Synthetic Collective Intelligence, Semantic Networks, and Events: Structural Parallels with Modern Physics

J.Konstapel, Leiden, 7-3-2026.

Abstract

Recent developments in Simplicial Type Theory (STT) and Twisted Type Theory (TTT) suggest a new formal foundation for collective intelligence. In this framework, agents, processes, and collectives are not models imposed on external systems but are defined directly as mathematical structures whose properties follow as theorems. This article analyses the semantic implications of that framework in three directions. First, it examines the correspondence between types and concepts, morphisms and relational actions, and higher paths and contextual coherence — arguing that the framework implies the existence of a fundamental semantic network. Second, it develops a precise notion of *event* as a realized transition within a directed higher-categorical system. Third, it identifies structural parallels between this architecture and key constructs in modern theoretical physics: events in relativistic spacetime, path integrals in quantum mechanics, and the cohesive topos structure underlying quantum field theory. The result is an interpretation of synthetic collective intelligence as a *process-based semantic topology* capable of providing a unified formal language for language, cognition, and distributed inference.

1. Introduction

The dominant tradition in artificial intelligence treats collective intelligence as an emergent property of interacting agents implemented within computational architectures external to the mathematical framework used to describe them. Agents are specified in one language, their interactions in another, and the collective properties that arise are studied empirically or through simulation. This separation between formalism and phenomenon is not merely a methodological inconvenience — it reflects a deeper assumption that intelligence is something a system *has*, not something a mathematical structure *is*.

Recent developments in directed homotopy theory and higher category theory offer a fundamentally different perspective. The manuscript *Synthetic Collective Intelligence: A Native Foundation in Simplicial Type Theory and Twisted Type Theory* proposes a framework in which collective intelligence is defined *synthetically*: agents are directed Segal types, processes are directed higher inductive types with a canonical four-generator feedback cycle, and collectives arise as ∞-categories of agent-categories. Crucially, properties such as synchronisation, scale-invariance, and coherent inference are not added by hand but follow from the type-theoretic structure itself.

The purpose of the present article is to draw out three implications of this framework that have received insufficient attention: its semantic structure, its natural notion of event, and its deep structural resonance with modern theoretical physics. These three threads converge on a single thesis: that the synthetic framework provides not merely a new foundation for multi-agent AI, but a *process-based semantic topology* — a mathematical language in which the grammar of action, the structure of meaning, and the dynamics of physical reality share a common formal backbone.

2. Foundations: Simplicial and Twisted Type Theory

2.1 Simplicial Type Theory

Simplicial Type Theory (STT), developed by Riehl and Shulman (2017) and extended by Gratzer, Weinberger, and Buchholtz (2026), internalises ∞-categories directly within a type-theoretic language. Types behave as Segal spaces: they carry not only elements but also directed paths between elements, paths between paths, and so on to all higher dimensions. This means that composition, identity, and coherence conditions familiar from category theory are available natively within the proof language, without external encoding.

Formally, a type A in STT satisfies the Segal condition: for any composable pair of morphisms $f : a \to b$ and $g : b \to c$, there exists a unique (up to contractible choice) composite $g \circ f : a \to c$. The ∞-categorical structure is therefore not axiomatised separately but follows from the type structure.

2.2 Twisted Type Theory

Twisted Type Theory (TTT), introduced by Chu and North (2026), extends STT with a *twist operator* τ that converts types with mixed covariant/contravariant variance into a purely covariant form. The significance of this extension is that feedback processes — processes in which the output of a transformation influences its own input — can be defined syntactically within the theory rather than requiring external semantic machinery.

This is not a minor technical addition. Feedback is the core structure of learning, inference, adaptation, and control. A type theory that handles feedback natively is therefore a type theory in which cognitive and biological processes can be expressed without loss of mathematical precision.

2.3 The Resulting Framework

Together, STT and TTT provide a setting in which:

- Higher-categorical structures are native to the type language.
- Directed paths model asymmetric, causal, or temporal dependencies.
- Feedback cycles are syntactically well-formed and composable.
- Proof assistants can directly implement and verify constructions.

