Toward Organizational Correspondence: A Conceptual Model for Diagnosing Cognitive Integrity in AI-Native Organizations

Jussi Nyfelt

Independent research, Correspondia

PDF edition. 2026-08-06.

Abstract

Organizations increasingly think, decide, remember, and act through evolving networks of persistent representations: strategies, decisions, policies, plans, models, work products, AI-generated recommendations, operational traces, outcomes, and organizational memory. Existing theories explain important aspects of organizational cognition, including learning, distributed cognition, sensemaking, memory, information processing, and feedback. Yet they do not explicitly explain how organizations preserve meaningful semantic and causal relationships across continuously evolving organizational representation networks. This paper develops Correspondence Theory, a representation-network-centered theory of organizational cognition. The theory introduces a representation-network ontology, defines Organizational Correspondence as the emergent organizational property by which meaningful semantic and causal relationships among organizational representations are preserved over time, defines Cognitive Integrity as the dynamic organizational condition associated with sustained Organizational Correspondence, and introduces the Correspondence Loop as the recursive, event-driven network maintenance mechanism through which representation networks are maintained, revised, restored, or degraded. The paper presents a proposition architecture and research agenda for future empirical work while preserving a clear distinction between theory, observable traces, indicators, variables, measurement, and validation. The central contribution is a conceptual architecture for studying whether organizational cognition remains meaningfully and causally connected as organizations become more distributed, AI-mediated, and representationally dense.

Keywords

Organizational cognition; organizational correspondence; cognitive integrity; representation networks; AI-native organizations; organizational learning; distributed cognition; organizational memory; sensemaking; knowledge management; conceptual theory

1. Introduction

Organizations increasingly think, decide, remember, and adapt through dense networks of persistent representations. Strategies are encoded in plans, priorities, roadmaps, policies, models, architectural diagrams, work items, meeting records, AI-generated recommendations, operational traces, and lessons learned. These artifacts are not merely administrative byproducts. They are among the principal media through which organizational cognition is expressed, stabilized, transferred, revised, and acted upon. As organizations become more digitally mediated and increasingly AI-enabled, cognition is distributed not only across people and teams, but also across systems, repositories, models, agents, and organizational memory (Faraj et al., 2018; Hutchins, 1995; Kellogg et al., 2020; Raisch and Krakowski, 2021; Walsh and Ungson, 1991).

This shift creates a practical problem that is becoming increasingly difficult to ignore. Organizations do not simply need more information, better documentation, or faster decision cycles. They need the meanings and causal assumptions embedded in their representations to remain coherent as those representations change. A strategic intent must remain meaningfully related to the decisions that interpret it, the work that enacts it, the outcomes that follow from it, and the lessons later drawn from it. Policies must remain connected to governance decisions. Plans must remain connected to operational events. AI-generated artifacts must remain connected to human intent, organizational memory, and accountable decision histories. When these relationships weaken, organizations may continue to operate while gradually losing the ability to explain how intent, decision, execution, outcome, and learning relate to one another.

The difficulty is especially visible in AI-native and AI-accelerated organizations, but it is not limited to them. AI systems increase the speed, volume, and distribution of organizational representation. They can generate plans, summaries, recommendations, classifications, decisions, and actions at a scale that exceeds traditional human review cycles. At the same time, many organizations already face increasing complexity, distributed decision making, fragmented knowledge systems, rapid strategic change, and uneven organizational memory. Under these conditions, semantic inconsistency and causal fragmentation can accumulate quietly. Actors may use the same terms to mean different things. Decisions may lose their relationship to prior intent. Outcomes may be recorded without clear causal connection to the actions that produced them. Organizational memory may persist as an archive while ceasing to function as a coherent cognitive resource.

Existing organizational theories provide important foundations for understanding these dynamics, including organization design, information processing, alignment, learning, memory, sensemaking, systems theory, and sociomaterial organizing, but they do not fully explain the phenomenon at issue here (Daft and Lengel, 1986; Galbraith, 1974; Henderson and Venkatraman, 1993; Orlikowski, 2007; Venkatraman, 1989). Organizational learning explains how organizations acquire, retain, and apply knowledge from experience. Distributed cognition explains how cognition is distributed across people, artifacts, and environments. Cybernetic and systems theories explain feedback, regulation, adaptation, and control. Organizational memory explains retention and retrieval of past knowledge. Sensemaking explains how actors construct meaning under ambiguity. These perspectives illuminate essential aspects of organizational cognition. Yet none explicitly explains how organizations preserve meaningful semantic and causal relationships across continuously evolving networks of organizational representations.

This claim is grounded in organizational learning, distributed cognition, systems and cybernetic theory, organizational memory, and sensemaking scholarship (Argyris and Schon, 1978; Ashby, 1956; Bertalanffy, 1968; Hutchins, 1995; March, 1991; Walsh and Ungson, 1991; Weick, 1995).

This paper addresses that gap. Its guiding research question is: How do organizations preserve meaningful semantic and causal coherence across continuously evolving organizational representation networks? The question is deliberately conceptual. It does not ask how organizations should design a particular software system, governance process, or diagnostic instrument. It asks what theoretical architecture is needed to explain the preservation, degradation, and maintenance of meaning-bearing and causally significant relationships among organizational representations over time.

To answer this question, the paper develops Correspondence Theory, a conceptual theory of organizational cognition centered on representation networks. The theory begins from the premise that organizations continuously create, maintain, transform, and discard representations, and that these representations are connected through semantic and causal relationships. The paper then develops a set of constructs and a mechanism for explaining whether these relationships remain coherent as organizational cognition unfolds.

The paper makes six contributions. First, it introduces a representation-network ontology that treats organizational cognition as expressed through persistent representations and the semantic and causal relationships among them. Second, it defines Organizational Correspondence, a construct describing the emergent property by which those relationships are preserved over time. Third, it defines Cognitive Integrity, a second construct describing the dynamic organizational condition associated with sustained Organizational Correspondence. Fourth, it introduces the Correspondence Loop, a recursive, event-driven mechanism explaining how the representation network is maintained, revised, or degraded. Fifth, it organizes the theory into a proposition architecture linking representation-network maintenance, Organizational Correspondence, Cognitive Integrity, and AI-native organizational cognition. Sixth, it establishes a cumulative research program that connects the theory to future diagnostics, measurement, and empirical validation without redefining the canonical constructs.

The remainder of the paper proceeds as follows. The next section develops the research problem and situates it within the changing conditions of organizational cognition. The literature foundations section then reviews adjacent theories and identifies the residual explanatory gap addressed by Correspondence Theory. The paper next introduces the representation-network ontology, followed by the two foundational constructs: Organizational Correspondence and Cognitive Integrity. It then defines the Correspondence Loop as the theory's canonical mechanism and integrates the ontology, constructs, and mechanism into a conceptual model. The subsequent section presents the proposition architecture, followed by a research agenda for future empirical work. The paper concludes by discussing theoretical implications, practical relevance for AI-native organizations, limitations, and directions for future research.

2. Research Problem

Organizations have always depended on representations of what they intend, decide, remember, and do. Strategies, policies, plans, decision records, work products, process descriptions, performance reports, governance artifacts, and organizational memory all provide durable forms through which organizational cognition persists beyond any single individual or moment. What is changing is not merely the presence of such representations, but their scale, speed, heterogeneity, and interdependence. In contemporary organizations, and especially in AI-native organizations, representations are created, modified, interpreted, and acted upon across distributed human and machine participants. They are no longer occasional records of organizational cognition. They increasingly constitute the medium through which organizational cognition unfolds.

This creates a theoretical problem that is easy to observe but difficult to explain with existing conceptual resources. Organizations may possess abundant representations and still lose coherence among them. A strategic intent may be translated into a decision, but the relationship between that decision and later execution may weaken. A policy may remain formally valid while the practices it governs evolve around it. An AI-generated recommendation may enter planning activity without being clearly related to prior assumptions, evidence, or accountability structures. A learning event may be recorded but not connected back to the decision or action that produced it. In each case, the issue is not simply whether an organization has information, knowledge, memory, or communication. The issue is whether meaningful semantic and causal relationships among organizational representations remain intact as the organization changes.

AI-native organizations intensify this problem because artificial intelligence increases the rate at which organizational representations can be generated and recombined. AI systems can produce analyses, plans, summaries, code, forecasts, risk assessments, customer responses, and operational recommendations at a scale that exceeds traditional human documentation practices. These outputs can be useful, but they also increase the burden of maintaining coherence. If AI-generated artifacts are not meaningfully related to strategic intent, prior decisions, organizational memory, human judgment, execution outcomes, and governance constraints, the organization may accumulate representations faster than it can preserve relationships among them. The result is not necessarily a lack of information. It may be an abundance of poorly connected representations.

Consider a simple example that will recur throughout the paper. An AI-native logistics organization adopts a strategic commitment to prioritize delivery reliability over route speed. Later, an AI-generated routing recommendation optimizes for short-term efficiency while remaining weakly connected to the updated strategic commitment and the decision record that approved it. Teams act on the recommendation, service outcomes shift, and complaints reveal that reliability has degraded. The organization has not lost all knowledge; rather, the relationships among strategic intent, AI-generated recommendation, operational action, outcome evidence, and organizational learning have weakened.

