Distinguishability All the Way Down: Why Multiscale Intelligence Needs a Finite-Dimensional Constraint Algebra

David Neale

Goleudy.ai, Rochester, New York, USA

ORCID: 0009-0006-7749-9826

April 2026 (revised)

Abstract

Multiple independent research programmes are converging on a shared insight: that the structure of distinguishability — not spacetime, not matter, not consciousness — is the primitive from which physics, agency, and intelligence can be derived. Friston's Free Energy Principle grounds self-organisation in the minimisation of surprise by statistically distinguishable systems. Hoffman and Prakash build observer-participants from measurable spaces of distinguishable experiences governed by Markov chains. Fields and Levin formalise collective intelligence as efficient search across multi-scale problem spaces. Each programme works in infinite-dimensional functional spaces — spaces of probability distributions, Markov kernels, or generative models — without specifying the finite-dimensional geometric structure that constrains how distinguishability actually operates. This perspective paper describes a complementary programme that derives, from a single modal logic axiom ($\Diamond N \to \neg N$: "if nothingness is possible, then nothingness does not obtain"), a specific finite-dimensional constraint algebra — five independent constraint dimensions encoded in a Clifford algebra — with empirically measurable signatures across cellular automata, Game of Life patterns, and reaction-diffusion systems. The constraint algebra provides what the converging programmes currently lack: a finite set of named dimensions whose independence is empirically verified, testable predictions about system behaviour under perturbation, and a geometric account of why persistent systems cluster near the vertices of the outcome simplex when perturbed. Crucially, the axiom itself predicts that specific measurement values will be observer-dependent while the dimensionality is invariant — a prediction confirmed by the empirical programme and formalised by recent work on structural information for computationally bounded observers (Finzi et al., 2026). The paper situates this contribution within the landscape of current work, identifies specific points of formal contact, and proposes collaborative directions.

---

1. The Convergence

Something unusual is happening across several fields simultaneously. Research programmes in theoretical neuroscience, consciousness science, developmental biology, and foundations of physics are independently arriving at the same structural claim: that the capacity to be *distinguished* — from an environment, from other systems, from nothingness — is not a derived property of physical systems but a foundational one from which other structure follows.

Friston (2010, 2019) begins with the observation that for something to exist as a *thing*, it must possess internal states that are statistically separable from external states via a Markov blanket. This statistical distinguishability, combined with the requirement that the system persist in a non-equilibrium steady state, yields the free energy principle: the dynamics of any such system can be described as minimising variational free energy, or equivalently, as minimising the surprise of its sensory encounters given its generative model. The recursive composition of Markov blankets — blankets within blankets at successively higher spatiotemporal scales — provides a multiscale architecture, and active inference describes how systems simultaneously update their models (perception) and change their environments (action) to maintain themselves.

Hoffman and Prakash (2014, 2024) take a different starting point — consciousness rather than statistics — but arrive at a structurally parallel architecture. Their conscious agents are minimal entities defined by measurable spaces of possible experiences and Markov kernels governing transitions between them. The recent "recursive trace logic" (Hoffman and Prakash 2024, in preparation) introduces a trace order on Markov chains that defines observation and agency within a non-Boolean logic. Like Friston's recursive Markov blankets, Hoffman's conscious agents compose: any collection of agents is itself an agent, and the dynamics at each scale have the same formal structure.

Fields and Levin (2022, 2025) approach multiscale intelligence empirically, through developmental biology and synthetic organisms. Levin's laboratory has demonstrated that cells, tissues, and organisms exhibit goal-directed problem-solving across scales — what they term the "multiscale competency architecture." Chis-Ciure and Levin (2025) formalise this as search efficiency: intelligence is the logarithm of the ratio between the cost of a random walk and the cost of the biological agent's actual trajectory through a problem space. Fields and Levin (2025) connect this to the free energy principle and the Conway-Kochen free will theorem, arguing that the same formal structures describe agency at molecular, cellular, and organismal scales.

