

Connection Theory: A General Framework for Adaptive Network Dynamics

From Temporally Extended Existence to Emergent Dynamic Properties

Kai Wang^{1*}

¹ Independent Researcher, Australia

* Correspondence: wangkai@connection-theory.org

Version 4.6 | July 2026

Abstract

The emergence of system-level properties from collections of interacting components is observed across biological, artificial, and social systems. Although emergence has been extensively investigated in many disciplines, existing explanations are often specific to particular physical mechanisms, leaving open the question of whether more general organizational principles exist across adaptive systems with fundamentally different implementations.

This paper proposes Connection Theory, a general theoretical framework for adaptive network dynamics. Rather than beginning from the physical mechanisms of individual systems, the framework starts from the observation that system-level properties are actively realized through organization operating over time, while some systems can preserve capability-relevant organizational state across bounded inactivity and later resume operation. Based on this perspective, we introduce temporally extended existence as the primary object of investigation and develop a conceptual progression from interaction, feedback, connection, and organization to continuous organizational dynamics, from which higher-level dynamic properties are proposed to emerge.

The framework further argues that many adaptive phenomena—including learning, adaptation, intelligence, and life—may be more appropriately investigated as properties realized through organizational dynamics rather than as properties of static structures alone. It distinguishes four questions that are often conflated: the active manifestation of a dynamic property, the retention of capability-relevant organizational state, the Organizational Continuity of a resumed lineage, and the current applicability of the retained or reconstructed organization. Active manifestation requires the relevant organization to operate; capability may, in some systems, survive bounded inactivity; Organizational Continuity tracks a non-forking causal lineage under a pre-specified individuation frame; and none of these conditions by itself guarantees that the reconstructed organization remains suited to a changed environment.

Connection Theory is intended to complement rather than replace existing mechanism-specific and process-oriented theories by providing a unified organizational perspective for investigating adaptive systems across diverse physical implementations. The paper identifies feedback selection, update non-commutativity, selective coordination beyond connection inventory, capability-state localization, and lineage continuity across interruption as testable problems. It also marks organization-conditioned evidence as a failure boundary that the four-stage formation sequence does not itself resolve. The proposed framework is intended as a conceptual foundation for future mathematical models of adaptive organizational dynamics.

1. Introduction

One of the most intriguing observations across science is that systems composed of relatively simple components can exhibit entirely new properties at the collective level (Anderson, 1972; Holland, 1998). Networks of neurons generate cognition, populations of cells form living organisms, and interactions among individuals give rise to languages, institutions, and civilizations. Similar phenomena are observed across many other adaptive systems despite their substantial differences in physical implementation. Why do fundamentally new system-level properties repeatedly emerge from collections of relatively simple components?

This question has appeared repeatedly in different scientific disciplines, often under different names, yet no broadly accepted cross-domain theoretical framework has emerged (Mitchell, 2009; Bar-Yam, 1997). In neuroscience, biology, complex systems, and artificial intelligence, the emergence of system-level properties has been investigated from many different perspectives. Considerable progress has been made in identifying the mechanisms underlying specific emergent phenomena, and the concept of emergence has become a widely accepted description for behaviors that cannot be understood solely by examining individual components. Yet the widespread use of the term "emergence" also raises another question. Does emergence itself constitute an explanation, or is it primarily a description of a phenomenon that still requires a more general theoretical account?

One possible approach is to investigate each emergent phenomenon independently, seeking explanations that are specific to its underlying mechanisms. This strategy has produced many important discoveries and remains essential

for understanding individual systems. An alternative possibility, however, is that at least some emergent phenomena may reflect organizational principles that are shared across otherwise very different adaptive systems. If such shared principles exist, how can they be identified without relying on the specific physical mechanisms of any individual system?

While this question originally arose from the study of emergence across natural systems, recent advances in artificial intelligence have made it increasingly difficult to ignore (LeCun et al., 2015; Vaswani et al., 2017). Artificial systems with fundamentally different physical implementations are beginning to exhibit capabilities once thought to be unique to biological intelligence, suggesting that the search for general organizational principles may now be both scientifically meaningful and practically important.

The search for such general principles is motivated by more than theoretical simplicity. If at least some emergent phenomena share common organizational principles, understanding these principles could provide a unified perspective across multiple disciplines. Rather than developing independent explanations for each individual phenomenon, it may become possible to explain diverse emergent behaviors within a common theoretical framework, complementing rather than replacing existing mechanism-specific theories. These considerations suggest that the central challenge may not be to explain each emergent phenomenon independently, but to understand how organization itself develops and gives rise to new system-level properties.

This possibility naturally leads to a further question: what kinds of principles could remain valid across systems whose physical implementations are fundamentally different? To address this question, we begin with the working hypothesis that adaptive systems with fundamentally different physical implementations may nevertheless share common organizational principles. Rather than focusing on mechanisms specific to neurons, genes, algorithms, or social institutions, we ask whether there exist higher-level principles governing the emergence and long-term organization of adaptive systems.

This line of reasoning also suggests a change in perspective. Across many areas of science, an important distinction already exists between structural properties and dynamic properties. Aerodynamics investigates moving air rather than stationary air, and fluid dynamics studies flowing liquids rather than static fluids, because many properties of interest are realized only through ongoing dynamics. This observation raises a natural question: could at least some higher-level properties of adaptive systems likewise be realized through organizational dynamics rather than through static structure alone, while capability-relevant organizational states may, in some systems, remain preserved across bounded interruptions of operation?