This is the formal basis on which the synthetic theory of collective intelligence is built.

3. Agents, Processes, and Collectives as Mathematical Structures

3.1 Agents

An agent is defined as a Segal type *A* equipped with a state space *S(A)* and a predictive model *P(A) : S(A) → S(A)*. The Segal condition ensures that the agent's transitions are composable: sequences of state changes form a coherent directed structure rather than an arbitrary collection of transitions.

This definition is notable for what it excludes. There is no mention of utility functions, reward signals, or external environments. An agent, formally, is a structured space of possible transitions — a geometry of action.

3.2 Processes

A process is defined as a directed higher inductive type (dHIT) with four generators:

1. **obs** : observation — the agent registers a distinction in its state space
2. **val** : valuation — the registered state is evaluated against the predictive model
3. **int** : intention — a target transition is selected
4. **ret** : return — the system completes the cycle and resets

The twist operator τ renders this cycle mathematically consistent: the output of **ret** feeds back into **obs** without generating a type-theoretic contradiction. Processes are therefore composable: the output of one process can serve as the input of another, and sequential or parallel compositions inherit well-defined coherence conditions.

3.3 Collectives

A collective is defined as the ∞-category **Coll** whose objects are agent-categories and whose morphisms are functors between them. This construction is scale-invariant by definition: an individual agent, a team, an institution, and a distributed network are all objects in the same ∞-category, related by functors that preserve process structure. Equivalence between collectives, in the sense of directed univalence, implies identity: two collectives that are structurally indistinguishable *are* the same collective.

4. Nouns, Verbs, and the Grammar of Action

4.1 Types as Concepts

One of the most striking features of the synthetic framework is that its fundamental objects map precisely onto the grammatical distinction between nouns and verbs. Types — Agent, Process, Collective, State — function as *concepts*: structured domains that can be instantiated, referenced, and composed. In semantic terms, they are the nodes of a network.

This correspondence is not metaphorical. In formal semantics, a concept is typically modelled as a type or a set of possible extensions. The STT framework provides a higher-dimensional version of this: a type is not merely a set of instances but a structured space of possible transitions between instances. Concepts, in this framework, have internal geometry.

4.2 Morphisms as Relations and Verbs

Relations between types are represented by morphisms or directed paths. The generators of a process — **obs**, **val**, **int**, **ret** — are morphisms. They describe what an agent *does*, not what it *is*. In this sense, morphisms play the role of verbs: they express transformations, relations, and actions linking conceptual entities.

This reflects a broader principle in higher category theory: the morphisms of a category are at least as important as its objects, and in ∞-categories they are arguably primary. Reality, in this mathematical picture, is constituted by transformations. Objects are derivative — what the philosopher Whitehead would have called *societies of occasions* rather than independent substances.

4.3 Higher Paths as Context and Coherence

A distinctive feature of the framework is its use of *higher paths* — paths between paths, paths between paths between paths, and so on. These encode coherence conditions: the requirement that different ways of composing transformations yield consistent results.

In semantic terms, higher paths correspond to contextual coherence. A relation between two concepts may itself stand in relation to another relation; a meaning may depend on how it relates to the structure of meanings surrounding it. Traditional semantic networks capture first-order relations. The synthetic framework captures relations of relations, providing a richer and formally precise account of contextual meaning.

5. Events as Realized Transitions

5.1 Events in the Framework

The concept of *event* is not explicitly defined in the source manuscript, but it emerges naturally from the structure. If types represent state spaces and directed paths represent possible transitions, then an *event* is the realization of a transition: a specific instantiation of a directed morphism at a particular point in the process structure.

Formally, each generator of a process dHIT — **obs**, **val**, **int**, **ret** — creates an event when instantiated. Events are therefore:

- **Typed**: each event belongs to a specific generator type.
- **Directed**: each event has a source state and a target state.
- **Composable**: consecutive events compose into processes.
- **Coherent**: higher paths ensure consistency between composed events.