This problem is partly semantic. Organizational representations must retain meaningful relationships of interpretation, classification, reference, and implication. When semantic relationships degrade, different parts of the organization may use the same terms differently, interpret priorities inconsistently, or attach incompatible meanings to shared artifacts. Strategic language may be repeated while its practical meaning shifts. A decision record may persist while its rationale becomes detached from later action. An AI-generated summary

may appear coherent while omitting distinctions that matter for organizational interpretation. Semantic inconsistency can therefore accumulate quietly, producing apparent agreement at the surface while weakening shared organizational understanding.

The problem is also causal. Organizations must preserve relationships among intentions, decisions, actions, outcomes, and learning. When causal relationships degrade, it becomes difficult to trace why a decision was made, which action followed from it, what outcome resulted, and how the organization should learn from that outcome. Execution may drift from intent without any single actor intending the drift. Accountability may weaken because representations no longer show how choices, actions, and consequences are connected. Organizational memory may retain records while losing the causal structure that makes those records useful. In AI-native settings, causal fragmentation may be amplified when machine-generated outputs influence action without clear traceability to assumptions, prompts, evidence, or prior organizational commitments.

The problem formulation draws on organizational learning, distributed cognition, cybernetic and systems theory, sensemaking, organizational memory, and knowledge management (Alavi and Leidner, 2001; Argyris and Schon, 1978; Ashby, 1956; Hutchins, 1995; Nonaka, 1994; Walsh and Ungson, 1991; Weick, 1995).

Existing organizational theories provide important foundations for studying this terrain, including information processing, organizational learning, memory, knowledge management, dynamic capabilities, and enterprise architecture, but they do not fully isolate this problem (Alavi and Leidner, 2001; Galbraith, 1974; Ross et al., 2006; Teece et al., 1997; Walsh and Ungson, 1991). Organizational learning explains how organizations acquire, retain, and modify knowledge through experience. Distributed cognition explains how cognition is spread across people, artifacts, and environments. Cybernetic and systems perspectives explain feedback, regulation, and dynamic system behavior. Sensemaking explains how actors interpret ambiguous cues. Organizational memory and knowledge management explain how knowledge is stored, retrieved, shared, and used. These traditions remain essential. Yet the central issue here is more specific: how organizations preserve meaningful semantic and causal relationships across continuously evolving organizational representations.

The distinction matters because organizations may appear capable in adjacent terms while still suffering from representational incoherence. An organization may learn rapidly but fail to connect learning back to the decisions that require revision. It may maintain extensive memory but lose the relationships that make memory actionable. It may exhibit feedback processes but regulate behavior without preserving the representational continuity needed to understand why behavior changed. It may make decisions quickly while weakening traceability between intent, action, and outcome. It may deploy advanced AI while increasing the speed at which disconnected representations enter organizational cognition. None of these failures is reducible to poor communication, weak knowledge management, inadequate documentation, or ineffective decision-making alone.

The theoretical challenge, then, is to explain coherence at the level of representations and their relationships. The relevant object is not an individual artifact, a single decision, a local workflow, or a software system. Nor is it simply the stock of organizational knowledge. The relevant object is the evolving set of representations through which organizational cognition becomes durable, and the semantic and causal relationships that connect those representations over time. When those relationships are preserved, organizational cognition can retain continuity across change. When they degrade, the organization may experience

drift, ambiguity, fragmentation, and loss of traceability even while continuing to generate more information.

This paper therefore begins from a different analytical premise. To understand the problem created by AI-native organizational cognition, it is necessary to study not only actors, processes, technologies, or knowledge assets, but the representational structure through which organizational cognition is expressed and maintained. The next step is ontological. Before defining the constructs that explain coherence or degradation, the paper must specify the substrate in which the problem arises: the evolving network of organizational representations and the semantic and causal relationships among them. This motivates the need for a representation-network ontology.

3. Literature Foundations

The problem developed in the preceding section sits near several major traditions in organizational and systems theory. Organizations have long been studied as learning systems, information-processing systems, decision systems, memory systems, adaptive systems, and distributed cognitive systems. Each tradition provides important resources for understanding how organizations sense, interpret, decide, act, remember, and adapt. The purpose of this section is not to displace those traditions. Rather, it is to clarify the specific explanatory gap that motivates Correspondence Theory: existing theories do not explicitly explain how organizations preserve meaningful semantic and causal relationships across continuously evolving networks of organizational representations.

This positioning is important for two reasons. First, Correspondence Theory depends on insights that adjacent theories have already established: cognition is distributed, learning unfolds over time, organizations regulate behavior through feedback, memory persists in artifacts and routines, and actors construct meaning under ambiguity. Second, the existence of these insights does not by itself explain the representational problem introduced in this paper. A theory can explain how organizations learn without explaining whether learning remains connected to prior decisions. A theory can explain feedback without explaining whether the representation of the feedback remains tied to intent, action, and memory. A theory can explain sensemaking without explaining whether the resulting interpretation is preserved across the representation network. The review below therefore treats existing theories as necessary foundations while identifying the residual explanatory gap they leave open.

Cybernetics and systems theory provide the foundations for feedback, regulation, interdependence, and system behavior (Ashby, 1956; Bertalanffy, 1968).

Cybernetics provides one of the deepest foundations for understanding organizational regulation. Cybernetic accounts explain how systems compare current states with reference conditions, detect deviation, and adjust behavior through feedback. This tradition is essential for explaining control, correction, error detection, and adaptive regulation. Yet cybernetic feedback primarily concerns the regulation of system behavior. It can explain how an organization corrects a process, responds to a signal, or adjusts behavior in relation to a goal. It does not, by itself, explain whether the representations through which goals, decisions, actions, outcomes, and memory are expressed remain semantically and causally connected over time. A system may regulate behavior while the organizational representations that explain why behavior is changing become fragmented or detached from prior intent.

This distinction matters because regulation and correspondence are not equivalent. A feedback mechanism may detect variance from a target and correct action, but the correction may not be represented in a way that preserves its relationship to strategic assumptions, decision history, or later learning. The organization may know that behavior changed without preserving why the change occurred, what representation was revised, or which causal assumptions were confirmed or weakened. Cybernetics therefore helps explain correction, but it does not fully account for the representational continuity through which correction becomes organizationally intelligible over time.

Systems theory and systems dynamics similarly provide powerful accounts of interdependence, feedback, delays, accumulations, and dynamic behavior. These perspectives help explain how organizational outcomes arise from interacting structures rather than isolated events. They are especially useful for modeling how variables affect one another over time, how feedback loops amplify or dampen behavior, and how unintended consequences emerge from complex systems. However, the primary objects of these theories are system structures, stocks, flows, feedback relations, and behavioral patterns. They do not directly explain the maintenance of semantic and causal relationships among organizational representations as such. A systems model may describe causal behavior, but the organization may still fail to preserve the representational continuity linking intent, decision, action, outcome, and learning.

The issue is not that systems theories ignore causality. On the contrary, they are among the strongest traditions for explaining causal interdependence. The difference concerns the level of analysis. Systems dynamics may represent causal structure in a model, while Correspondence Theory asks whether the organization's own representations of causality remain meaningfully connected across work. The model may describe a system; the organization may still lose coherence in the representations by which it remembers, communicates, governs, and revises its understanding of that system.

Distributed cognition and boundary-object scholarship establish why cognition can be analyzed across actors, artifacts, tools, and environments (Hutchins, 1995; Star and Griesemer, 1989).

Distributed cognition offers another important foundation. It challenges individualistic accounts of cognition by showing how cognitive activity is distributed across people, artifacts, tools, routines, and environments. This is highly relevant for contemporary organizations, where cognition is rarely contained within individual minds. Organizational work depends on documents, models, decision records, procedures, analytic outputs, governance artifacts, and increasingly AI-generated artifacts. Distributed cognition therefore helps establish why artifacts and environments matter for organizational cognition. Yet its central contribution is to explain the distribution of cognition, not the recursive maintenance of representational coherence. It can show that cognition is spread across actors and artifacts, but it does not fully specify how the resulting representations remain meaningfully and causally connected as they evolve.

In AI-native organizations, this distinction becomes sharper as algorithmic systems alter work, control, and managerial decision-making without eliminating the need for accountable human organizational cognition (Faraj et al., 2018; Kellogg et al., 2020; Raisch and Krakowski, 2021). Cognition may be distributed not only across humans and artifacts, but also across machine-generated outputs, automated recommendations, and persistent organizational memory. Distributed cognition helps make this distribution visible. Correspondence Theory addresses a further question: once cognition is distributed across

such representations, what preserves the relationships among them? Distribution explains where cognition occurs. It does not fully explain how the semantic and causal structure of represented cognition is maintained.

Organizational learning scholarship explains error correction, learning from experience, and exploration-exploitation dynamics (Argyris and Schon, 1978; March, 1991).