To explore this possibility, we propose Connection Theory, a general theoretical framework for adaptive network dynamics. In this paper, adaptive denotes the capacity for consequence-sensitive organizational change, not a guarantee that every change improves environmental fit. Rather than explaining individual emergent phenomena in isolation, Connection Theory seeks to identify organizational processes that may be shared across diverse adaptive systems and to generate experimentally testable hypotheses across biological, artificial, and social systems.

Specifically, this paper makes five primary contributions. First, it defines adaptive network dynamics as a common research object across diverse systems capable of consequence-sensitive organizational change. Second, it proposes a conceptual progression from interaction, feedback, and stabilized connection to continuously coordinated organization. Third, it distinguishes active manifestation, retained capability, lineage continuity, and current applicability from the outset. Fourth, it separates the lineage-bearing substrate of Organizational Continuity from property-relative capability-relevant persistence state. Finally, it derives a compact set of testable predictions that distinguish temporal organizational dynamics from static structure and cross-domain analogy alone.

The present framework builds upon previous work in emergence, complex systems, cybernetics, network science, systems theory, process-oriented biology, dynamical approaches to cognition, enactivism, and artificial intelligence, while proposing a common object of investigation: adaptive organization as it operates, changes, preserves capability-relevant state, and continues or branches across interruption.

2. What Should Be Explained?

If emergent properties cannot be fully understood by studying individual components alone, then the object of explanation must lie elsewhere. But where? One possibility is that the answer lies in the organization of the system itself. Before asking how organization gives rise to new system-level properties, we must first clarify what organization actually refers to.

Traditionally, organization is often understood as a relatively stable arrangement of components. Biological organisms possess anatomical structures, neural systems exhibit characteristic connectivity, and human societies develop institutions and hierarchies. Such observations naturally encourage the view that organization is primarily a

structural property. But is a static description sufficient to explain systems that continuously change while remaining identifiable over time?

Closer examination of adaptive systems reveals a striking observation. The molecular composition of living organisms changes continuously. Neural connections are constantly modified through learning. Organizations gain and lose members, while artificial systems continuously update their internal parameters and representations. Across these examples, one feature appears repeatedly: the organization persists despite continuous changes in its components. These observations suggest that organization cannot be fully characterized as a static arrangement of components.

Static observations capture only instantaneous states of adaptive systems. Observation across time reveals the organizational processes through which a system changes while remaining historically identifiable.

Figure 1 - Snapshot Observation vs Temporal Observation




Figure 1. From snapshot observation to temporally extended organization. A snapshot captures an instantaneous configuration, whereas observation across time reveals the organizational process and historical relations that a static view omits.

A possible source of this intuition is that observations are often made as snapshots. A snapshot naturally presents an adaptive system as if it were a static object. However, the apparent stability observed in a single snapshot does not imply that the underlying organization is static.

A useful way to think about active realization is to consider a vortex in a flowing river. A vortex appears to be a stable object, yet none of the water molecules composing it remain the same over time. What persists while the vortex is active is not a fixed collection of matter, but a pattern continuously realized through flow.

The vortex establishes a claim about active manifestation, not a universal requirement of uninterrupted execution. If the sustaining flow ceases and no vortex-specific state remains from which that same organized pattern can resume, the vortex disappears. Other systems may suspend active operation while preserving organization-relevant state in a durable substrate.

Four questions must therefore be separated from the outset.

Active manifestation asks whether the relevant organization is operating now. Retained capability asks whether capability-relevant persistence state survives an interruption and can support later reconstruction. Organizational Continuity asks whether the resumed system remains connected to the earlier system by a qualifying non-forking causal lineage under a pre-specified individuation frame. Current applicability asks whether the retained or reconstructed organization remains suited to present environmental conditions. These questions can receive different answers for the same system.

Systems vary in their dependence on uninterrupted operation. A vortex lies near an operation-dependent endpoint: its active organization is inseparable from current flow, and little individualized dormant state remains after complete cessation. Many engineered systems lie closer to a state-retentive endpoint, where active execution can

stop while parameters, routing state, memory, or other capability-relevant organization remain encoded. Biological systems occupy multiple intermediate positions. This spectrum concerns how organizational capability can persist; it does not make static structure alone sufficient to explain an actively realized property.

From this perspective, an adaptive system is better understood not as a static object that merely changes over time, but through a temporally extended organizational history. Active system-level properties are realized through operation. Some systems also preserve capability-relevant state across bounded inactivity. Whether a resumed system retains a capability, remains in the same lineage, or remains applicable in a changed environment must be evaluated separately.

This account of temporally extended existence is not offered as a new claim that adaptive organization persists despite continuous component turnover—closely related observations have been made before, most notably in the concept of operational closure developed in the theory of autopoiesis (Maturana & Varela, 1972, 1980; discussed in full in Section 8.3). Here, “temporally extended” means that the explanatory object is a history rather than a single snapshot; it does not require uninterrupted execution. What remains open is a different question: through what process does such organization come to exist in the first place?

The remainder of this paper develops one candidate answer to this generative question. Sections 3 through 6 introduce four such stages—interaction, feedback, connection, and organization—using a derivation strategy described below.

Figure 2a - Ontological Foundation: From Reality to Temporally Extended Existence




Figure 2a. The ontological foundation of Connection Theory. Reality is described through continuous processes rather than static configurations, from which temporally extended existence is derived as an ontological status. The concept identifies an adaptive system through an organizational history rather than any one snapshot; it does not require uninterrupted execution. Figure 2b addresses the formation of organization, not the later questions of continuity or applicability.

