

Polyphonic Intelligence: Constraint-Based Emergence, Pluralistic Inference, and Non-Dominating Integration

Alexander D. Shaw

Computational Psychiatry & Neuropharmacological Systems (CPNS) Lab

Department of Psychology, Faculty of Health & Life Sciences

University of Exeter, UK

<https://cpnslab.com>

Abstract

Across neuroscience, artificial intelligence, physics, and mathematics, dominant theories of intelligence tend to privilege convergence. Intelligence is frequently framed in terms of a single objective function, a unique solution, or a centralised controller that resolves uncertainty through optimisation. Here we propose an alternative framing, termed *polyphonic intelligence*, in which coherent behaviour and meaning emerge from the alignment of multiple weakly coupled processes operating under shared constraints. Rather than eliminating ambiguity through dominance, polyphonic systems sustain plural explanations, allowing structure to arise through soft coordination, constraint satisfaction, and transient synchronisation. We argue that this perspective unifies principles from predictive processing, dynamical systems theory, statistical physics, control theory, and interpretable artificial intelligence, and reframes intelligence as coordination without command.

1 Introduction

Across disciplines concerned with intelligence such as neuroscience, artificial intelligence, cognitive science, control theory, and statistical physics, there exists a striking convergence of assumptions, even when surface-level formalisms differ. Intelligence is repeatedly characterised as the ability to reduce uncertainty, minimise error, optimise performance, or converge upon a solution. Whether framed as Bayesian inference, optimal control, reinforcement learning, or efficient coding, the underlying logic is one of progressive collapse: uncertainty is something to be removed, multiplicity is something to be resolved, and ambiguity is something to be eliminated [1, 2, 3, 4].

This framing has proven extraordinarily powerful; it underlies much of modern machine learning, signal processing, and computational neuroscience [2, 1, 4]. Yet it also embeds a normative claim that is rarely made explicit: that intelligence is fundamentally convergent. But does this claim deserve scrutiny?

Biological intelligence routinely violates this assumption, for example neural systems maintain redundant and overlapping representations, competing perceptual interpretations coexist across cortical hierarchies and action selection unfolds in the presence of unresolved alternatives [5, 6, 7]. Learning does not instantaneously collapse uncertainty, but reshapes it gradually (and often incompletely) [8, 3]. At every scale, cognition appears to tolerate, or even exploit, multiplicity [9, 7].

These observations suggest that convergence may be a special case of intelligent behaviour, rather than its defining feature. In this paper, we argue that intelligence is more accurately understood as a process of coordination among multiple interacting processes than as the optimisation of a single objective [10, 9]. We refer to this framing as *polyphonic intelligence*.

The term “polyphonic” is chosen deliberately to emphasise the coexistence of multiple semi-independent processes whose interactions, rather than their suppression, give rise to coherent structure. Importantly, this is not merely a metaphor. As we show, the same mathematical and dynamical principles recur across inference, control, neural systems, and physical models of collective behaviour [11, 7, 12].

1.1 Dominance vs. Alignment

A defining assumption of dominant theories of intelligence is that uncertainty should be resolved as efficiently as possible. This assumption is typically formalised as optimisation over a scalar objective:

$$\theta = \arg \min_{\theta} \mathcal{L}(\theta) (1)$$

where \mathcal{L} may represent prediction error, negative reward, control cost, surprise, or free energy [13, 3, 4].

Although these quantities differ semantically, they share a common mathematical structure in which intelligence is equated with movement toward the extreme of a global function. This structure appears across reinforcement learning as reward maximisation, in control theory as trajectory optimisation, in variational inference as posterior concentration, and in neuroscience as error minimisation or efficient coding [2, 1, 3, 4]. Across these domains, the convergence-based framing tends to privilege a single objective and a single solution, with alternative representations treated as temporary errors rather than as meaningful possibilities. While this approach is computationally convenient and often analytically tractable, it embeds a strong assumption about how intelligent systems are expected to behave [4, 8].