These programmes share a deep structural commitment: distinguishability is primitive; structure is derived; the derivation is recursive across scales. They disagree about what is fundamental (statistics, consciousness, or physics) but agree about the *form* of the answer.

And they share a common gap.

2. The Gap

Each of these programmes works in infinite-dimensional spaces. Friston's generative models are probability distributions over potentially unbounded state spaces. Hoffman's conscious agents are defined by measurable spaces with arbitrary cardinality. Fields and Levin's problem spaces have dimensionality determined by the specific biological system under study. The free energy principle tells you *that* a system minimises surprise, but not *how many independent ways* there are to be surprised. The conscious agent framework tells you *that* agents compose, but not *what specific structure* governs the composition. The search efficiency metric tells you *how much better* an organism is than a random walk, but not *what dimensions* define the space being searched.

This is not a criticism — infinite-dimensional generality is a feature when the goal is universality. But it leaves three questions unanswered:

1. **How many independent constraints does distinguishability impose?** Is the constraint space two-dimensional? Ten-dimensional? Infinite?
2. **What is the relationship between constraints?** Do they operate independently, or do they interact? If they interact, is the interaction a universal property of the constraints or a substrate-specific feature of the measurement?
3. **Why do persistent systems exhibit a small number of qualitatively distinct strategies?** Levin observes a "multiscale competency architecture." Wolfram identifies four classes of cellular automaton behaviour. Pattern classification schemes across fields typically converge on three to four categories. Is this a coincidence of observation, or does it follow from the geometry of constraint space?

3. The Constraint Algebra

The programme described here — developed under the name "Being from Nothingness" (BFN) — addresses these questions by taking the shared starting point (distinguishability as primitive) and pushing it through a specific algebraic derivation.

3.1 The Axiom

The formal starting point is a single proposition in modal logic:

$\Diamond \mathbf{N} \rightarrow \neg \mathbf{N}$

Read: "if nothingness is possible, then nothingness does not obtain." Equivalently: to exist is to be distinguishable from nothingness. This is not an empirical claim but a logical constraint on what it means for anything to be the case. It shares its logical structure with Friston's observation that "for something to exist it must possess states that can be separated statistically from external states" (Friston 2019, §1), but it is stated as a modal principle rather than a statistical one.

3.2 Five Constraints from Categorical Exhaustion

The axiom, when iterated across increasing numbers of mutually distinguishable elements, generates constraints by categorical exhaustion — each new element that must be distinguished from all existing elements introduces a new *type* of relational requirement that cannot be reduced to previous types.

N (distinguishable elements)	New constraint type	Label
1	Distinction from environment	β (boundary)
2	Distinction between elements	κ (pattern)
3	Orientation / chirality	ρ (resource/complexity)
4	Integration / non-decomposability	λ (integration)
5	Ordering / directed change	τ (ordering)

At N = 5 the algebra closes: no new independent constraint type emerges at N = 6 or beyond. This closure is the algebraic content of the claim that five constraints are both necessary and sufficient.

The categorical exhaustion argument is theoretical. The *independence* of the five constraints is an empirical question, addressed in Section 3.4.

3.3 Clifford Algebra Structure

The five constraints are encoded as generators of a Clifford algebra. The theoretical framework proposes Cl(3,2) with mixed signature — β and τ as the timelike (negative-signature) generators and κ, ρ, λ as spacelike — based on the expectation that boundary maintenance and change generation are in tension. However, the specific signature is an empirical question tied to the sign structure of the constraint coupling matrix, which proves to be observer-protocol-dependent (see Section 3.5).