Concept Derivation Strategy

This paper adopts a simple derivation strategy for introducing each foundational concept in Sections 3–6. Rather than presenting interaction, feedback, connection, and organization as definitions supplied in advance, each concept is introduced only after the explanatory capacity of the concept preceding it reaches an explicit, storable boundary.

Accordingly, Sections 3–6 each follow the same structure:

Observation. What is observed about adaptive systems at this stage.

Current Concept. What the concept established in the previous section can explain.

Failure Boundary. A specific phenomenon the current concept cannot explain.

Termination Check. Can the explanatory gap be resolved without introducing a qualitatively new explanatory object—for instance, by extending the current concept with a parameter, rather than requiring a new kind of state? If yes, the current concept is retained in extended form and the derivation stops here. If no, proceeding to a new concept is warranted.

Canonical Instantiation. A brief indication that systems matching the failure boundary exist in the world. This step establishes existence, not validation.

Name. Only after the preceding steps is the new concept named.

This strategy concerns what Reichenbach (1938) called the context of discovery—how a concept comes to be proposed—rather than the context of justification—whether a proposed concept survives systematic attempts at falsification. The latter question is addressed separately by a deliberately adversarial counterexample search. Briefly, this procedure proceeds by (i) identifying the weakest implicit assumption behind a candidate concept, (ii) predicting, in advance of any test, what a counterexample would look like if the concept were false, and (iii) deliberately selecting a well-documented reference system chosen because it is likely to expose that predicted failure—not a system chosen merely to illustrate the concept—and checking whether the concept survives. A concept that survives is treated as having withstood one round of validation; a concept that fails is revised or, if the failure is severe enough, withdrawn.

A concept introduced through the derivation strategy below therefore carries a specific, limited status: it has a stated reason for being introduced (discovery), not yet a demonstrated resistance to a deliberately adversarial counterexample search (justification). Section 8.5 notes that none of the four concepts derived in Sections 3–6 has yet undergone such independent adversarial testing—the derivation below establishes only that each concept has a reason to exist, not that it has been shown to survive scrutiny.

The Termination Check step borrows, as a heuristic intuition only, the distinction between parameter extension and state-space expansion in dynamical systems. It is not applied here in a formal mathematical sense; a rigorous dynamical treatment is left to future work (Section 8.5).

3. Interaction

Observation. Adaptive systems are never observed in isolation. Living organisms exchange matter and energy with their environment; neural systems receive and transmit signals; social systems exchange information and resources across individuals and institutions; artificial systems receive inputs and produce outputs continuously during operation.

Current Concept. At this point, the only established concept is temporally extended existence itself (Section 2)—the claim that an adaptive system must be understood across an organizational history rather than reduced to a single instantaneous state.

Failure Boundary. Temporally extended existence, on its own, does not explain how a system acquires organization that is sensitive to conditions beyond an isolated instantaneous state. A closed system can continue to evolve according to its internal dynamics without external exchange, but autonomous state evolution alone does not account for the acquisition of new information about environmental conditions or for causal coordination among otherwise distinct components. Adaptive organization therefore requires causal exchange among components and, for open adaptive systems, across the system boundary.

Termination Check. Could causal exchange be treated merely as an unstated background condition of temporally extended existence? We do not think so. Temporally extended existence identifies the explanatory object across a history, whereas interaction identifies a relation through which the state of one component or an environment can alter the state of another. Distinguishing endogenous state evolution from relational exchange requires an explanatory object not supplied by historical individuation alone.

Canonical Instantiation. Synaptic transmission between neurons, metabolic exchange across a cell membrane, and message-passing between agents in a multi-agent system are all instances of this ongoing exchange, despite substantial differences in physical implementation.

Name. We refer to this ongoing exchange between an adaptive system and other components or its environment as interaction.

Interaction explains how components and environments can causally influence one another and how an adaptive system can acquire information not contained in an isolated instantaneous state. It does not explain why some interactions leave a lasting trace on future behavior while most do not. This boundary is addressed in Section 4.

4. Feedback

Observation. Adaptive systems experience a very large number of interactions over their lifetime, yet only a small fraction appear to have any lasting influence on future behavior. A neuron fires in response to countless inputs, but

only specific, repeated patterns of input measurably alter its future response. A person has many brief exchanges, but only a few reshape how they act afterward.

Current Concept. Interaction (Section 3) explains that systems can change state in response to exchange. It does not specify which interactions matter.

Failure Boundary. If every interaction contributed equally and permanently to a system's state, adaptive systems would accumulate unlimited, undifferentiated change and rapidly lose stable organization. Interaction alone offers no account of why some interactions are retained in future behavior and others are not.

Termination Check. Could this be handled by simply extending interaction with a decay parameter—every interaction leaves a trace that fades at some fixed rate? This resolves part of the problem but not all of it: canonical systems (below) show that whether an interaction is retained depends not only on time elapsed but on whether it is compatible with, or reinforced by, the system's existing state—a history-dependent selection process, not uniform decay. This requires the update to depend on more than a fixed decay of each interaction—specifically, on the differential influence of past interactions on future ones. Whether that dependence is best represented as an additional state variable or as an explicitly history-dependent update is a modeling choice deferred to future work (Section 2; 8.5); the Termination Check establishes only that a single decay parameter is insufficient.

Canonical Instantiation. The distinction between reflexive responses that leave no lasting trace (e.g., the patellar reflex) and habituation or sensitization in *Aplysia*, where repeated stimulation measurably and persistently changes the strength of future responses, cleanly separates interactions with no lasting influence from interactions that reshape future behavior.

Name. We refer to this history-dependent, selective influence of past interactions on a system's future organization as feedback. Feedback names a selection criterion—which past interactions continue to influence future organization—not a rule for how the selected influence updates organization; the latter is left open (Section 2; 8.5).