Organizational learning explains how organizations acquire, retain, interpret, and revise knowledge through experience. Learning theories help explain how experience becomes knowledge, how routines change, how errors are corrected, and how organizations adapt in response to feedback. These accounts are central to any theory of organizational cognition. Still, learning can occur without preserving the relationships that make learning coherent across the organization. A lesson may be recorded without being linked to the decision that produced it. An outcome may be interpreted without revising the assumptions that shaped the original action. A new practice may emerge without updating the strategic or policy representations it affects. Organizational learning explains knowledge change; it does not fully explain whether that change remains connected within the broader representation network.

This limitation is especially visible when learning is local. A team may learn from an execution outcome and alter its behavior, while the organization's strategy, governance records, decision rationale, and memory remain unchanged. From one perspective, the organization has learned. From another, its representation network has become less coherent because the learning is not related to the representations it should modify. Correspondence Theory therefore treats learning events as potentially important representation-network events, but it does not equate learning with correspondence.

Adaptive systems perspectives also contribute significantly. They explain how organizations and other systems adjust to environmental pressure, variation, uncertainty, and internal complexity. Adaptation may involve behavioral change, structural change, learning, experimentation, or reconfiguration. These accounts are particularly useful for understanding organizations under uncertainty. However, adaptation can occur at the level of behavior while representational relationships degrade. An organization may adapt quickly to market or technological change while losing traceability between its new actions and earlier strategic commitments. It may reorganize in response to pressure while leaving policies, decisions, plans, and memory in inconsistent states. Adaptive systems theory explains adaptation; it does not fully explain representational maintenance during adaptation.

The distinction is not merely semantic. Organizations often need to adapt while preserving enough continuity to remain intelligible to themselves. They must know which assumptions were abandoned, which commitments remain, which outcomes motivated the change, and which future decisions should inherit the revised understanding. Without this representational continuity, adaptation can produce local responsiveness but organization-level fragmentation.

Sensemaking scholarship explains how actors construct meaning under ambiguity and enact plausible accounts of organizational situations (Weick, 1995).

Sensemaking theories explain how actors interpret ambiguous cues, construct meaning, and enact plausible accounts of organizational situations. Sensemaking is crucial because organizations do not merely process information; they interpret events, negotiate meaning, and construct shared understandings. This tradition helps explain ambiguity, interpretation,

identity, narrative, and collective meaning-making. Yet sensemaking primarily concerns the construction and revision of meaning. It does not fully account for the ongoing maintenance of semantic and causal relationships among persistent organizational representations after interpretation occurs. An organization may make sense of an event locally while failing to update the relationships among strategies, decisions, actions, outcomes, and memory. The residual problem is not only meaning construction, but representational coherence over time.

Sensemaking also tends to foreground interpretive activity. Correspondence Theory accepts the importance of interpretation but asks what happens to the representations after interpretation is stabilized, revised, or contested. Does the new interpretation become connected to the relevant decision records? Does it revise prior assumptions? Does it alter how outcomes are understood? Does it remain linked to the causal chain that produced the event? These questions concern the maintenance of a representation network rather than interpretation alone.

Organizational and transactive memory scholarship explains retention, retrieval, distributed expertise, and remembered organizational experience (Walsh and Ungson, 1991; Wegner, 1987).

Organizational memory addresses how knowledge, experience, routines, and records persist beyond individual actors. It explains retention, retrieval, and the organizational consequences of remembered experience. This tradition is highly relevant because Correspondence Theory also concerns persistent organizational representations. However, memory persistence is not the same as correspondence among representations. An organization may retain decision histories, lessons, policies, and outcome records while the relationships among them become incomplete, outdated, contradictory, or invisible. Memory can persist while causal traceability collapses. Correspondence Theory addresses the relationship integrity of memory representations, not merely their retention.

Knowledge-management scholarship explains knowledge creation, codification, storage, retrieval, and reuse as organizational knowledge practices (Alavi and Leidner, 2001; Nonaka, 1994).

Knowledge management explains how knowledge is captured, stored, shared, retrieved, and reused. It focuses on knowledge assets, repositories, transfer practices, codification, and access. These are important concerns, but knowledge availability is not equivalent to Organizational Correspondence. Knowledge can be available but disconnected. A repository may contain strategy documents, lessons learned, and AI-generated recommendations without preserving how those representations relate to one another semantically or causally. Correspondence Theory therefore addresses a different problem: whether represented knowledge remains meaningfully connected within organizational cognition.

This is a decisive distinction for organizations that are rich in stored information. Accumulated knowledge may create the appearance of cognitive maturity while increasing the burden of relationship maintenance. The more representations an organization retains, the more important it becomes to preserve relationships among them. In this sense, organizational memory can become both a resource and a site of degradation: it preserves representations, but it does not automatically preserve correspondence among them.

Organizational information-processing theory explains how uncertainty, equivocality, and information requirements shape organizational design and communication capacity (Daft and Lengel, 1986; Galbraith, 1974; March and Simon, 1958).

Organizational information processing explains how organizations gather, process, interpret, and distribute information under uncertainty and complexity. It is especially useful for understanding information needs, processing capacity, and the fit between uncertainty and organizational arrangements. Correspondence Theory does not replace this view. It asks a different question. Even when information is processed efficiently, the relationships among resulting representations may degrade. An organization may process a new market signal, generate a decision, and communicate an action while failing to preserve how that signal relates to prior intent, assumptions, outcomes, and learning.

Sociomaterial and boundary-object perspectives support the paper's treatment of artifacts and technologies as entangled with organizing while remaining distinct from the representation-network property theorized here (Orlikowski, 2007; Star and Griesemer, 1989).

Sociomaterial approaches explain how social and material arrangements are entangled in organizational practice. They help show that technologies, artifacts, routines, and actors are not separable in actual organizing. Correspondence Theory is compatible with this insight but focuses on a narrower theoretical target. It does not primarily explain sociomaterial entanglement. It explains whether the representations produced, modified, and stabilized within such entanglements preserve meaningful semantic and causal relationships over time.

Process improvement and decision-cycle models, including PDCA and OODA, add another layer of comparison. PDCA explains iterative improvement through planning, doing, checking, and acting. OODA explains decision and action under changing conditions through observation, orientation, decision, and action. Both are valuable for understanding organizational action under uncertainty and correction over time. But neither is primarily a theory of representation-network maintenance. PDCA can improve a process without ensuring that process changes remain connected to strategic intent, causal assumptions, or organizational memory. OODA can accelerate decision cycles while leaving decision traces, action outcomes, and learning artifacts disconnected. These cycles explain improvement or decision tempo; they do not fully explain the preservation of semantic and causal relationships across evolving organizational representations.

PDCA and OODA also illustrate why the Correspondence Loop should not be understood as another execution loop or decision cycle. The problem addressed in this paper is not how organizations plan, act, check, decide, or orient more effectively. It is how the representations produced by such activities remain connected to the wider network of organizational cognition. A plan, action, check, observation, orientation, decision, or lesson may all become representations. The theoretical question is whether their semantic and causal relationships are preserved as the organization continues to change.

Table 2 - Existing Theory Comparison

Table 2 - Existing Theory Comparison

Table 2 - Existing Theory Comparison

Section 3 - Literature Foundations

Theory	Primary unit of analysis	Primary explanatory focus	Mechanism	Limitation	Residual gap addressed by Correspondence Theory
Organizational learning	Experience, knowledge, routines	Knowledge acquisition and behavioral change	Learning from experience, feedback, reflection, and adaptation	Learning may remain local or disconnected from prior representations	Whether learning remains semantically and causally connected to intent, decision, action, outcome, and memory
Distributed cognition	People, artifacts, tools, environments	Distribution of cognition across actors and artifacts	Cognitive activity distributed across social and material arrangements	Distribution does not explain maintenance of relationships among representations	How distributed representations preserve meaningful semantic and causal relationships over time
Organizational memory	Retained knowledge, records, routines	Persistence and retrieval of organizational knowledge	Retention, storage, retrieval, and reuse	Memory may persist while relationships among memories degrade	Whether memory representations remain connected to decisions, outcomes, and learning
Knowledge management	Knowledge assets and repositories	Capture, storage, sharing, retrieval, and application of knowledge	Codification, transfer, repository use, sharing practices	Available knowledge may be fragmented or disconnected	Whether represented knowledge retains relationship integrity within the representation network
Sensemaking	Cues, interpretation, narratives	Meaning construction under ambiguity	Interpretation, enactment, narrative construction	Local meaning may not revise or connect persistent representations	Whether interpretations become integrated into the evolving representation network
Organizational information processing	Information needs, processing capacity, uncertainty	Information processing under uncertainty	Gathering, processing, distributing, and interpreting information	Efficient processing does not ensure representational continuity	Whether processed information remains connected to prior and future organizational representations
Focal contribution					
Correspondence Theory	Organizational representations and relationships	Preservation of semantic and causal coherence across the representation network	Correspondence Loop: Representation Update, Relationship Update, Correspondence Evaluation, Correspondence Revision	Does not explain all organizational performance or all cognition	Provides the focal account of representation-network correspondence

Table 2 compares adjacent theoretical traditions with the residual representation-network gap addressed by Correspondence Theory.

