arsiwalla, Pedro A.M. Mediano, and Paul F.M.J. Verschure. Spectral Modes of Network Dynamics Reveal Increased Informational Complexity Near Criticality. *Procedia Computer Science*, 108:119–128, 2017.

[5] Joseph B. Bak-Coleman, Mark Alfano, Wolfram Barfuss, Carl T. Bergstrom, Miguel A. Centeno, Iain D. Couzin, Jonathan F. Donges, Mirta Galesic, Andrew S. Gersick, Jennifer Jacquet, Albert B. Kao, Rachel E. Moran, Pawel Romanczuk, Daniel I. Rubenstein, Kaia J. Tombak, Jay J. Van Bavel, and Elke U. Weber. Stewardship of global collective behavior. *Proceedings of the National Academy of Sciences*, 118(27):e2025764118, July 2021.

[6] Eric Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the national academy of sciences*, 99(suppl_3):7280–7287, 2002.

[7] Benoît Bonnet, Nastassia Pouradier Duteil, and Mario Sigalotti. Consensus formation in first-order graphon models with time-varying topologies. *Mathematical Models and Methods in Applied Sciences*, 32(11):2121–2188, 10 2022.

[8] Katy Börner, Soma Sanyal, and Alessandro Vespignani. Network science. *Annual Review of Information Science and Technology*, 41(1):537–607, 1 2007.

[9] Neil Cammardella, Ana Bušić, and Sean P. Meyn. Kullback–Leibler-Quadratic Optimal Control. *SIAM Journal on Control and Optimization*, 61(5):3234–3258, oct 25 2023.

[10] Oliver Cliff, Mikhail Prokopenko, and Robert Fitch. Minimising the Kullback–Leibler Divergence for Model Selection in Distributed Nonlinear Systems. *Entropy*, 20(2):51, jan 23 2018.

[11] B. COURCELLE. *THE EXPRESSION OF GRAPH PROPERTIES AND GRAPH TRANSFORMATIONS IN MONADIC SECOND-ORDER LOGIC*, pages 313–400. WORLD SCIENTIFIC, 2 1997.

[12] Allan Dafoe, Edward Hughes, Yoram Bachrach, Tantum Collins, Kevin R. McKee, Joel Z. Leibo, Kate Larson, and Thore Graepel. Open Problems in Cooperative AI, December 2020.

[13] Matthias Dehmer. Information processing in complex networks: Graph entropy and information functionals. *Applied Mathematics and Computation*, 201(1-2):82–94, 2008.

[14] Gábor Fáth and Miklos Sarvary. A renormalization group theory of cultural evolution. *Physica A: Statistical Mechanics and its Applications*, 348:611–629, March 2005.

[15] Karl J. Friston, Thomas Parr, and Bert de Vries. The graphical brain: Belief propagation and active inference. *Network Neuroscience*, 1(4):381–414, 12 2017.

[16] Karl J. Friston, Thomas Parr, Conor Heins, Axel Constant, Daniel Friedman, Takuya Isomura, Chris Fields, Tim Verbelen, Maxwell Ramstead, John Clippinger, and Christopher D. Frith. Federated inference and belief sharing. *Neuroscience & Biobehavioral Reviews*, 156:105500, 1 2024.

[17] Karl J Friston, Maxwell JD Ramstead, Alex B Kiefer, Alexander Tschantz, Christopher L Buckley, Mahault Albarracin, Riddhi J Pitliya, Conor Heins, Brennan Klein, Beren Millidge, Dalton AR Sakthivadivel, Toby St Clere Smithe, Magnus Koudahl, Safae Essafi Tremblay, Capm Petersen, Kaiser Fung, Jason G Fox, Steven Swanson, Dan Mapes, and Gabriel René. Designing ecosystems of intelligence from first principles. *Collective Intelligence*, 3(1):26339137231222481, January 2024.

[18] Karl J Friston, Maxwell JD Ramstead, Alex B Kiefer, Alexander Tschantz, Christopher L Buckley, Mahault Albarracin, Riddhi J Pitliya, Conor Heins, Brennan Klein, Beren Millidge, Dalton AR Sakthivadivel, Toby St Clere Smithe, Magnus Koudahl, Safae Essafi Tremblay, Capm Petersen, Kaiser Fung, Jason G Fox, Steven Swanson, Dan Mapes, and Gabriel René. Designing ecosystems of intelligence from first principles. *Collective Intelligence*, 3(1), 1 2024.

[19] Paul W. Goldberg. A Survey of PPAD-Completeness for Computing Nash Equilibria, March 2011.