Regardless of the specific signature, the Clifford algebra provides:

- **Grade structure** corresponding to N-thresholds: grade-1 (individual constraints), grade-2 (pairwise couplings), grade-3 (three-body interactions including τ_circ as a trivector), grade-4 (λ-mediated collective patterns), grade-5 (the pseudoscalar $I_5 = \beta \wedge \kappa \wedge \rho \wedge \lambda \wedge \tau$).
- **A geometric account of system behaviour**: trajectories in five-dimensional constraint space, rather than in configuration space, describe how systems maintain themselves.
- **A closure property**: the algebra closes at grade 5, setting an upper bound on the independent dimensions of collective structure. There is no grade 6 in a five-dimensional Clifford algebra, corresponding to the prediction that no sixth independent constraint type exists.

3.4 Empirical Validation: Five Independent Dimensions

The theoretical prediction of five independent constraint dimensions has been validated empirically across two substrates with maximally different architectures.

Gray-Scott reaction-diffusion (222 regimes, two-dimensional continuous dynamics): All five constraints are statistically independent, with maximum pairwise correlation $|r| = 0.206$ and all variance inflation factors between 1.03 and 1.08. Principal component analysis yields a nearly uniform variance distribution (27%, 23%, 19%, 16%, 14% for PC1–PC5). The fifth component is not a residual — it carries 14.2% of total variance.

Elementary cellular automata (256 rules, one-dimensional discrete dynamics): After correcting for an observer-protocol artefact (Section 3.5), all five constraints are statistically independent, with maximum VIF = 1.40 and PC5 = 10.2% of variance.

In both substrates, the integration constraint (λ) is notably independent: VIF = 1.05 (GS) and 1.12 (CA), with maximum correlation to any other constraint below 0.26. This consistent independence of the most global measurement — spatial coherence across the system — across substrates with different dimensionalities, state types, and dynamical mechanisms supports the claim that the five dimensions capture genuinely distinct aspects of structural information.

3.5 Observer-Dependent Measurement as a Prediction

The axiom states that distinguishability is relational — it depends on both the system and the observer. This predicts that specific constraint values will depend on the observation protocol: change the observer's measurement procedure and the apparent geometry of the constraint space changes, even for the same underlying systems.

This prediction is confirmed by the empirical programme, most strikingly in a controlled experiment on the ordering constraint τ. The measurement of τ involves computing a trivector from edge vectors in constraint space, where each edge-vector component is the change in one constraint between consecutive timesteps. In one-dimensional CA, one constraint (κ) fluctuates much more than the others between timesteps, producing edge vectors dominated by the κ-component. The resulting τ_circ tracks κ by proxy, creating an apparent correlation of $r = 0.778$ and reducing the effective dimensionality to four.

Normalising the edge-vector components by their temporal variances — so that each constraint axis contributes equally — eliminates the coupling entirely ($r = -0.011$) and restores the five-dimensional independence. The same systems, the same five measurements, a single methodological choice, and the apparent geometry of constraint space changes from four-dimensional to five-dimensional.

This is not a measurement failure. It is the axiom in action. Distinguishability is relational; measurement is a relation between system and observer; therefore measurement outcomes are observer-dependent. What is *invariant* across all protocol variations tested is the dimensionality: five independent axes, however they are measured.

Recent work by Finzi et al. (2026) provides the formal machinery for this observation. Their "epiplexity" framework demonstrates that the structural information extractable from a dynamical system depends on the computational resources of the observer — the same system contains different amounts of learnable structure for observers with different time budgets or measurement protocols. Our result is a specific instance: the same systems present different apparent constraint geometries to observers with different measurement protocols, but the dimensionality of the underlying structure is invariant.

The earlier version of this paper reported specific coupling values (β–τ = −0.58, ρ–λ = +0.83) as established facts. These values were obtained from CA measurements with a particular observation protocol and are not reproduced in the Gray-Scott data, where all pairwise correlations are below 0.21. The coupling structure is protocol-dependent; the independence structure is not. The question of whether Cl(3,2) with its mixed signature captures the *universal* constraint geometry, or only the geometry as seen by specific observation protocols on specific substrates, remains open. What is established is the dimensionality: five.