The comparison is conceptual and complementary; it does not claim replacement of adjacent theories.

Table 2 compares adjacent theoretical traditions with the residual representation-network gap addressed by Correspondence Theory. The comparison is conceptual and complementary; it does not claim replacement of adjacent theories.

Across these traditions, the recurring pattern is complementarity rather than replacement. Existing theories explain regulation, system behavior, distributed cognition, learning, adaptation, meaning construction, memory persistence, knowledge flow, process improvement, and decision cycles. Correspondence Theory builds from these foundations but asks a narrower and underdeveloped question: how does an organization preserve meaningful semantic and causal relationships across the representations through which organizational cognition is expressed, remembered, revised, and acted upon?

This question becomes especially important in AI-native organizations. AI systems can generate, transform, summarize, classify, recommend, and act upon representations at high speed and scale. Existing theories can explain parts of this phenomenon: AI may accelerate learning, distribute cognition across human and machine systems, alter feedback dynamics, support sensemaking, or expand organizational memory. Yet the representational problem remains. As the volume and speed of representations increase, organizations must preserve relationships among intent, decision, action, outcome, memory, policy, and machine-generated artifacts. Without such preservation, cognition may become faster but less coherent.

The residual explanatory gap is therefore specific. Existing organizational and systems theories explain cognition, learning, regulation, memory, adaptation, decision-making, interpretation, and knowledge flow, but none explicitly explain the recursive maintenance of the evolving organizational representation network from which Organizational Correspondence emerges. This gap is ontological before it is mechanistic: the paper must first specify the substrate in which correspondence can be preserved or lost. This motivates the representation-network ontology developed in the next section and the constructs and mechanism that follow.

4. Representation-Network Ontology

The preceding sections identify a gap in how organizational theory explains coherence across evolving organizational representations. To address that gap, the paper first requires an ontology: an account of the basic entities and relationships through which the phenomenon should be studied. Correspondence Theory begins from the premise that organizational cognition can be examined through the representations organizations create, maintain, transform, relate, evaluate, revise, and discard. These representations are not incidental by-products of organizing. They are among the durable media through which organizational cognition persists across actors, systems, decisions, and time.

The first element of the ontology is the organizational representation. Organizational representations are persistent artifacts, statements, models, records, memories, outputs, or structured expressions through which organizational cognition is represented. They may include strategic intent, decisions, plans, work products, observations, outcomes, policies, governance records, AI-generated artifacts, and organizational memory. The category is intentionally broad because organizations do not express cognition through a single medium. A strategy deck, product roadmap, risk register, architectural model, operating policy, decision log, customer analysis, and machine-generated recommendation may all function as organizational representations when they express something the organization knows, intends, decides, observes, or remembers.

The concept of organizational representation does not imply that all organizational cognition is explicit, a boundary consistent with distributed cognition, sensemaking, and

sociomaterial perspectives on situated cognition and artifact-mediated organizing (Hutchins, 1995; Orlikowski, 2007; Weick, 1995). Organizations also rely on tacit knowledge, embodied practice, informal judgment, and situated interpretation. Correspondence Theory does not claim to capture the whole of organizational cognition. Its focus is the portion of organizational cognition that becomes represented, related, revised, remembered, or acted upon. Tacit or hidden cognition matters when it affects the representation network. This boundary is important because the theory studies represented organizational cognition without reducing organizations to documents or databases.

This boundary also protects the ontology from becoming too broad. If every thought, habit, feeling, informal exchange, or practice were treated as a representation, the concept would lose analytical precision. The theory is concerned with organizational cognition once it takes a durable or structured form that can be related to other organizational expressions. A conversation may matter because it produces a decision record, revises a plan, changes an interpretation, or alters how an outcome is remembered. In such cases, the theoretical focus falls on the represented consequence and its relationships within the network.

The second element is the representation relationship. Organizational representations become theoretically significant not only because they exist, but because they stand in meaningful semantic and causal relationships to one another. Semantic relationships concern meaning, reference, interpretation, classification, implication, or consistency among representations. For example, a decision may interpret a strategic intent, a policy may define the meaning of a governance commitment, or a project plan may instantiate a broader priority. Causal relationships concern dependency, influence, consequence, traceability, or explanation. For example, a decision may lead to an action, an action may produce an outcome, an outcome may trigger learning, and learning may revise a future decision.

These relationships are central because organizational coherence depends on more than the presence of representations. An organization may have many representations and still lose coherence if relationships among them are missing, outdated, contradictory, or invisible. A plan disconnected from intent, an outcome disconnected from action, a lesson disconnected from decision history, or an AI-generated artifact disconnected from governance constraints can all weaken organizational cognition. The ontology therefore treats relationships as foundational rather than secondary.

The ontology distinguishes semantic and causal relationships because organizations require both. Semantic relationships preserve meaning among representations: whether a policy expresses a principle, whether a plan interprets a strategy, whether a term is used consistently, or whether an artifact retains its intended significance. Causal relationships preserve traceability among representations: whether a decision explains an action, whether an action produced an outcome, whether an outcome generated learning, or whether learning revised future intent. Organizational coherence depends on the interaction of these two relationship types.

The third element is the representation network. The representation network is the evolving network of organizational representations and their semantic and causal relationships. It is not a software graph, database graph, documentation system, or implementation model. It is an ontological substrate: the structure through which represented organizational cognition can be studied. The network evolves as organizations create new representations, retire old ones, revise existing ones, reinterpret prior meanings, add or remove relationships, and learn from outcomes. Because organizations change continuously, the representation network is dynamic rather than static.

The representation network therefore should not be understood as a complete map that an organization already possesses. It is an analytical object. Some relationships may be explicit, some implicit, some contested, and some missing. Some representations may be formal and durable, while others may be provisional or weakly institutionalized. The theory does not require perfect visibility into the network. It requires only that organizations have persistent representations and meaningful relationships among them that can be preserved, degraded, revised, or inferred.

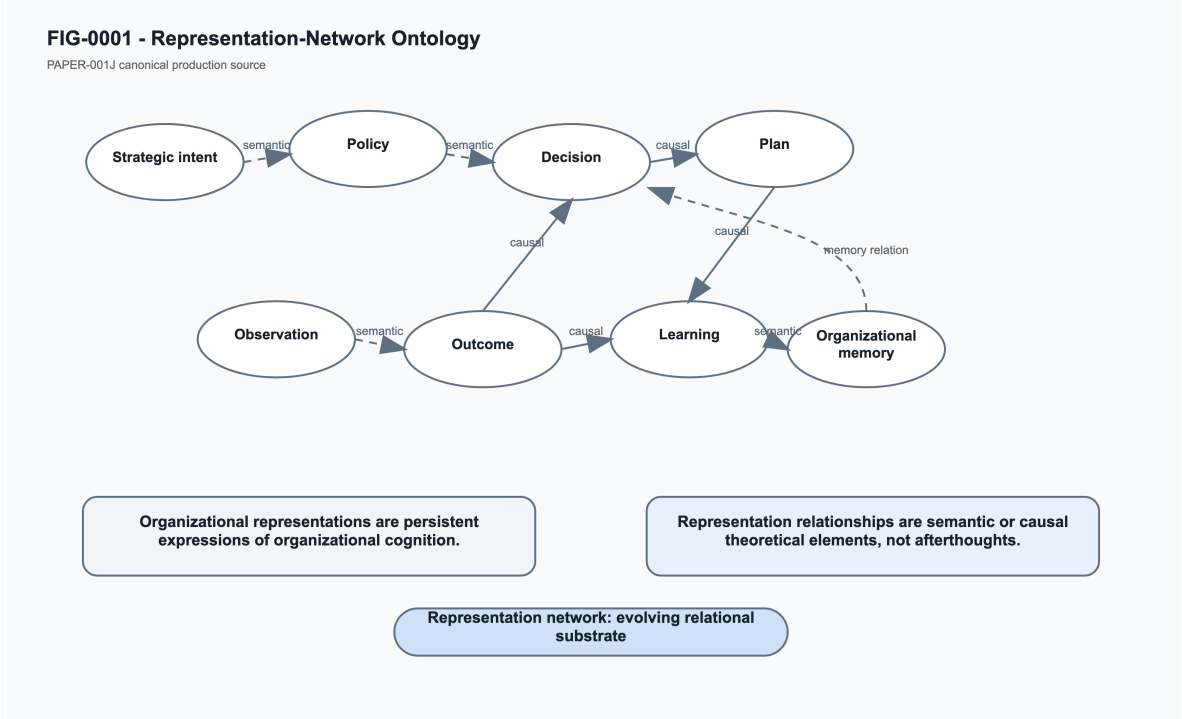

Figure 1 - Representation-Network Ontology. The representation-network ontology treats persistent organizational representations and their semantic and causal relationships as the relational substrate of Correspondence Theory.