That assumption sits awkwardly with empirical observations of biological cognition because real-world environments are non-stationary, partially observable, and time-constrained, and objectives shift continually with context. The associated landscapes are typically high-dimensional and non-convex, meaning that strong pressure toward rapid convergence might be maladaptive. Under these conditions, early commitment often increases brittleness, reduces flexibility, and amplifies the cost of error [10, 7, 9].

Polyphonic intelligence replaces dominance with *alignment* as the organising principle. That is, rather than forcing agreement, multiple processes evolve in parallel, each governed by its own dynamics, state variables, and partial objectives. Coherence arises when these processes become sufficiently compatible to support coordinated behaviour [9, 7, 12].

Alignment differs from convergence in several important ways: it tends to be graded rather than all-or-nothing, local rather than global, and reversible rather than final. From a dynamical perspective, alignment is closer to synchronisation in coupled systems, where coherence can emerge temporarily without suppressing individual dynamics [11, 9]. From a geometric perspective, it corresponds to trajectories becoming mutually consistent within a shared constraint manifold, rather than collapsing onto a single solution [14].

Once dominance is abandoned, intelligence can no longer be defined by the elimination of alternatives. Instead, it must be understood as the management of relations among coexisting processes [10, 15, 5].

1.2 Pluralistic Inference

The move from dominance to alignment necessitates a corresponding shift in how inference is understood; if intelligent systems are not required to converge upon a single explanation, then inference cannot be equated with posterior collapse [1, 2].

In classical Bayesian formulations, inference is often approximated as unimodal,

$$q(z) \approx \mathcal{N}(\mu, \Sigma), \quad (2)$$

an assumption that enforces representational singularity. While computationally convenient, this approximation encodes a strong bias toward explanatory dominance [2, 1].

Polyphonic intelligence instead permits pluralistic inference,

$$q(z) = \sum_{k=1}^K \pi_k q_k(z), \quad \sum_k \pi_k = 1, \quad (3)$$

a form that preserves internal structure within the posterior rather than collapsing it prematurely [1, 16].

Each component q_k represents a locally coherent hypothesis that remains viable unless rendered inconsistent by evidence or constraints. Importantly, these hypotheses are not merely samples from a distribution, but structured explanatory modes, often associated with distinct dynamical regimes, interpretations, or action policies [2, 4].

This structure occurs across several disciplines, for example, in statistical physics systems near criticality occupy ensembles of metastable states rather than unique equilibria [9, 11]. In ensemble filtering and particle-based methods, multiple trajectories are propagated in parallel to hedge against uncertainty [16]. In Bayesian model comparison, posterior mass is distributed across competing generative models rather than collapsed onto a single explanation [1].

Neural systems exhibit analogous behaviours, for example, parallel cortico-thalamic loops encode competing interpretations of sensory input, which are maintained until alignment pressures arising from sensory evidence, action consequences, or contextual priors favour one interpretation over others [5, 7]. Even then, alternatives are often suppressed rather than erased, remaining latent and recoverable [6].

Pluralistic inference is therefore not *indecision*, but a structurally rational response to uncertainty in complex environments. It allows systems to remain adaptable, robust, and context-sensitive [10]. Crucially, pluralism is not unconstrained; hypotheses compete through alignment pressures rather than dominance, and those that fail to remain compatible with the broader system lose influence gradually rather than catastrophically [15].

In this sense, pluralistic inference is the inferential counterpart of alignment in that it formalises how multiple explanatory processes can coexist without requiring a final arbiter [1, 9].

1.3 Non-Dominating Integration

If inference is pluralistic, then integration cannot be fully hierarchical or centrally controlled. Instead, polyphonic intelligence requires a form of integration in which multiple subsystems interact without any single component imposing a global solution. In this setting, coherence arises through regulated interaction rather than command [15, 5].