3.6 Inevitable Structural Information Accumulation

The deductive chain from the axiom to observable consequences can now be completed using the epiplexity framework:

1. **Persistence requires distinguishability** (the axiom).
2. **Distinguishability requires iteration** — a system must keep producing distinguishable states.
3. **Iteration is computation** — any system evaluating a successor function is computing.
4. **Computation creates structural information for bounded observers** (Finzi et al., 2026).
5. **Accumulated structural information has five independent dimensions** (empirical validation, Section 3.4).

This chain is deductive: each link follows from the previous. The consequence is that any persistent system *necessarily* accumulates structural information along five independent axes. A bounded observer watching this accumulation *necessarily* infers memory (persistent patterns), learning (refined patterns), and adaptive behaviour (context-dependent patterns). These are not metaphors applied to simple systems — they are the correct inferences for any observer below the computational threshold needed to simulate the system exactly.

The MLD paper ("Memory Without Storage, Learning Without a Learner"), a companion to this work, demonstrates this empirically: 287 systems across four substrates produce reproducible habituation and sensitisation patterns (Spearman $\rho \approx 0.75$) when probed by bounded observers. The patterns are not designed into the systems; they are inevitable consequences of structural information accumulation under bounded observation.

4. Points of Formal Contact

4.1 Friston's Free Energy Principle

The constraint framework stands in a specific formal relationship to the FEP. The variational free energy F = Complexity − Accuracy can be mapped to constraint-space quantities: Complexity corresponds to deviation from the system's characteristic constraint profile, and Accuracy corresponds to how well the current constraint configuration explains the system's sensory encounters.

What the constraint algebra adds: FEP tells you that systems minimise surprise. The constraint algebra tells you that surprise has *five independent dimensions*, and that the observer-dependent coupling between those dimensions determines which persistence strategies are available to which observers. The Markov blanket — Friston's fundamental construct — corresponds to the β constraint. The epiplexity framework (Finzi et al., 2026) formalises what the FEP leaves implicit: that the *amount* of structural information a system accumulates depends on the observer's computational budget, not just on the system's dynamics.

4.2 Hoffman and Prakash's Conscious Agents

Hoffman and Prakash's trace logic and the constraint algebra share a formal ancestor in the notion that composition of observers generates new structure at each level. The N-threshold hierarchy (new constraint types at $N = 1$ through 5) is structurally parallel to the recursive composition of conscious agents, where new collective properties emerge at each level of Markov chain composition. The constraint algebra provides a specific prediction: the recursive composition *closes* at five independent dimensions. Whether the trace logic's open-ended recursive structure is consistent with this closure, or whether it generates independent structure beyond $N = 5$, is an open question that could be formally investigated.

The key philosophical difference: Hoffman takes consciousness as fundamental and derives physics as a user interface. The constraint framework takes distinguishability as fundamental and is agnostic about whether the resulting structure is experienced. The MLD results show that observer-constructed categories like "memory" and "learning" are inevitable inferences by bounded observers watching persistent systems — they do not require that the systems be conscious, nor do they rule it out. The epiplexity framework adds precision: what an observer can learn from a system depends on the observer's computational resources, not on whether the system is "aware" in any deeper sense.

4.3 Fields and Levin's Search Efficiency

Chis-Ciure and Levin's (2025) search efficiency metric — the log ratio of random-walk cost to biological-agent cost — can be decomposed in constraint space. A system's search efficiency in a given problem space depends on which structural information dimensions dominate its profile. Systems dominated by boundary stability are efficient in stable environments but fragile to novelty; systems dominated by temporal ordering are efficient in rapidly changing environments but cannot consolidate gains. Systems with balanced profiles sacrifice peak efficiency for robustness across problem types.