This ontology is representation-network-centric. It does not deny the importance of actors, teams, processes, technologies, governance practices, or AI systems. Rather, it locates their theoretical relevance in relation to the representation network. Actors, teams, processes, documents, systems, governance mechanisms, and AI agents matter because they create, modify, interpret, evaluate, or revise organizational representations and relationships. They are not the primary substrate of the theory. The primary substrate is the evolving network of representations and relationships through which organizational cognition becomes durable enough to be preserved, degraded, restored, or studied.

This distinction is especially important for AI-native organizations. AI systems can generate, modify, summarize, recommend, classify, and act upon organizational representations at high speed and scale. Yet AI does not define the ontology. AI intensifies the need for it. As machine-generated representations enter organizational cognition, the organization must preserve their semantic and causal relationships to prior intent, decisions, policies, evidence, outcomes, and memory. Without an ontology centered on representation networks, the theoretical problem risks being misread as a problem of tooling, automation, communication, or knowledge storage alone.

The ontology also clarifies the sequence of the theory. Organizational representations form the most basic layer. Representation relationships connect them. The representation

network emerges from those connected representations. Organizational Correspondence, introduced in the next section, is an emergent property of that network: the preservation of meaningful semantic and causal relationships among organizational representations over time. Cognitive Integrity is then defined as the dynamic organizational condition associated with sustained Organizational Correspondence. The Correspondence Loop later explains how the representation network is recursively maintained.

The representation-network ontology therefore performs three functions in the paper. First, it defines the object of analysis. The paper is not primarily concerned with isolated documents, individual cognition, local workflows, or software artifacts, but with the evolving representational structure of organizational cognition. Second, it establishes the substrate from which the paper's constructs emerge. Organizational Correspondence and Cognitive Integrity are not free-standing concepts; they depend on the representation network. Third, it preserves implementation independence. The theory can apply to human-only, hybrid, and AI-native organizations because the ontology concerns represented organizational cognition rather than any specific technology.

By grounding the theory in organizational representations, representation relationships, and representation networks, the paper can now define the core property at stake. The central question is not merely whether organizations possess representations, knowledge, or memory. It is whether meaningful semantic and causal relationships among those representations are preserved as organizational cognition unfolds over time. That property is Organizational Correspondence.

5. Organizational Correspondence

The representation-network ontology establishes the substrate of the theory: organizations create, maintain, transform, and discard representations, and those representations are connected through semantic and causal relationships. The first formal construct of Correspondence Theory identifies the property that becomes possible within that substrate. Organizational Correspondence is the foundational construct of the theory because it names the condition of relationship preservation within the evolving representation network.

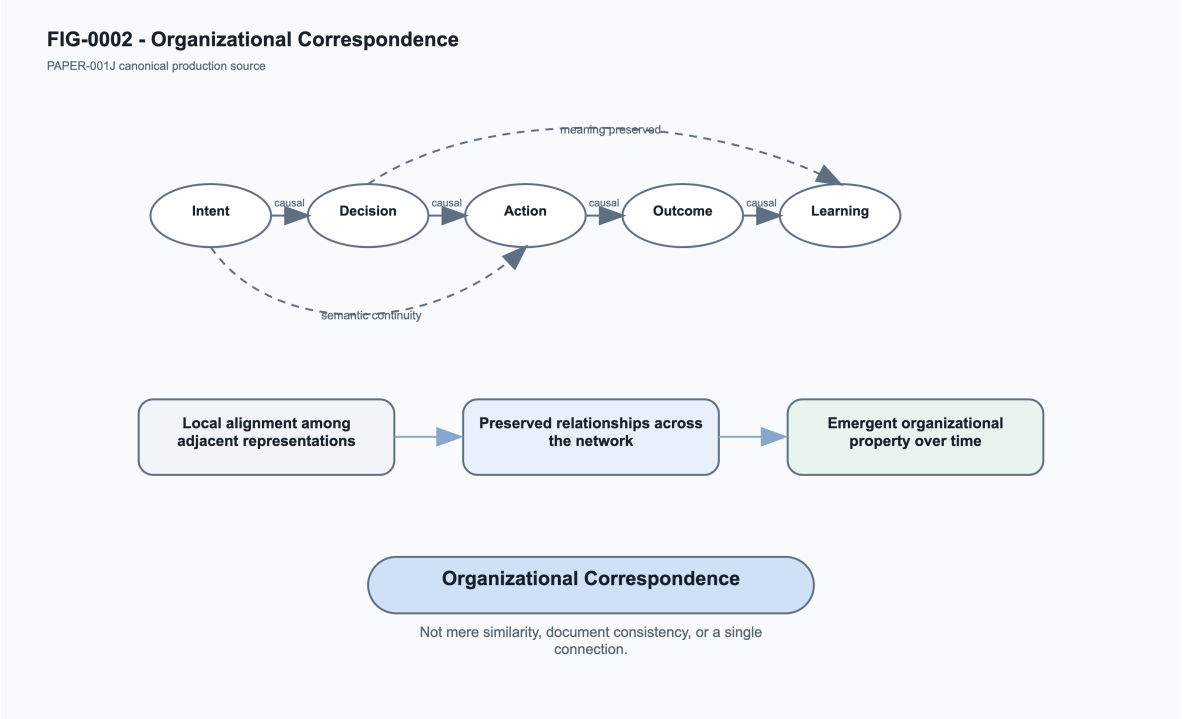

Figure 2 - Organizational Correspondence. Organizational Correspondence is the emergent organizational property by which meaningful semantic and causal relationships among organizational representations are preserved over time.

Organizational Correspondence is defined as follows: Organizational Correspondence is the emergent organizational property by which meaningful semantic and causal relationships among organizational representations are preserved over time. This definition is intentionally minimal. It specifies the construct's ontological character without embedding mechanisms, management practices, technologies, diagnostic criteria, or performance outcomes. The construct does not describe a process that organizations perform. It describes a property that emerges when the representation network preserves meaningful relationships across time.

Table 1 - Construct Definitions

Table 1 - Construct Definitions

Table 1 - Construct Definitions

Section 5 - Organizational Correspondence

Construct	Canonical definition	Ontological role	Analytical level	What the construct is not
Representation-network ontology				
Organizational Representation	Persistent artifact, statement, model, record, memory, output, or structured expression through which organizational cognition is represented.	Basic ontological element	Representation-level	Not all cognition, not tacit practice itself, and not limited to documents or databases
Semantic Relationship	Relationship concerning meaning, reference, interpretation, classification, implication, or consistency among organizational representations.	Relationship type within the representation network	Relationship-level	Not causal traceability, communication volume, or agreement by itself
Causal Relationship	Relationship concerning dependency, influence, consequence, traceability, or explanation among organizational representations.	Relationship type within the representation network	Relationship-level	Not correlation alone, performance attribution, or a validated causal model
Representation Network	The evolving network of organizational representations and their semantic and causal relationships.	Ontological substrate	Network-level	Not a software graph, database graph, documentation system, or implementation model
Formal constructs				
Organizational Correspondence	The emergent organizational property by which meaningful semantic and causal relationships among organizational representations are preserved over time.	Formal construct and emergent property	Organization-level property of the representation network	Not a process, capability, governance mechanism, communication framework, information system, or software architecture
Cognitive Integrity	The dynamic organizational condition associated with sustained Organizational Correspondence.	Formal construct and dependent dynamic condition	Organization-level condition	Not a capability, process, mechanism, performance outcome, or governance property
Canonical mechanism				
Correspondence Loop	The Correspondence Loop is a recursive, event-driven network maintenance mechanism that operates on the organizational representation network through Representation Update, Relationship Update, Correspondence Evaluation, and Correspondence Revision to preserve or restore the conditions under which Organizational Correspondence is sustained.	Canonical mechanism	Mechanism operating on the representation network	Not a formal construct, linear workflow, governance model, software system, process prescription, or execution cycle

Table 1 defines the ontological substrate, formal constructs, and canonical mechanism used by Correspondence Theory.

Canonical source: PAPER-001 construct and mechanism files.

Table 1 defines the ontological substrate, formal constructs, and canonical mechanism used by Correspondence Theory. Canonical source: PAPER-001 construct and mechanism files.

Several features of this definition are important. First, Organizational Correspondence is emergent. It does not reside in any single artifact, decision, actor, model, policy, or system. A strategy document may be carefully written, a decision record may be complete, and an outcome report may be accurate, but correspondence depends on whether meaningful relationships among such representations are preserved. The property arises from the network, not from an isolated representation.

Second, Organizational Correspondence is organizational rather than individual. It concerns collective organizational cognition as represented across durable artifacts and relationships. Individual actors may interpret, remember, decide, or revise, but the construct refers to the organization-level property by which represented organizational cognition remains meaningfully and causally connected. This distinction allows the theory to study organizations whose cognition is distributed across people, teams, tools, routines, and AI systems without reducing the phenomenon to individual cognition.