[20] Lewis Hammond, Alan Chan, Jesse Clifton, Jason Hoelscher-Obermaier, Akbir Khan, Euan McLean, Chandler Smith, Wolfram Barfuss, Jakob Foerster, Tomáš Gavenčiak, The Anh Han, Edward Hughes, Vojtěch Kovařík, Jan Kulveit, Joel Z. Leibo, Caspar Oesterheld, Christian Schroeder de Witt, Nisarg Shah, Michael Wellman, Paolo Bova, Theodor Cimpeanu, Carson Ezell, Quentin Feuillade-Montixi, Matija Franklin, Esben Kran, Igor Krawczuk, Max Lamparth, Niklas Lauffer, Alexander Meinke, Sumeet Motwani, Anka Reuel, Vincent Conitzer, Michael Dennis, Iason Gabriel, Adam Gleave, Gillian Hadfield, Nika Haghtalab, Atoosa Kasirzadeh, Sébastien

Krier, Kate Larson, Joel Lehman, David C. Parkes, Georgios Piliouras, and Iyad Rahwan. Multi-Agent Risks from Advanced AI, February 2025.

[21] Su Han. Commitment in networks. *Journal of Economic Theory*, 193:105226, 2021.

[22] Conor Heins, Beren Millidge, Lancelot Da Costa, Richard P. Mann, Karl J. Friston, and Iain D. Couzin. Collective behavior from surprise minimization. *Proceedings of the National Academy of Sciences*, 121(17), apr 17 2024.

[23] Crawford S Holling. Resilience and stability of ecological systems. *Annual review of ecology and systematics*, 4(1):1–23, 1973.

[24] Edward Hughes, Joel Z Leibo, Matthew G Phillips, Karl Tuyls, Edgar A Duéñez-Guzmán, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin R McKee, Raphael Koster, et al. Inequity aversion resolves intertemporal social dilemmas. *arXiv preprint arXiv:1803.08884*, 2018.

[25] Marco A Janssen, John M Anderies, and Elinor Ostrom. Lab experiments for the study of social-ecological systems. *Science*.

[26] Rafael Kaufmann, Pranav Gupta, and Jacob Taylor. An Active Inference Model of Collective Intelligence. *Entropy*, 23(7):830, jun 29 2021.

[27] Raphael Koster, Daniel Application, Computer Science from Degrees, et al. Model-based planning in discrete and continuous actions spaces. In *NeurIPS 2020 Workshop: Deep Reinforcement Learning*, 2020.

[28] Magnus Koudahl, Thijs van de Laar, and Bert de Vries. Realising Synthetic Active Inference Agents, Part I: Epistemic Objectives and Graphical Specification Language. *arXiv.org*, 2023.

[29] Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent reinforcement learning in sequential social dilemmas. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 464–473, 2017.

[30] Patrick McMillen and Michael Levin. Collective intelligence: A unifying concept for integrating biology across scales and substrates. *Communications Biology*, 7(1):1–17, March 2024.

[31] Wout Merbis and Manlio de Domenico. Emergent information dynamics in many-body interconnected systems. *Physical Review E*, 108(1), jul 31 2023.

[32] Elinor Ostrom. *Governing the commons: The evolution of institutions for collective action.* Cambridge university press, 1990.

[33] Francesco Palmieri, Krishna Pattipati, Giovanni Di Gennaro, Amedeo Buonanno, and Martina Merola. Multiple Agents Interacting via Probability Flows on Factor Graphs. In *AHFE International.* AHFE International, 2023.

[34] Giorgio Piatti, Zhijing Jin, Max Kleiman-Weiner, Bernhard Schölkopf, Mrinmaya Sachan, and Rada Mihalcea. Cooperate or Collapse: Emergence of Sustainable Cooperation in a Society of LLM Agents, December 2024.

[35] Jaime Ruiz-Serra, Patrick Sweeney, and Michael S. Harré. Factorised Active Inference for Strategic Multi-Agent Interactions. *arXiv.org*, 2024.

[36] Bornholdt S. and Schuster H. Handbook of Graphs and Networks. nov 8 2002.

[37] Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H Van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

[38] Marten Scheffer, Steve Carpenter, Jonathan A Foley, Carl Folke, and Brian Walker. Catastrophic shifts in ecosystems. *Nature*, 413(6856):591–596, 2001.

[39] David I Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. (arXiv:1211.0053v2), March 2013.

[40] D. H Wolpert, K. R Wheeler, and K Tumer. Collective intelligence for control of distributed dynamical systems. *Europhysics Letters (EPL)*, 49(6):708–714, mar 15 2000.

[41] Cha Zhang, Dinei Florencio, and Philip A. Chou. Graph signal processing – a probabilistic framework. 2013.