Fields and Levin's (2025) connection between the free energy principle and the Conway-Kochen free will theorem touches directly on the constraint framework's Bridge 4 programme, which investigates the relationship between constraint coupling at $N \geqslant 3$ and non-classical correlations. The Conway-Kochen theorem establishes that if experimenters have free choice of measurement settings, then particles' responses cannot be determined by prior information. In constraint-algebra terms, this connects to properties of the coupling structure at $N \geqslant 3$ — a formal relationship that is under active investigation.

4.4 Levin's Multiscale Competency Architecture

Levin's empirical observation that biological systems exhibit competency at every scale — from molecular networks to cells to tissues to organisms to swarms — maps onto the N-threshold hierarchy. The constraint framework predicts that each scale of organisation generates new collective constraint types at specific thresholds ($N = 3$ for chirality, $N = 4$ for non-decomposable integration, $N = 5$ for full independence). Levin's neural cellular automata work (Hartl, Levin, and Pio-Lopez, 2025) uses the same computational substrate (cellular automata) that provides empirical validation for the constraint algebra. The question of whether the constraint signatures measured in elementary CA correspond to the competency signatures measured in Levin's biological and synthetic systems is a concrete, testable point of contact.

5. What Is at Stake

The convergence described here is not merely a philosophical observation. It identifies a specific empirical programme that could resolve a standing question in the multiscale intelligence literature: **is the structure of distinguishability finite-dimensional?**

If the answer is yes — if five independent constraint dimensions capture the essential geometry of how systems maintain themselves — then several consequences follow:

- **Vertex clustering is geometrically inevitable.** Systems perturbed from their attractors concentrate near the vertices of the outcome simplex rather than distributing uniformly (Neale, companion paper). This is a geometric consequence of dimensional independence: balanced investment across five independent axes is rare, so most systems are dominated by one or two structural information dimensions, producing specialised perturbation responses. The "three or four categories" observed in classification schemes across biology, physics, and complexity science are not imposed by observers — they are carved by the geometry of the information space itself.
- **Structural information accumulation is inevitable.** Any system that persists necessarily computes, and computation creates structural information for bounded observers (Finzi et

al., 2026). What bounded observers call "intelligence" — memory, learning, adaptation — is the inevitable appearance of accumulated structure from the outside. The question of whether a system is "really" intelligent dissolves: for any observer below the simulation threshold, the intelligence is as real as the measurement.

- **Observer-dependent measurement is not a limitation but a prediction.** The axiom predicts that distinguishability is relational. The epiplexity framework formalises this: structural information content depends on the observer's computational budget. The empirical programme confirms it: measurement geometry varies with protocol while dimensionality is invariant. Any programme that takes distinguishability as fundamental should expect — and welcome — observer-dependent measurement.
- **The free energy principle acquires specific internal structure**: surprise is not a scalar but decomposes along five independent axes, and different components of surprise couple differently depending on the substrate and observation protocol.
- **Multiscale collective intelligence has a specific algebraic signature**: the emergence of new constraint types at N-thresholds corresponds to the appearance of new grades in the Clifford algebra, and the closure at grade 5 sets an upper bound on the independent dimensions of collective competency.

If the answer is no — if the constraint space is genuinely infinite-dimensional, or if the specific five-dimensional structure fails to reproduce in new substrates — then the constraint algebra is wrong in a specific, identifiable way, and the convergence between programmes is shallower than it appears.

Either outcome is informative. The point of this perspective is that the question is now well-posed and empirically addressable.

6. Available Resources and Invitation

The empirical programme behind the constraint algebra has produced reproducible computational results across three substrates:

- **MLD paper** ("Memory Without Storage, Learning Without a Learner"): demonstrates that bounded observers necessarily infer memory and learning from persistent systems, with Spearman $\rho \approx 0.75$ correlation between constraint profiles and observer-attributed intelligence. Confirmed reproducible; submission-ready for *Adaptive Behavior*.
- **Persistence Fingerprints paper**: perturbation-response classification across 256 CA rules, 27 GoL patterns, and 14 active Gray-Scott regimes, demonstrating vertex clustering in the outcome simplex and its substrate-specific modulation. First draft complete; targeted for *Physical Review E*.