Third, Organizational Correspondence is dynamic. It concerns preservation over time, not static consistency at a single moment. Organizations continuously produce new representations and revise old ones. Strategies change, decisions accumulate, policies evolve, actions generate outcomes, and learning modifies future interpretation. Correspondence is therefore not equivalent to perfect agreement or frozen alignment. It describes whether semantic and causal relationships remain meaningful as the representation network changes.

Fourth, the construct is representation-centric. The primary unit of analysis is organizational representations and the semantic and causal relationships among them. Examples include strategic intent, policies, decisions, plans, work items, observations, outcomes, organizational memory, governance records, AI-generated artifacts, and future forms of organizational representation. The theory remains open to different media and organizational forms because it does not define correspondence by reference to any specific repository, process, system, or document type.

The semantic dimension of Organizational Correspondence concerns meaning. Representations must remain interpretable in relation to one another. A decision should retain a meaningful relationship to the intent it enacts. A policy should preserve the meaning of the principle or constraint it expresses. A plan should remain interpretable in light of the decision, priority, or assumption it implements. When semantic relationships degrade, organizations may continue to use the same words while attaching different meanings to them. Apparent agreement may conceal interpretive fragmentation.

The causal dimension concerns traceability and consequence. Representations must preserve relationships among intentions, decisions, actions, outcomes, and learning. A decision should be causally related to the action it authorizes. An outcome should remain traceable to the action and assumptions that produced it. Learning should revise the relevant representations rather than remain as disconnected memory. When causal relationships degrade, organizations may lose the ability to explain why a course of action occurred, what it produced, and what should change as a result.

In the logistics example introduced earlier, Organizational Correspondence weakens when the strategic commitment to delivery reliability, the AI-generated routing recommendation, the operational actions taken from that recommendation, and the resulting service complaints no longer remain meaningfully related. The problem is not that the organization lacks representations. It has a strategy, a recommendation, actions, outcomes, and memory. The problem is that the semantic and causal relationships among

those representations have degraded.

Organizational Correspondence complements existing organizational theories without replacing them. It is distinct from organizational alignment, organizational cognition, organizational memory, sensemaking, cybernetics, dynamic capabilities, knowledge management, and organizational routines. Those theories explain important phenomena involving actors, knowledge, regulation, interpretation, capabilities, or recurring activity. Organizational Correspondence explains a more specific property: the integrity of meaningful semantic and causal relationships among organizational representations over time.

The distinction from alignment is especially important. Alignment often refers to consistency among goals, structures, incentives, resources, or activities. Organizational Correspondence is narrower and more representational. It does not ask whether all parts of the organization are aligned in a managerial or strategic sense. It asks whether the semantic and causal relationships among organizational representations are preserved. An organization may be strategically aligned in broad terms while losing correspondence among the representations that explain how intent, decisions, actions, outcomes, and learning relate to one another.

The distinction from communication is equally important. Communication may transmit information among actors, but transmission does not guarantee preservation of relationships across the representation network. A decision may be communicated widely while its rationale remains disconnected from prior assumptions. A policy may be announced clearly while its relationship to existing practices remains ambiguous. A lesson may be shared without revising the representations it should affect. Organizational Correspondence concerns the integrity of relationships among represented organizational cognition, not the mere movement of messages.

This distinction matters because organizations may appear strong in adjacent terms while still losing correspondence. They may possess extensive documentation but weak relationship integrity. They may learn quickly but fail to connect learning to prior decisions. They may make decisions rapidly while weakening traceability between intent, action, and outcome. They may use sophisticated AI systems while increasing the volume of representations faster than the organization can preserve relationships among them. Organizational Correspondence names the property at risk in these cases.

The construct applies wherever persistent organizational representations and meaningful relationships among them exist. Artificial intelligence may change the scale and speed of the representational problem, but it does not define the construct. Organizational Correspondence is not directly observable. Researchers infer it through traces of organizational representations and the evolving relationships among them. This does not make the construct speculative; it places it in the same methodological family as many organization-level theoretical constructs that are studied through observable evidence. Future empirical work may investigate dimensions such as semantic continuity, causal traceability, temporal continuity, representation integrity, and learning continuity. In this paper, however, the construct is introduced conceptually rather than operationalized.

The temporal dimension is also essential. Correspondence is preserved or degraded as organizational cognition unfolds. A representation may be coherent when created and incoherent later if the relationships around it change. A decision may accurately reflect intent at one moment but become detached as conditions, actions, or interpretations shift. An

outcome may initially be linked to an action but later be remembered without its causal context. Organizational Correspondence therefore concerns continuity across change, not merely consistency at the moment of production.

The theoretical role of Organizational Correspondence is foundational. It explains whether the meaningful relationships among organizational representations remain intact as organizational cognition unfolds. It also provides the parent construct for Cognitive Integrity, introduced in the next section, and the property whose conditions are maintained by the Correspondence Loop. Without Organizational Correspondence, the theory would lack its root explanatory object. The next step is to explain the organizational condition associated with sustained correspondence over time.

6. Cognitive Integrity

Organizational Correspondence identifies the emergent property by which meaningful semantic and causal relationships among organizational representations are preserved over time. The second foundational construct of Correspondence Theory names the organizational condition associated with that sustained property. Cognitive Integrity is defined as follows: “Cognitive Integrity is the dynamic organizational condition associated with sustained Organizational Correspondence.”

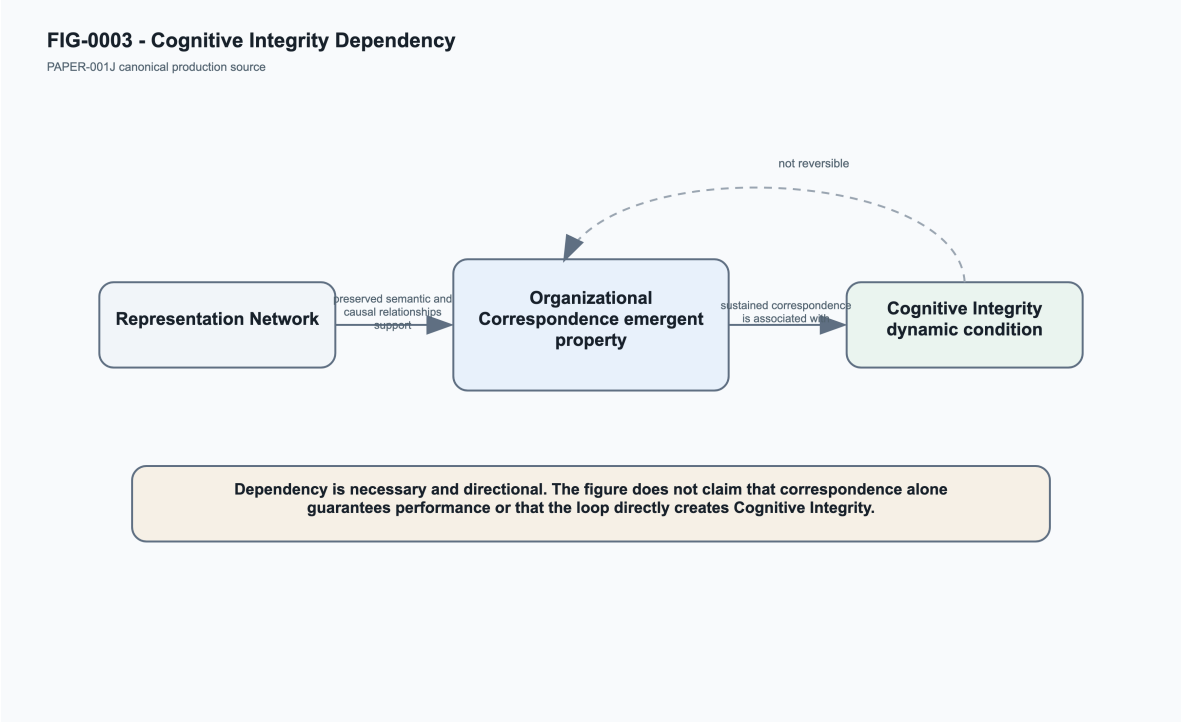

Figure 3 - Dependency Between Organizational Correspondence and Cognitive Integrity. Cognitive Integrity is the dynamic organizational condition associated with sustained Organizational Correspondence.

This definition is intentionally concise. It identifies Cognitive Integrity as a condition rather than a mechanism, process, capability, governance practice, performance outcome, or software feature. It also preserves the dependency between the two foundational constructs. Organizational Correspondence is the parent construct. Cognitive Integrity depends on sustained Organizational Correspondence and should not be defined independently of it.

The dependency sequence is:

Representation Network

□

Organizational Correspondence

□

Cognitive Integrity

The representation network is the substrate. Organizational Correspondence is the emergent property of that network when meaningful semantic and causal relationships among organizational representations are preserved over time. Cognitive Integrity is the dynamic organizational condition associated with that sustained property. The dependency is not reversible. Organizational Correspondence can be defined without Cognitive Integrity; Cognitive Integrity cannot be defined without Organizational Correspondence.