4.1 Feedback Availability and Selective Sampling

Feedback can act only through consequences that become causally available to the system. A consequence that never returns to the system, leaves no retained proxy, and is not reconstructed by another process cannot influence later organization through feedback.

The available feedback is generally not a neutral sample of environmental outcomes. Existing organization influences which interactions occur, which consequences return to the system, how quickly and reliably those consequences are coupled back, and which signals are retained. Outcomes of unchosen actions, silently lost participants, delayed effects, and consequences prevented by the current organization are therefore systematically underrepresented unless an additional process actively samples or reconstructs them.

Feedback selection consequently operates on an organization-conditioned, partially available, and potentially biased sample rather than on a complete record of environmental value. This limitation does not change the definition of feedback, but it establishes a boundary on what feedback alone can correct.

Feedback explains why some interactions leave a lasting influence. It does not explain how that influence becomes a stable, standing relationship rather than a residual trace that could still be overwritten by the next interaction. This boundary is addressed in Section 5.

5. Connection

Observation. Adaptive systems develop relationships among their components that persist across many subsequent interactions, resisting erasure by ordinary noise or a single conflicting input—a habit, a synapse potentiated over repeated use, a long-standing collaboration between two people.

Current Concept. Feedback (Section 4) explains that some interactions leave a lasting influence on future behavior. It does not explain why that influence, once left, becomes stable rather than being overwritten by the next feedback-bearing interaction.

Failure Boundary. If feedback only ever produced a transient adjustment with no further consolidation, every new feedback event would be equally capable of overwriting the last, and no relationship between components could accumulate enough stability to persist over an extended timescale. Feedback alone gives no account of why some influences become resistant to subsequent perturbation while others remain as fragile as the interaction that produced them.

Termination Check. Could stability be captured simply by a strong-enough single feedback event, rather than a separate concept? Canonical evidence below indicates otherwise: stability in these systems arises from a second, distinct process—a consolidation step operating on a different timescale from the initiating feedback event, and one

that can be selectively blocked without preventing the initial feedback effect. This is a qualitatively distinct mechanism, not a stronger instance of feedback, and requires its own state (consolidated vs. unconsolidated) that feedback's transient trace does not carry.

Canonical Instantiation. Long-term potentiation and synaptic tagging-and-capture show exactly this two-tier structure: an initial, transient synaptic tag can decay within hours unless a separate, slower stabilization process consolidates it into a lasting structural change.

Name. We refer to such a stabilized, feedback-derived relationship between components as a connection.

A connection explains how a single relationship between components becomes stable. It does not explain how a large number of independently stabilized connections come to behave as a single, coordinated system rather than an unstructured accumulation of stable but unrelated relationships. This boundary is addressed in Section 6.

6. Organization

Observation. Adaptive systems exhibit coordinated, system-level behavior—an organism's physiology, a neural circuit's function, a firm's operations—that is not adequately captured by simply listing the stable connections present in the system.

Current Concept. Connection (Section 5) explains how individual relationships between components become stable. It does not explain how a collection of such connections comes to function as an integrated whole.

Failure Boundary. A system could, in principle, accumulate many stable connections operating independently of one another—analogous to a set of separately reinforced habits with no coordination between them. The presence, number, or weights of connections do not by themselves specify how interactions are routed, timed, synchronized, amplified, sequenced, and combined across those connections. The operationally relevant dimensions differ by system, and any one dimension may be absent or degenerate. Connection alone therefore cannot distinguish an uncoordinated collection from a system whose stable relationships jointly produce coherent behavior.

Termination Check. Could coordination be treated as an automatic side effect of accumulating enough connections, without a further concept? The evidence below suggests not: coordination can be selectively lost or reconfigured while the underlying connection set remains largely intact, indicating that coordination is a separable property of the system, determined by more than the mere presence, count, or weights of stable connections. Whether this additional dependence is represented as an extra state variable is a modeling choice deferred to future work (Section 2; 8.5).

Canonical Instantiation. Structural and functional connectomics show that overlapping sets of stable anatomical connections can support markedly different, task-dependent functional configurations—a degeneracy in which the same connection set is coordinated differently depending on context. This indicates that the coordination layer is not reducible to the connection set itself.

Name. We refer to the continuously coordinated, state-dependent integration of stable connections as organization. Depending on the system, this coordination may be expressed through routing, timing, sequencing, gain, synchronization, combination, or other relations. A distribution over temporally ordered traversals is one possible operational representation, not the definition of organization itself. In an order-insensitive system, the sequencing dimension may be degenerate without eliminating organization in other dimensions.

6.1 Non-Commutative Updating

Where an update produced by an interaction depends on the current organizational state, order can matter because a later update acts on a state already altered by an earlier one. Let U_i and U_j denote two update operations applied to organizational state O . A local measure of update non-commutativity can be written as $N_{ij}(O) = d[U_j(U_i(O)), U_i(U_j(O))]$, where d is a pre-specified distance between resulting organizational states.

An aggregate non-commutativity measure X can be estimated from independent calibration interventions across initial states and update pairs, then used prospectively to predict the magnitude of held-out causal order effects. Where X is approximately zero, matched order swaps should not systematically change the resulting organization. Where X is larger, otherwise matched changes in order should produce correspondingly larger organizational differences.

6.2 Failure Boundary: Organization-Conditioned Evidence

Existing organization influences which interactions occur, which consequences become causally available as feedback, and which feedback-bearing changes are consolidated. The evidence that later supports an organization is therefore partly conditioned by that organization.