One way to formalise this is through a non-dominating integration scheme,

$$\mathcal{F} = \sum_i \mathcal{F}_i + \sum_{i \neq j} \lambda_{ij} \mathcal{C}(z_i, z_j), \quad (4)$$

where each \mathcal{F}_i represents a local objective or free-energy-like quantity, \mathcal{C} encodes consistency relations between subsystems, and the coupling coefficients λ_{ij} modulate the strength of these interactions. Importantly, these coupling terms do not enforce agreement. Instead, they shape the degree to which subsystems influence one another, allowing coordination to emerge gradually [3, 15].

This formal structure also recurs across a range of disciplines: in decentralised control, it appears in consensus and coordination algorithms that rely on local interactions rather than central planners [12, 17]. In statistical physics, similar forms arise in models of weakly interacting systems, where global behaviour reflects collective dynamics rather than imposed order [11]. In neural systems, the same principle maps naturally onto precision-weighted prediction errors and modulatory gain, which regulate the influence of different signals without enforcing uniformity [13, 7].

The behaviour of such systems depends, critically, on the strength of coupling. When interactions are too weak, subsystems drift apart and coherence breaks down. When interactions are too strong, diversity is suppressed and the system collapses toward a single mode. Polyphonic intelligence operates between these extremes, in a regime where coordination is sufficient to support coherent behaviour but not so strong as to eliminate plurality [9, 15].

1.4 Dynamics Under Constraints

Taken together, the preceding sections point toward a deeper organising principle - in polyphonic intelligence, behaviour is not organised around the pursuit of explicit goals in the classical sense. Instead, it is shaped by constraints that define which states and trajectories are viable [15, 14].

Rather than specifying target trajectories, the system evolves within a constrained state space,

$$\mathcal{M} = \{x \mid g_k(x) \leq 0 \ \forall k\}, \quad (5)$$

where the constraints delimit a set of permissible behaviours rather than prescribing a particular outcome. The task of the system is therefore not to optimise within \mathcal{M} , but to remain within it while responding to ongoing perturbations and changes in context [14, 10].

This notion of viability appears across cybernetics, ecology, and dynamical systems theory, and stands in contrast to optimisation-based control frameworks [15, 10]. Within a feasible manifold, behaviour can take many forms. Synchronisation and desynchronisation, transient alignment, and reconfiguration of interacting processes can all emerge without central orchestration [9, 11, 7]. From this perspective, intelligence consists in maintaining coherence within an evolving set of constraints, rather than in achieving a fixed endpoint [14].

1.5 Coordination in Distributed Systems

The metaphor of musical counterpoint provides a useful way of characterising the form of coordination described throughout this paper, but it is best understood as a structural analogy rather than a decorative one. In counterpoint, independent voices retain their own identity and dynamics while operating under shared constraints that limit how they may interact. Coherence arises not because one voice dominates the others, but because their relationships remain mutually compatible over time [9, 11].

This structure closely mirrors the formal framework developed above, where individual voices correspond to subsystems or hypotheses that evolve according to their own local dynamics. Harmony might correspond to alignment under shared constraints, while dissonance reflects misalignment that prompts reconfiguration (rather than outright suppression). Importantly, these relationships are dynamic, such that alignment can strengthen, weaken, or dissolve as conditions change [7, 10].

Seen in this way, the counterpoint metaphor captures a concrete mechanism for distributed coordination. Similar principles appear in neural circuits, where parallel pathways interact through modulatory signals [5, 13], in multi-agent systems that rely on local interactions rather than global planners [12, 17], and in complex adaptive networks whose collective behaviour emerges from relational structure [15, 11]. In all cases, coherence is achieved through regulated interaction rather than command.

2 Implications

Framing intelligence as a polyphonic process has several implications that cut across existing theoretical divides. First, it reframes the long-standing tension between interpretability and performance: in systems organised around alignment rather than dominance, structure is not an overhead imposed for the sake of explanation. Instead, structured interaction between subsystems is what enables coherent behaviour in the first place. Interpretability and performance are therefore not competing objectives, but mutually reinforcing properties of systems that rely on coordinated, constraint-governed dynamics [10, 5, 4].