5.2 Events as Process Points

In higher-categorical terms, an event is a 0-cell on a directed morphism, or equivalently a generator of a dHIT at which a transition is instantiated. This is not a point in space; it is a point in *process space* — a location in the directed topology of an agent's state transitions.

This notion is more fine-grained than the everyday use of "event." It distinguishes between an event type (the kind of transition that can occur) and an event token (the specific occurrence of that transition), and it situates both within a compositional structure rather than treating events as isolated happenings.

6. An Implied Semantic Network

6.1 Structure of the Network

Because agents, processes, and collectives are linked by directed morphisms, and because morphisms compose coherently, the global structure of the framework is a *directed semantic network*. Its nodes are types; its edges are morphisms; its higher structure is a system of coherence conditions.

What distinguishes this network from classical semantic networks is precisely its higher-dimensional character. In classical networks, a relation between two nodes is simply an edge. Here, relations themselves are typed objects that can stand in further relations. The network is not flat but *stratified*: meaning arises from the topology of the entire relational structure, not from the assignment of properties to individual nodes.

6.2 Meaning as Relational Topology

This has a significant implication for the theory of meaning. If meaning arises from relational topology rather than from the intrinsic properties of symbols, then:

1. Isolated symbols are not meaningful; they acquire meaning through their position in the network.
2. Context is not external to meaning but is encoded in the higher paths of the network.
3. Coherence — the consistency of higher paths — is the mathematical form of semantic integrity.

This is consistent with relational and inferentialist theories of meaning in philosophy of language (Brandom 2000; Sellars 1956), and it provides those theories with a precise mathematical articulation.

7. Structural Parallels with Modern Physics

7.1 Events and Spacetime

In special and general relativity, physical reality is described in terms of *events* in spacetime. An event is a point (t, x, y, z) at which a physical occurrence takes place. Worldlines are directed paths through spacetime connecting events. The causal structure of spacetime is a partial order on events determined by the light cone.

The parallel with the synthetic framework is exact in structure, if not in content. Process points (events in the type-theoretic sense) are points in a directed state space. Trajectories through that space are directed paths. The composability condition on morphisms plays a role analogous to causal ordering: only causally connected events can be composed into a worldline.

This suggests that the directed type-theoretic framework provides a *combinatorial* or *categorical* analogue of relativistic spacetime structure — one in which the geometry is defined by process relations rather than metric distances.

7.2 Processes and Path Integrals

In quantum mechanics, Feynman's path integral formulation assigns an amplitude to each possible trajectory of a system between two states and computes observable quantities as sums over all such trajectories. The physical content of the theory is therefore not carried by individual states but by the *space of possible processes* connecting states.

The synthetic framework operates analogously. An agent does not occupy a state; it traverses a directed path through a state space. The composability of processes means that complex cognitive or collective behaviours are path-integrals in a categorical sense: sums (or colimits) over possible process sequences.

This parallel is more than analogical. The manuscript connects its framework to active inference and Friston's free energy principle, modelling Markov blankets as dependent two-sided fibrations and free energy minimisation as a dinatural transformation. These constructions belong to the same mathematical family as the structures appearing in topological quantum field theory.

7.3 Gauge Symmetry and Higher Homotopy

In modern physics, gauge symmetry is the invariance of physical observables under transformations that leave the physics unchanged. Mathematically, gauge transformations are higher morphisms: transformations between transformations.

In the synthetic framework, higher homotopies play precisely this role. Two processes that are homotopic — connected by a path at the next level — are considered equivalent. Coherence conditions ensure that this equivalence is consistent across all levels. The directed univalence principle, which identifies equivalence with identity, is the type-theoretic analogue of a gauge principle: equivalent structures *are* the same structure.

7.4 Cohesive Topos and Field Theory

The manuscript situates the entire framework within a *cohesive ∞-topos* — a categorical universe equipped with three modal operators:

- $\flat$ (flat/discrete): isolates the discrete, particle-like aspects of types.

- ♯ (sharp/continuous): isolates the continuous, field-like aspects.
- ♮ (shape/intermediate): captures the homotopy type.