Nothing in Interaction, Feedback, Connection, or Organization distinguishes persistence supported by broader environmental fit from persistence produced by control over the evidential sample. Uniform decay does not resolve this boundary because it weakens environmentally supported and self-confirming organization indiscriminately.

This is a stated failure boundary of the foundational sequence, not a fifth formation stage or a new Connection Theory primitive. The four-stage sequence supplies a necessary organizational scaffold for consequence-sensitive updating, not a sufficient condition for successful adaptation or continued environmental fit. The long-term dynamics of self-confirming closure and corrective re-adaptation remain outside the present derivation.

Relationship to operational closure

The organization defined here is closely related to operational closure as developed in the theory of autopoiesis (Maturana & Varela, 1972), where a system's identity is maintained by a network of processes that regenerate the relations that produced them (see Section 8.3). We take these notions to be compatible rather than competing: operational closure specifies a criterion for self-maintaining organization, while the derivation above offers one account of the more elementary stages—interaction, feedback, and connection—through which coordinated organization may be assembled. Section 6.2 marks a boundary that this formation account leaves unresolved.

Organization, so derived, provides the coordinated basis for consequence-sensitive adaptive behavior. It does not by itself establish that the resulting organization is optimal, environmentally successful, or protected from self-confirming feedback.




Figure 2b. Derivation of organization from interaction, feedback, and connection. Each transition passes through an explicit Termination Check: if the explanatory gap can be resolved by extending the current concept, the derivation halts; otherwise, a new concept is introduced. Organization-conditioned evidence is a post-formation failure boundary, not a fifth formation stage. The resulting notion of organization is compatible with, rather than a replacement for, operational closure.

7. Organizational Continuity

Observation. Adaptive systems are commonly regarded as “the same” system after an interruption of the processes that constitute their Organization—a computer is considered the same computer after a reboot, a desiccated tardigrade the same organism after revival, a river the same river after a seasonal dry-up ends. Yet Organization (Section 6) is a continuously coordinated integration that, by its own definition, does not operate during an interruption.

Current Concept. Connection (Section 5) explains how a stabilized relationship can persist without continuous reinforcement. Organization (Section 6) explains how connections are coordinated into integrated behavior while operating. Neither concept specifies what licenses calling the system that resumes after an interruption the same system, rather than a new instance that merely resembles it.

Failure Boundary. Two systems can share an identical Connection-layer encoding without occupying the same system lineage. A type-identical substitute or independently initialized copy can reproduce capability without inheriting the actual provenance of the earlier system. Conversely, copying dormant state does not by itself create a second system lineage unless more than one independently continuable successor is realized. Resemblance, duration, and copying alone therefore do not resolve continuity.

Termination Check. Could this be handled by extending Connection with a persistence-duration parameter, or by requiring uninterrupted operation? Neither resolves the lineage problem. What is required is a pre-specified individuation frame and a method for tracing whether a qualifying causal path through a continuity-bearing substrate connects the earlier and resumed system without forking.

Canonical Instantiation. Across ischemia tolerance, seasonal river desiccation, unpowered digital storage, and tardigrade anhydrobiosis, continuity judgments track actual provenance and causal participation in resumption. At a cell-level individuation frame, mitosis produces two successor cell lineages rather than one unique continuation. In engineered systems, a single checkpoint restoration after the predecessor has terminated may preserve a non-forking path, whereas parallel restorations from the same ancestral state create successor branches.

Name. We refer to this lineage criterion as Organizational Continuity. Given a pre-specified individuation frame—comprising a system boundary, lineage level, and adjudication interval—Organizational Continuity holds where the resumed system is connected to the earlier system by a qualifying non-forking causal lineage carried through an identifiable continuity-bearing substrate that causally contributes to resumption. Within a fixed frame, the judgment is property-invariant: losing or recovering a particular capability does not by itself determine lineage continuity.

7.1 Operational Identification Across Interruption

The individuation frame is stipulated before continuity is evaluated. Its selection is methodological, but the provenance facts within that fixed frame are empirically assessable. Different lineage levels may yield different judgments without contradiction; for example, cell-level continuity and DNA-molecule continuity are different questions. Organizational Continuity is therefore boundary-conditional rather than boundary-free, while remaining independent of which capability is chosen for testing.

A candidate continuity-bearing substrate is identified by tracing custody, provenance, and causal participation across the interruption. Where feasible, a lineage-preserving transformation is compared with a lineage-external substitution matched for type, state, and capability. A candidate set is lineage-bearing where substituting matched states without causal descent from the earlier system cuts all qualifying provenance paths to the resumed system. Distributed cases may require identifying a minimal lineage cut set rather than a unique component.

Copying dormant state is not itself a fork. A fork occurs when two or more independently continuable bounded systems are instantiated from the same ancestral state within the adjudication interval. A single restoration can therefore remain non-forking where the predecessor has terminated and the restored system inherits a qualifying provenance path; concurrent or later parallel realizations form successor branches.

The no-fork condition has conceptual parallels with non-branching clauses in philosophical accounts of personal identity (Nozick, 1981; Parfit, 1984). Here it is formulated as a substrate-general, operationally traceable criterion for organizational lineage across biological, engineered, and institutional systems. Because Organizational Continuity is explicitly a lineage relation, its dependence on facts elsewhere in the lineage is not a defect: lineage relations are inherently external and branch-sensitive. The judgment is indexed to the pre-specified adjudication interval; extending that interval may change the verdict where a later branch is instantiated, which is a feature of a branching-sensitive lineage criterion rather than an inconsistency.