- **Five Dimensions paper**: PCA, VIF, and correlation analysis across 222 Gray-Scott regimes and 256 CA rules demonstrating five-dimensional independence, including the normalisation experiment that confirms observer-protocol dependence of coupling values. First draft complete; targeted for *Physical Review E*.
- **Bridge papers** connecting the algebraic structure to established physics (three-dimensional geometry, classical-quantum regimes, unitarity, Bell inequalities) are in various stages of development.

All simulation code is self-contained Python (numpy, scipy only) and runs in under 15 minutes on a standard laptop. Code and data are available at goleudy.ai.

This is an explicit invitation to researchers working in the programmes described above — and in related areas including integrated information theory, computational mechanics, and structural realism — to examine the constraint algebra, attempt to reproduce or falsify its empirical predictions, and explore whether the finite-dimensional structure it identifies corresponds to or conflicts with the structures arising in their own frameworks. The converging insight that distinguishability is foundational is, I believe, correct. Whether the specific algebraic structure proposed here is the right formalisation of that insight is an open question that benefits from multiple perspectives.

References

Chis-Ciure, R. and Levin, M. (2025). Cognition all the way down 2.0: neuroscience beyond neurons in the diverse intelligence era. *Synthese* 206(257). doi:10.1007/s11229-025-05319-6.

Fields, C. and Levin, M. (2022). Competency in navigating arbitrary spaces as an invariant for analyzing cognition in diverse embodiments. *Entropy* 24, 819.

Fields, C. and Levin, M. (2025). Life, its origin, and its distribution: a perspective from the Conway-Kochen Theorem and the Free Energy Principle. *Communicative & Integrative Biology* 18, 2466017.

Fields, C. and Levin, M. (2025). Thoughts and thinkers: On the complementarity between objects and processes. *Physics of Life Reviews* 52, 256–273.

Fields, C., Hoffman, D.D., Prakash, C. and Singh, M. (2018). Conscious agent networks: Formal analysis and application to cognition. *Cognitive Systems Research* 47, 186–213.

Finzi, M., Qiu, S., Jiang, Y., Izmailov, P., Kolter, J.Z. and Wilson, A.G. (2026). From Entropy to Epiplexity: Rethinking Information for Computationally Bounded Intelligence. *arXiv:2601.03220v2*.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience* 11(2), 127–138.

Friston, K. (2019). A free energy principle for a particular physics. arXiv:1906.10184.

Hartl, B., Levin, M. and Pio-Lopez, L. (2025). Neural cellular automata: Applications to biology and beyond classical AI. *Physics of Life Reviews* 56.

Hoffman, D.D. and Prakash, C. (2014). Objects of consciousness. *Frontiers in Psychology* 5, 577.

Hoffman, D.D., Prakash, C. and Chattopadhyay, S. (2024). Traces of consciousness. *Preprints.org* 202410.1305.

Hoffman, D.D., Prakash, C. and Prentner, R. (2023). Fusions of consciousness. *Entropy* 25(1), 129.

Levin, M. (2024). Collective intelligence: A unifying concept for integrating biology across scales and substrates. *Communications Biology* 7, 378.

Levin, M. (2025). The multiscale wisdom of the body: Collective intelligence as a tractable interface for next-generation biomedicine. *BioEssays* 47(3), e202400196.

Neale, D. (2025). Being from Nothingness: Deriving the Structure of Existence from a Single Axiom. Preprint, goleudy.ai.

Zhang, T., Goldstein, A. and Levin, M. (2024). Classical sorting algorithms as a model of morphogenesis. *Adaptive Behavior* 32(6). doi:10.1177/10597123241269740.

Goleudy.ai research programme. Correspondence: goleudy.ai. ORCID: 0009-0006-7749-9826.