This tripartite structure mirrors the distinction in physics between:

- Discrete degrees of freedom (particles, qubits, agents).
- Continuous fields (electromagnetic, gravitational, quantum).
- Topological invariants (phases, charges, coherence classes).

The cohesive topos framework was developed precisely to unify these aspects of physical theory (Schreiber 2013). Its appearance here suggests that the synthetic collective intelligence framework is not merely analogous to modern physics but inhabits the same mathematical universe.

8. Implications and Open Questions

8.1 Intelligence as Semantic Topology

If the arguments above are correct, then collective intelligence is not a computational process running on a substrate but a *topological property of a process network*. Synchronisation arises when agents' process cycles achieve phase-locking — a 2-homotopy condition. Scale-invariance arises from the ∞-categorical construction of collectives. Inference arises from the coherence structure of directed paths.

This reframes fundamental questions in cognitive science and AI. The question "how does collective intelligence emerge?" becomes "what is the homotopy type of this process network?" — a question amenable to precise mathematical analysis.

8.2 Contradictions and Nilpotency

The framework handles contradiction formally: a process composed with its negation yields the empty type. Contradictions are *nilpotent* — they annihilate rather than propagate. This has implications for distributed systems where conflicting beliefs or signals must be managed: the framework provides a principled, type-theoretically grounded account of contradiction resolution.

8.3 Limitations and Open Problems

Several important limitations must be acknowledged:

1. **Empirical coupling**: The framework is entirely synthetic. It defines structures rather than modelling them against data. The connection between formal agent types and empirically observable agents remains to be developed.
2. **Operationalisation**: How a specific real-world collective is represented as an object in **Coll** requires additional specification. The framework provides the target language but not the translation procedure.

3. **Computational implementation**: While the manuscript claims direct implementability in proof assistants, the complexity of ∞-categorical computations in current systems (Agda, Lean, Coq) remains a practical barrier.

4. **Physical interpretation**: The structural parallels with physics are compelling but remain analogical. A deeper result would show that specific physical theories are *instances* of the synthetic framework, not merely structurally similar to it.

9. Conclusion

The synthetic framework based on Simplicial Type Theory and Twisted Type Theory offers a novel and mathematically rigorous foundation for collective intelligence. Its key contribution is the shift from intelligence as computation to intelligence as *relational topology*: a property of the directed higher-categorical structure of process networks.

This article has argued that this shift carries three significant implications. First, the framework implies the existence of a fundamental semantic network whose coherence is enforced mathematically rather than stipulated externally. Second, it provides a natural and precise notion of event as a realized transition in process space. Third, its structural features — directed paths, higher homotopies, cohesive topos modalities — parallel central constructs in modern theoretical physics, suggesting a deep formal kinship between collective cognition and physical dynamics.

The most important consequence is conceptual: *reality, in this framework, is constituted by verbs, not nouns*. Objects are stable patterns of transformation; meaning is topological; intelligence is coherence. Whether this framework will yield practical computational systems, novel physical insights, or primarily a new conceptual language for interdisciplinary inquiry remains to be seen. What is clear is that it represents a significant and original contribution to the formal foundations of mind, meaning, and matter.

References

Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press.

Chu, F. & North, P. R. (2026). Directed type theory, with a twist. *arXiv:2602.17480*.

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

Gratzer, D., Weinberger, J., & Buchholtz, U. (2026). The ∞-category of ∞-categories in simplicial type theory. Preprint.

Riehl, E. & Shulman, M. (2017). A type theory for synthetic ∞-categories. *Higher Structures*, 1(1), 147–224.

Schreiber, U. (2013). *Differential cohomology in a cohesive ∞-topos*. arXiv:1310.7930.

Sellars, W. (1956). Empiricism and the philosophy of mind. *Minnesota Studies in the Philosophy of Science*, 1, 253–329.

The Univalent Foundations Program (2013). *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study.

Synthetic Collective Intelligence: A Native Foundation in Simplicial Type Theory and Twisted Type Theory (source manuscript, 2026).