Cognitive Integrity is an organizational-level construct. It does not describe the mental integrity of an individual actor, nor does it refer to individual rationality, attention, memory, or expertise. It describes the condition of organizational cognition as represented across the organization. This matters because organizations may rely on many competent individuals while still suffering from incoherent organizational cognition. If decisions, policies, plans, outcomes, and memories no longer preserve meaningful semantic and causal relationships, the organization-level condition may degrade even when individual actors remain capable.

Cognitive Integrity is also dynamic. It concerns the condition of organizational cognition as correspondence is sustained, weakened, restored, or degraded over time. The construct therefore differs from static coherence at a single point. An organization may appear coherent in one moment but lose integrity as representations change, decisions accumulate, AI-generated artifacts enter the network, or learning fails to revise relevant relationships. Conversely, an organization may experience local ambiguity or contradiction while preserving broader integrity if those issues are recognized, related, and revised within the representation network.

The construct is emergent because it depends on the sustained preservation of relationships within the representation network. It cannot be reduced to any single document, decision, artifact, actor, system, or governance rule. A well-maintained knowledge repository alone does not establish Cognitive Integrity. Nor does a sophisticated decision process, a strong leadership team, or an advanced AI capability. Cognitive Integrity is associated with the organization-level condition created when the representation network continues to preserve meaningful semantic and causal relationships over time.

This distinction helps explain why organizations can possess many apparently valuable resources while still experiencing cognitive degradation. They may have abundant documentation but weak causal traceability. They may have formal governance but inconsistent interpretation of decisions. They may have advanced analytics but poor connection between analysis, action, outcome, and learning. They may have rapid AI-generated output but insufficient preservation of relationships among that output and the organization's prior representations. In such cases, the issue is not simply absence of information. It is degradation in the condition of organizational cognition associated with weakened Organizational Correspondence.

In the logistics example, Cognitive Integrity declines when the organization can no longer treat the reliability strategy, routing recommendation, operational action, customer complaints, and learning record as a coherent set of related representations. The organization may still possess all of these artifacts, and individual actors may still understand

their local tasks. The organization-level condition is weakened because sustained Organizational Correspondence has been interrupted.

Cognitive Integrity should not be confused with performance, dynamic capabilities, strategic alignment, or enterprise-architecture maturity (Ross et al., 2006; Teece et al., 1997; Venkatraman, 1989). High Cognitive Integrity does not guarantee market success, competitive advantage, effective strategy, or operational excellence. Organizations may maintain coherent representations and still choose poorly, face adverse markets, lack resources, or encounter external shocks. Conversely, organizations may perform well for a period despite degraded correspondence. The construct's theoretical purpose is narrower: it describes the condition of organizational cognition associated with sustained Organizational Correspondence. Any relationship to performance remains downstream and requires separate empirical study.

The relationship between Cognitive Integrity and the Correspondence Loop must also remain precise. The Correspondence Loop does not directly create Cognitive Integrity. It recursively maintains the organizational representation network. That maintenance preserves or restores the conditions from which Organizational Correspondence emerges. Sustained Organizational Correspondence is associated with Cognitive Integrity. This separation keeps the theory architecturally clean: the representation network is the substrate, Organizational Correspondence is the emergent property, Cognitive Integrity is the dynamic condition, and the Correspondence Loop is the mechanism that maintains the representation-network conditions under which the constructs are sustained.

Cognitive Integrity is independent of any specific technical or organizational form. AI may increase an organization's capacity to generate representations, but a larger volume of representations does not itself imply higher Cognitive Integrity. The relevant question is whether those representations remain connected to intent, decisions, policies, outcomes, and memory. Cognitive Integrity is not directly observed. It is inferred from evidence concerning representation-network conditions, Organizational Correspondence, and the activities that preserve or degrade correspondence over time. This preserves the distinction between theory and later research methodology.

Cognitive Integrity therefore gives Correspondence Theory its second foundational layer. Organizational Correspondence explains the emergent property that must be preserved. Cognitive Integrity explains the dynamic organizational condition associated with that sustained property. The remaining theoretical task is to explain how organizations maintain, degrade, revise, or restore the representation-network conditions under which these constructs are sustained. That task belongs to the Correspondence Loop.

7. Correspondence Loop

The preceding sections establish the ontology and constructs of Correspondence Theory. Organizational representations and their semantic and causal relationships form the representation network. Organizational Correspondence is the emergent property by which meaningful relationships among those representations are preserved over time. Cognitive Integrity is the dynamic organizational condition associated with sustained Organizational Correspondence. The remaining theoretical requirement is a mechanism explaining how the representation network is maintained, revised, or degraded as organizational cognition unfolds.

The Correspondence Loop is defined as follows: “The Correspondence Loop is a recursive, event-driven network maintenance mechanism operating on the organizational representation network through four core micro-mechanisms: Representation Update, Relationship Update, Correspondence Evaluation, and Correspondence Revision.”

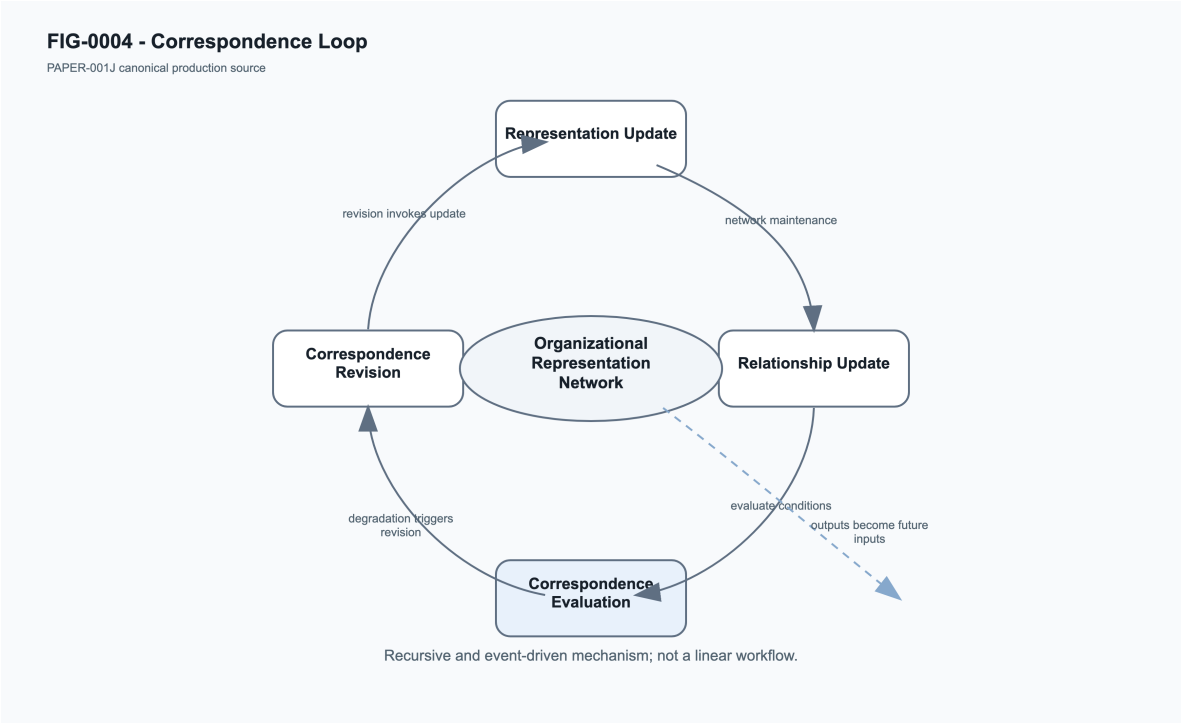

Figure 4 - Correspondence Loop. The Correspondence Loop is a recursive, event-driven network maintenance mechanism operating through four canonical micro-mechanisms.

This definition classifies the Correspondence Loop as a mechanism, not a construct. It does not name a property of the organization or a condition of organizational cognition. It explains how an underlying structure changes, is maintained, degrades, or is revised over time. The underlying structure is the organizational representation network. The mechanism is therefore representation-network-centric rather than actor-centric, process-centric, software-centric, or document-centric.

This classification preserves the architecture of the theory. Organizational Correspondence defines the property that may be preserved or degraded. Cognitive Integrity defines the organizational condition associated with sustained correspondence. The Correspondence Loop explains activity directed toward the representation network. Without a distinct mechanism, the theory would identify what must be preserved but would not explain how preservation, degradation, revision, or restoration occurs through ongoing organizational cognition.

The Correspondence Loop is recursive because outputs of one activity may become inputs to later activity. A representation update may require relationship updates. A relationship update may require correspondence evaluation. Correspondence evaluation may trigger revision. Revision may require further representation or relationship updates. The mechanism is event-driven because local activity is activated by organizational events that affect, or may affect, the representation network. Such events may include decisions, representation changes, relationship changes, execution outcomes, contradictions, learning events, governance decisions, AI-generated artifacts, memory updates, or external disruptions.