Under this criterion, a single migration of dormant organizational state to a new physical carrier can preserve lineage where the predecessor is terminated, no parallel realization is instantiated within the adjudication interval, and the resumed system inherits a qualifying provenance path. Continuity is carried by the non-forking causal provenance of the migration, not by persistence of any particular physical carrier.

Capability-relevant persistence state answers a different question. For a pre-specified property P, interruption condition, and reconstruction protocol, selective removal, randomization, or replacement can identify which retained states causally support reconstruction of P. Because such state may be distributed or redundant, the result may be a family of minimal sufficient sets rather than one individually necessary component.

Capability interventions establish causal relevance to reconstruction of P; they do not establish lineage. Conversely, provenance and no-fork evidence establish Organizational Continuity within the fixed frame; they do not establish retention of every capability. Identity labels used in ordinary or legal contexts may add further criteria beyond this empirical lineage relation.

Organizational Continuity concerns the lineage of the resumed system; it is not necessary for the reappearance of an equivalent capability. A copied or independently reconstructed system may exhibit equivalent performance without continuing the earlier lineage. Conversely, behavioral equivalence alone does not establish Organizational Continuity.

Organizational Continuity does not weaken the claim that active manifestation requires the relevant Organization to operate. It explains how a lineage may pass through an interval in which that manifestation is suspended, provided that a qualifying lineage-bearing substrate persists within its tolerance window.

Organizational Continuity is a lineage relation under a fixed individuation frame, not a criterion of performance, fitness, or current applicability. A system may remain lineage-continuous and faithfully reconstruct its previous organization while becoming poorly suited to an environment that changed during the interruption.

8. Discussion

8.1 What Is the Scope of Connection Theory?

Throughout this paper, we have argued that adaptive systems with fundamentally different physical implementations may nevertheless be investigated through common organizational questions. Connection Theory asks whether a sequence of organizational relations—interaction, feedback selection, stabilized connection, and coordinated organization—together with a separate lineage criterion across interruption, can clarify similarities and differences among biological, artificial, and social systems.

Scientific disciplines routinely distinguish between structural properties and dynamic properties. Aerodynamics investigates moving air because lift is realized through motion; fluid dynamics studies flowing liquids because vortices and turbulence cannot be inferred from material composition alone. In such cases, static structure remains causally important but is not by itself the phenomenon to be explained.

As established in Section 2, four questions must be separated: whether a property is actively manifested, whether the state required to recover that capability remains retained, whether a resumed system continues the earlier lineage under a fixed individuation frame, and whether the retained or reconstructed organization remains applicable in the current environment.

These distinctions also prevent biological life and computational intelligence from being forced into the same interruption model. Biological life ordinarily depends on ongoing material and energetic self-maintenance, so interruption may rapidly damage the substrate required for resumption. A computational system may instead preserve capability-relevant parameters, routing state, memory, and other organization in a persistent medium, suspend execution, and reconstruct the corresponding dynamics later. The common claim is that life and intelligence are realized dynamically, not that every realizing system must execute continuously in the same physical sense.

Section 6.2 identifies a further limit of the foundational sequence: an organization may condition the evidence later available to evaluate it. Connection Theory states this boundary but does not attempt to derive the long-term corrective dynamics required to overcome it.

From this perspective, the principal contribution of Connection Theory is a proposed research direction rather than a claim of priority over process or dynamical accounts. It redirects attention from the question of how static structures mysteriously possess higher-level properties to the question of how organization operates, changes through feedback, preserves capability-relevant state, and resumes or diverges across interruption. The value of this direction depends on whether it yields discriminative, falsifiable results, addressed in Section 8.4.

8.2 Temporal Structure, Observation, and Intervention

One implication of the present framework concerns how adaptive systems are observed. Discrete measurements and structural descriptions are indispensable for identifying components, connections, and instantaneous organization. They are insufficient when the explanatory object includes transitions, feedback-dependent updating, non-commutative order effects, interruption, reconstruction, and lineage.

Continuous observation can reveal the trajectory through which one organizational state becomes another: which interactions were selected by feedback, which consequences became available, which changes persisted, how coordination was redistributed, and how a system responded after interruption. The scientific value of such observation is not merely that it supplies more samples, but that it can preserve temporal order and transition structure.

Accordingly, a fair comparison between structural and dynamical explanations should control for information quantity, observation duration, and model capacity. Update non-commutativity should be estimated from matched order interventions rather than inferred merely from historical association.

Continuous observation does not replace structural observation. Structural measurements specify the current organization and the substrate that constrains possible dynamics; temporally ordered measurements specify how that organization is maintained, updated, suspended, and reconstructed. The two descriptions are complementary rather than competing.

The distinction is therefore not simply between a photograph and a video. It is between a state-only account and an account that represents causally relevant transitions under matched evidential conditions.

Figure 3 - Four Distinct Questions Across Interruption




Figure 3. Active manifestation, retained capability, Organizational Continuity, and current applicability. Operation is required for current expression; capability-relevant state may persist during inactivity; lineage continuity depends on a qualifying non-forking causal path under a fixed individuation frame; and successful reconstruction does not guarantee continued environmental fit.

8.3 Relationship to Existing Frameworks

The present framework is intended to complement rather than replace existing theories of adaptive systems. Neuroscience, evolutionary biology, artificial intelligence, control theory, and complex systems research each investigate mechanisms that are specific to particular classes of adaptive systems. Connection Theory addresses a different question: whether these diverse systems may also share common organizational principles governing the formation, active realization, retained capability, and lineage continuity of higher-level dynamic properties.

