

Architecture of Wholeness

A Candidate Framework for the Relational Organisation of Complex Systems

Research Note	RN001
Program	Architecture
Status	Working Hypothesis
Version	0.9
Date	June 2026
Related Research Notes	RN002 — Meteorology of Complex Systems RN003 — Human–AI Iterative Synthesis
Part of	Architecture of Wholeness Research Program

Research Programme Notice

This document forms part of an ongoing research programme. It presents working research hypotheses intended for discussion, testing, critique, and possible falsification. It should not be interpreted as an established scientific theory.

Context

This document did not originate as a planned research project.

It emerged during the development of the Ethical Companies project, when the original question — "How can we recognise an ethical company?" — gradually expanded into a much broader one:

What kinds of relational organisation enable healthy, adaptive, and long-term sustainable systems?

Over time, the same question began appearing across multiple domains, including companies, communities, non-profit organisations, human teams, AI Councils, autonomous AI agents, mycelial networks, and complex ecosystems.

This observation led to the working hypothesis that these may not represent fundamentally different phenomena, but rather different manifestations of the same underlying relational principle.

Working Hypothesis

Architecture of Wholeness may serve as a candidate framework for understanding the relational organisation of complex systems.

In other words, universal principles governing the organisation of relationships may recur across different classes of systems regardless of their substrate.

These principles could potentially be observed in biological ecosystems, human organisations, economic networks, open communities, autonomous AI societies, and hybrid human–AI systems.

This hypothesis does not claim that these systems are identical.

It proposes only that they may share similar topological principles governing the organisation of relationships.

A Shift in Perspective

Original question

How can we build better AI?

Emerging question

What kind of relational organisation is most appropriate for a given class of problems?

This perspective suggests that no single organisational structure is universally optimal.

Instead, the most appropriate topology is likely to depend on the nature of the problem being addressed.

Hierarchical structures may be advantageous for linear problems.

Network-based structures may better support complex adaptive problems.

Research-oriented environments may benefit from high cognitive diversity.

Ethical decision-making may require organisational forms that encourage deliberation, trust, and distributed perspectives.

Core Assumption

The primary object of inquiry should not be individual agents or isolated nodes.

Instead, the focus should be placed on the architecture of relationships between them.

The central question therefore shifts from:

"How intelligent is this agent?"

to:

"What kinds of relationships between agents enable new systemic properties to emerge?"

Proposed Experiments

Experiment 1 — Incomplete World

Each agent receives only a partial view of the available information.

No individual agent has access to the complete picture.

The objective is to observe whether cooperation emerges spontaneously and how this need becomes reflected in the evolving topology of the network.

Experiment 2 — Paradox Injection

A problem is introduced that cannot be solved using the group's existing organisational structure.

The experiment investigates whether a new relational organisation emerges spontaneously in response.

Experiment 3 — Problem → Topology

The same group of agents is asked to solve different categories of problems.

The objective is to investigate whether different classes of problems naturally give rise to different relational organisations.

Note on Experimental Design

The implementation environment for these experiments has not yet been determined.

Multi-agent AI collaboration platforms represent one possible candidate, but the operationalisation of these experiments remains an open research question.

Candidate Metrics

Relationship Innovation Rate

How frequently do genuinely new relationship types or functional roles emerge without being explicitly designed?

Novel Concept Birth

Does the system generate genuinely new concepts or modes of reasoning, rather than merely producing new combinations of existing language?

Topological Plasticity

How rapidly can the system reorganise its relational structure in response to changing conditions?

Bridge Emergence

How readily do bridging nodes emerge between previously disconnected clusters?

The operational definition and measurement of these metrics remain subjects for future research.

Methodological Considerations

This hypothesis is speculative.

It does not constitute an established scientific theory.

Care must be taken to distinguish analogy from genuine regularities, anthropomorphic interpretation from measurable phenomena, and correlation from causation.

Likewise, it remains necessary to determine whether the observed patterns emerge spontaneously or merely reflect structures already embedded within AI training data.

An equally important question remains unanswered:

What empirical result would falsify this hypothesis?

What evidence would demonstrate that Architecture of Wholeness is not a useful scientific framework, but merely an insightful metaphor?

Answering this question is essential before the framework can evolve from a research note into a testable scientific hypothesis.

Open Questions

Do universal relational principles genuinely exist across different substrates?

Can Architecture of Wholeness be translated into formally testable hypotheses?

Which existing disciplines—such as network science, complexity science, evolutionary biology, organisational theory, sociology, or computer science—already investigate similar questions?

Where does this framework contribute genuinely new ideas, and where does it simply provide new terminology for existing concepts?

Purpose

The purpose of this document is not to propose a new scientific theory.

Its purpose is to formulate a research hypothesis that is sufficiently explicit to be criticised, tested, refined, and potentially falsified.

If similar relational principles were ultimately shown to operate across multiple classes of complex systems, this framework could provide a promising direction for future interdisciplinary research within the Ethimind project.