Second, this perspective alters how familiar trade-offs such as exploration and exploitation should be understood. Rather than being treated as discrete modes that must be switched between, these behaviours emerge naturally from the management of multiple hypotheses under uncertainty. Maintaining plural explanatory modes supports exploration, while transient alignment among a subset of these modes supports exploitation. The balance between the two is regulated continuously through alignment pressures rather than by explicit control policies [1, 16, 9].

A similar reframing applies to the relationship between stability and flexibility; in optimisation-based systems, these properties are often in tension, requiring careful tuning to avoid rigidity or instability. In constraint-driven systems, however, stability arises from remaining within viable regions of state space, while flexibility is afforded by the multiplicity of admissible trajectories within those regions. Constraint-based dynamics therefore support adaptive behaviour without requiring explicit mechanisms to toggle between stable and flexible modes [15, 14, 10].

Looking forward, the framework of polyphonic intelligence suggests several directions for future exploration. In neuroscience, it motivates models that prioritise distributed coordination, modulatory coupling, and plural representations over hierarchical control [13, 7, 5]. In artificial intelligence, it points towards architectures built from interacting generative modules, structured uncertainty, and soft consensus rather than simple error-based objectives [4, 12, 17].

3 Conclusion

Throughout this paper, we have argued that many prevailing theories of intelligence are organised around a commitment to convergence - whether expressed as optimisation, error minimisation, or reward maximisation, intelligence is often equated with the capacity to resolve uncertainty and settle on a solution [2, 3, 4]. While this framing has clear practical advantages, it also imposes a narrow view of how intelligent systems must behave.

Polyphonic intelligence offers an alternative; rather than privileging dominance by a single objective, representation, or controller, it emphasises alignment among multiple weakly coupled processes operating under shared constraints. Coherence arises not through suppression of alternatives, but through regulated interaction among them [15, 9, 7]. Inference remains pluralistic, integration is non-dominating, and behaviour unfolds within constrained but flexible state spaces.

From this perspective, intelligent systems do not primarily ask which explanation is correct in an absolute sense. They ask which explanations can coexist without contradiction, given the current constraints and context [1, 14]. Alignment is therefore provisional, dynamical, local, and reversible, allowing systems to adapt without collapsing prematurely onto brittle solutions.

Seen in this light, intelligence is best understood as a property of *process* rather than outcome. It consists not in reaching a final answer, but in sustaining coherence under uncertainty. By shifting the focus from convergence to coordination, the framework of polyphonic intelligence provides a unifying lens for understanding intelligent behaviour across biological, artificial, and physical systems [10, 11, 5].

References

- [1] David J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- [2] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [3] Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- [4] Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- [5] Olaf Sporns. *Networks of the Brain*. MIT Press, 2011.
- [6] Giulio Tononi, Olaf Sporns, and Gerald M. Edelman. A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences*, 91(11):5033–5037, 1994.
- [7] Michael Breakspear. Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20:340–352, 2017.
- [8] Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3):181–204, 2013.
- [9] J. A. Scott Kelso. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. MIT Press, 1995.
- [10] Randall D. Beer. Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3):91–99, 2000.
- [11] Steven H. Strogatz. *Sync: The Emerging Science of Spontaneous Order*. Hyperion, 2003.
- [12] Reza Olfati-Saber, J. Alex Fax, and Richard M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [13] Karl Friston, James Kilner, and Lee Harrison. A free energy principle for the brain. *Journal of Physiology-Paris*, 100(1–3):70–87, 2006.
- [14] Jean-Pierre Aubin. *Viability Theory*. Birkhäuser, 1991.
- [15] W. Ross Ashby. *An Introduction to Cybernetics*. Chapman & Hall, 1956.

- [16] Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010.
- [17] Mehran Mesbahi and Magnus Egerstedt. *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.