Process-oriented philosophy of biology explicitly treats organisms as processes stabilized through continuous change rather than as substances defined by fixed material composition (Dupré & Nicholson, 2018). Connection Theory does not claim priority for that ontological orientation. Its narrower proposal is to connect process realization to a staged organizational vocabulary and to distinguish active manifestation, recoverable capability, lineage continuity, and environmental applicability.

Dynamical systems approaches to cognition and development likewise analyze cognitive behavior as temporally extended dynamics rather than as a sequence of static representations (Thelen & Smith, 1994; van Gelder, 1998). Enactivist work emphasizes cognition as embodied, history-dependent coupling between an organism and its environment (Varela et al., 1991; Thompson, 2007). Connection Theory is compatible with these traditions while extending the comparison to artificial and social systems and emphasizing that dynamical realization need not imply uninterrupted execution where capability-relevant state can be preserved and reconstructed.

The theory of autopoiesis (Maturana & Varela, 1972, 1980) proposed operational closure as the defining organizational property of living systems: an autopoietic system is a network of processes that continuously regenerates the relations that produced it. This account and the temporally extended existence introduced in Section 2 address closely related territory, and we do not claim priority over the observation that adaptive organization persists through component turnover. Connection Theory instead asks through what more elementary organizational stages such coordinated organization may be assembled.

Complex systems research has demonstrated that interactions among many components can produce collective behavior that cannot be understood solely from individual elements (Holland, 1998; Mitchell, 2009). Connection Theory is fully consistent with this perspective. However, rather than treating emergence as the endpoint of explanation, the present framework attempts to further examine how continuously operating adaptive organization gives rise to dynamic system-level properties.

Likewise, network science provides powerful mathematical descriptions of relationships among components (Newman, 2010; Barabási, 2016). Connection Theory does not replace these representations. Instead, it proposes that stable connections become scientifically significant primarily through the state-dependent organizational dynamics that they collectively support.

Cybernetics emphasizes feedback, regulation, and control within adaptive systems (Wiener, 1948; Ashby, 1956). Connection Theory incorporates feedback as one of its fundamental organizational processes while emphasizing that available feedback is selectively sampled and may be conditioned by the organization being updated.

Evolutionary theory explains how adaptive systems change across generations through variation and selection (Darwin, 1859; Maynard Smith, 1982). Connection Theory focuses on a different but complementary level of analysis: the organizational dynamics through which adaptive systems maintain themselves and continuously develop during their own existence.

Rather than competing with existing mechanism-specific disciplines, Connection Theory seeks to provide a common organizational perspective across them. Its purpose is not to replace domain-specific explanations, but to identify general organizational questions that may remain meaningful despite substantial differences in physical implementation.

8.4 Testable Predictions

A general framework becomes scientifically useful only where its claims can fail. The compact set below tests operational dependence, retained capability, capability-state localization, causal order sensitivity, coordination beyond connection inventory, and graded feedback availability.

Prediction 1 — Operational dependence of active manifestation

For a system-level property operationally specified in advance and linked to an identified organizational process, experimentally interrupting or selectively suppressing that process while preserving the relevant structural substrate should reduce or suspend the active manifestation of the property on a timescale corresponding to the interrupted process. The appropriate comparison is a targeted-interruption condition against sham interruption and structural-damage controls. This prediction is not supported if the property remains unchanged despite verified cessation of the proposed realizing process, or if gross structural preservation alone predicts the outcome equally well.

Prediction 2a — Retained capability across inactivity

In a system capable of storing recovery-relevant organizational state, inactivity by itself should not eliminate retained capability. Among matched systems inactive for the same duration, a system whose capability-relevant state is preserved should resume closer to its pre-interruption performance and should require less new feedback to recover its prior adaptive trajectory than a system whose corresponding state is reset, corrupted, or omitted. A copied system may exhibit equivalent capability without satisfying Organizational Continuity; identity and capability are therefore evaluated separately. This prediction is not supported if inactivity causes the same capability loss regardless of the integrity of the preserved state.

Prediction 2b — Localization of capability-relevant persistence state

For a pre-specified organizational property P, interruption condition, and reconstruction protocol, graded removal, randomization, or replacement of candidate retained states—and of candidate combinations—should produce a reproducible reconstruction-loss profile. Where capability-relevant persistence state is distributed or redundant, the result may be a family of minimal sufficient component sets rather than a unique individually necessary component. Support requires controls distinguishing stored organizational state from generic reconstruction machinery and nonspecific damage. The result identifies causal relevance to reconstruction of P; it does not by itself identify the continuity-bearing substrate or establish a qualifying non-forking provenance path.

Prediction 3 — Update non-commutativity predicts causal order effects

Estimate update non-commutativity X on calibration interventions by comparing the resulting organizational states after matched update pairs $U_i \rightarrow U_j$ and $U_j \rightarrow U_i$. On held-out update pairs or initial states, larger X should predict a larger causal order effect Y_{causal} under a pre-specified organizational distance, whereas X approximately equal to

zero should predict no systematic order effect. Support requires the same update multiset, matched exposure, and separation of order from total interaction quantity.

Prediction 4 — Selective coordination under matched connection structure

Two systems with sufficiently matched component and connection inventories but different coordination states should exhibit discriminably different organized behavior. Conversely, selective disruption of routing, timing, gain, synchronization, sequencing, or combination should abolish or alter an organized capability while leaving much of the connection inventory intact. Support requires intervention on coordination rather than merely observing correlation, together with verification that the relevant connection structure remains substantially preserved.

Prediction 5 — Graded feedback availability

Holding outcome content, objective value, and interaction frequency constant within a pre-specified unsaturated range, experimentally increasing the accessibility, immediacy, or reliability of a consequence signal should produce a correspondingly greater causal contribution of that consequence to a bounded-horizon organizational update. The dimensions should be manipulated separately where possible. Manipulations of reliability must also control for the signal's informativeness about the outcome; otherwise, increased weight may reflect normative updating rather than availability-conditioned selection. This prediction is not supported if verified changes in signal availability produce no systematic change in update influence, after controlling for signal content and exposure. It tests feedback selection, not the long-term organizational drift consequences of biased sampling.

These predictions do not claim that every property conventionally labeled life, intelligence, learning, or adaptation will satisfy the same operational test. Each experiment must specify the property, the proposed realizing process, the comparison condition, and the criterion for failure before data are evaluated.

8.5 Limitations and Future Directions

Limitations

The present work establishes a conceptual organizational framework rather than a complete explanatory theory of adaptive systems. Several important questions remain outside the scope of the current paper.

The derivation strategy used in Sections 3–6 (Section 2) establishes a stated reason for introducing each foundational concept but does not itself constitute validation: it confirms that each concept has a demonstrated reason to exist (via the Termination Check), not that the concept withstands a deliberate, adversarial search for counterexamples. None of the four foundational concepts—interaction, feedback, connection, organization—has yet undergone such a dedicated adversarial validation pass. This is listed as a priority item for future work, alongside a formal dynamical treatment of the Termination Check heuristic used in Section 2.

Section 6.2 identifies organization-conditioned evidence only as a failure boundary of the foundational sequence. The present paper does not model the long-term dynamics of self-confirming closure or corrective re-adaptation.

Connection Theory does not attempt to replace the domain-specific mechanisms investigated in neuroscience, evolutionary biology, artificial intelligence, or other disciplines. Instead, it operates at a higher organizational level intended to complement these mechanism-specific explanations.

The present framework also does not claim that life, intelligence, learning, and adaptation have identical interruption requirements. It distinguishes active manifestation, retained capability, Organizational Continuity, and current applicability, but does not yet provide domain-independent quantitative criteria for identifying the organizational state sufficient to recover each property. The predictions in Section 8.4 therefore apply only after the property and its proposed realizing process have been operationally specified.

Continuity-bearing substrate and capability-relevant persistence state require different evidence. Continuity is assessed from provenance, causal participation, and no-fork conditions under a pre-specified individuation frame. Capability-state localization is property-relative and may require joint interventions because of distributed or redundant encoding.

Feedback availability and update non-commutativity are conditional claims. Not every consequence relevant to an observer becomes feedback for the system, and not every update process is order-sensitive.

Although the present framework identifies organizational dynamics as its primary object of investigation, it does not yet provide a complete mathematical description of organizational evolution, interruption, reconstruction, lineage, or non-commutative updating.

Likewise, the predictions proposed here require systematic experimental evaluation across biological, artificial, and social adaptive systems. Failure in one domain may delimit the scope of the proposed cross-domain motif rather

than determine the status of every other claim, so the scope and readiness of each result should be reported separately.

Future Directions

Future work may investigate how selected interactions produce persistent organizational change, how capability-relevant states are encoded across different substrates, how lineage can be measured across complex interruptions, and how reconstruction relates to resumed adaptation rather than behavioral replay.

The long-term consequences of organization-conditioned feedback availability—including self-confirming closure and corrective re-adaptation—remain a subject for future work.

The present framework also raises the possibility that adaptive systems may require organizational measures of change that complement conventional physical time, providing another direction for future theoretical development.

More generally, establishing quantitative descriptions of adaptive organizational dynamics, update non-commutativity, capability-relevant persistence state, and lineage continuity represents an essential next step toward transforming the present framework into a predictive scientific theory.

Whether the proposed framework proves broadly applicable remains an empirical question. We hope that Connection Theory provides a useful organizational perspective for investigating adaptive systems and for generating experimentally testable questions across biological, artificial, and social domains.

References

- Anderson, P. W. (1972). More is different. *Science*, 177(4047), 393–396.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Barabási, A.-L. (2016). *Network Science*. Cambridge University Press.
- Bar-Yam, Y. (1997). *Dynamics of Complex Systems*. Addison-Wesley.
- Bertalanffy, L. von. (1968). *General System Theory: Foundations, Development, Applications*. George Braziller.
- Darwin, C. (1859). *On the Origin of Species*. John Murray.
- Dupré, J., & Nicholson, D. J. (Eds.). (2018). *Everything Flows: Towards a Processual Philosophy of Biology*. Oxford University Press.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Haken, H. (1983). *Synergetics: An Introduction* (3rd ed.). Springer.
- Holland, J. H. (1998). *Emergence: From Chaos to Order*. Oxford University Press.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Luhmann, N. (1995). *Social Systems*. Stanford University Press.
- Maturana, H. R., & Varela, F. J. (1972). *De Máquinas y Seres Vivos*. Editorial Universitaria.
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- Mitchell, M. (2009). *Complexity: A Guided Tour*. Oxford University Press.
- Newman, M. (2010). *Networks: An Introduction*. Oxford University Press.
- Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Prigogine, I., & Stengers, I. (1984). *Order Out of Chaos*. Bantam Books.
- Reichenbach, H. (1938). *Experience and Prediction*. University of Chicago Press.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482.
- Thelen, E., & Smith, L. B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21(5), 615–628.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
- Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- von Foerster, H. (1974). *Cybernetics of Cybernetics*. University of Illinois.
- Watts, D. J. (1999). *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press.
- Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press.