The term loop should therefore not be read as a simple sequence. The mechanism is not a linear workflow in which each step always follows the same order. Some events may begin with evaluation, such as when a contradiction is detected. Others may begin with representation update, such as when a new decision is recorded. Others may begin with relationship update, such as when an existing artifact is reinterpreted in light of an outcome. The loop is recursive because activities can revisit and reshape earlier representations and relationships.

The mechanism operates through four canonical micro-mechanisms.

Table 3 - Four Core Micro-Mechanisms

Table 3 - Four Core Micro-Mechanisms

Table 3 - Four Core Micro-Mechanisms

Section 7 - Correspondence Loop

Micro-mechanism	Trigger or condition	Core activity	Representation-network effect	Possible observable trace
Representation Update	A new, obsolete, duplicated, malformed, reclassified, merged, or reinterpreted organizational representation affects the network	Creates, modifies, retires, merges, reclassifies, or reinterprets organizational representations	Keeps the representation network connected to ongoing organizational cognition by ensuring it contains representations capable of supporting meaningful relationships	Representation change, new decision record, revised plan, retired artifact, memory update, AI-generated artifact
Relationship Update	New, modified, or existing representations require semantic or causal connection, reconnection, disconnection, or reinterpretation	Creates, modifies, retires, or reinterprets semantic and causal relationships among organizational representations	Maintains the connective structure from which Organizational Correspondence can emerge	Relationship change, traceability link, revised rationale, updated intent-to-decision or action-to-outcome connection
Correspondence Evaluation	Representation-network conditions may preserve, weaken, contradict, fragment, or restore meaningful semantic and causal relationships	Assesses whether representation-network conditions preserve, weaken, contradict, fragment, or restore meaningful semantic and causal relationships	Detects whether representation-network change has preserved or degraded correspondence conditions	Evaluation record, contradiction signal, ambiguity flag, missing-correspondence finding, confirmed-continuity note
Correspondence Revision	Degradation, contradiction, ambiguity, fragmentation, or missing correspondence is detected or inferred	Revises representations or relationships when degradation, contradiction, ambiguity, fragmentation, or missing correspondence is detected	Restores local coherence, reduces degradation, resolves contradiction, or defers change when evidence or authority is insufficient	Revision event, corrected representation, repaired relationship, deferred revision record, updated evidence or authority note

Table 3 distinguishes the four activities through which the Correspondence Loop operates on the representation network.

Observable traces are possible evidence sources, not validated measures.

Table 3 distinguishes the four activities through which the Correspondence Loop operates on the representation network. Observable traces are possible evidence sources, not validated measures.

Representation Update creates, modifies, retires, merges, reclassifies, or reinterprets organizational representations. Its purpose is to keep the representation network connected to ongoing organizational cognition. Organizations continuously generate new expressions of intent, decision, observation, action, outcome, and memory. If these representations are missing, obsolete, duplicated, or malformed, correspondence conditions weaken. Representation Update ensures that the network contains representations capable of supporting meaningful relationships.

Relationship Update creates, modifies, retires, or reinterprets semantic and causal relationships among organizational representations. This activity maintains the connective structure of the representation network. It is central because Organizational Correspondence emerges from preserved meaningful semantic and causal relationships. A new representation is not sufficient by itself. It must be related to prior intent, decision, action, outcome, policy, learning, or memory where such relationships are meaningful. Relationship Update therefore preserves the network conditions from which correspondence can emerge.

Correspondence Evaluation assesses whether representation-network conditions preserve, weaken, contradict, fragment, or restore meaningful semantic and causal relationships. This activity detects whether representation-network change has preserved or degraded correspondence conditions. Evaluation may occur before, during, or after representation and relationship updates. It may identify contradiction, ambiguity, missing correspondence, fragmentation, or confirmed continuity. Without evaluation, degradation may remain latent while organizational activity continues.

Correspondence Revision revises representations or relationships when degradation, contradiction, ambiguity, fragmentation, or missing correspondence is detected. It is the corrective activity of the mechanism. Revision may restore local coherence, reduce degradation, resolve contradiction, or defer changes when evidence or authority is insufficient. It is recursive rather than terminal because revision can trigger further representation updates, relationship updates, or evaluation.

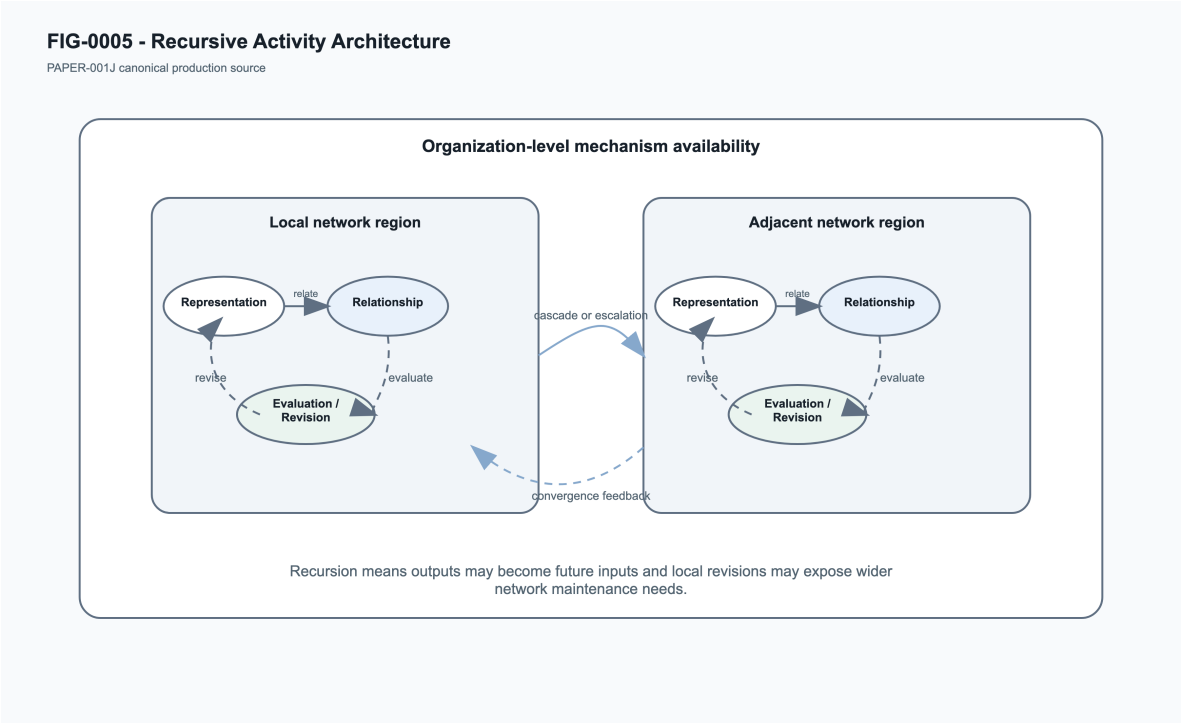

Figure 5 - Recursive Activity Architecture. Recursive activity architecture explains how local Correspondence Loop instances operate within nested representation-network regions.

The four micro-mechanisms should not be interpreted as a linear workflow. They form a recursive activity architecture. Representation Update may require Relationship Update because new or modified representations must be connected to the representation network. Relationship Update may trigger Correspondence Evaluation because new or revised relationships may preserve or degrade correspondence conditions. Correspondence Evaluation may trigger Correspondence Revision when degradation is detected. Correspondence Revision may recursively invoke Representation Update and Relationship Update. The activities may operate in different orders, concurrently, conditionally, and repeatedly depending on the triggering event.

In the logistics example, a service complaint may activate Correspondence Evaluation by revealing a contradiction between the reliability strategy and the AI-generated routing recommendation. Correspondence Revision may then require updating the recommendation representation, revising the relationship between routing criteria and strategic intent, and reconnecting the outcome evidence to organizational learning. This illustrates the mechanism without turning it into a fixed workflow: the activities may begin from evaluation, revision, or update depending on where degradation becomes visible.

This recursive architecture is particularly important in complex organizations. A single decision may affect several representations at once: strategic intent, delivery plans, policy commitments, accountability structures, and expected outcomes. Revising one representation may expose missing relationships elsewhere. Updating a causal relationship may require reinterpretation of an earlier decision. In these cases, the mechanism does not run once and stop. It moves through the representation network as the implications of change become visible.

This architecture clarifies how the Correspondence Loop relates to Organizational Correspondence. The mechanism does not directly manipulate Organizational Correspondence. Organizational Correspondence is an emergent property of the

representation network. The Correspondence Loop operates on the representation network by updating representations, updating relationships, evaluating correspondence conditions, and revising degraded structures. When these activities preserve meaningful semantic and causal relationships, the conditions under which Organizational Correspondence emerges are preserved. When they fail, Organizational Correspondence may degrade.

The relationship to Cognitive Integrity is similarly indirect. The Correspondence Loop does not directly produce Cognitive Integrity. Cognitive Integrity is the dynamic organizational condition associated with sustained Organizational Correspondence. The mechanism supports Cognitive Integrity by supporting the sustained preservation of Organizational Correspondence. If representation-network maintenance preserves correspondence conditions, Cognitive Integrity may be sustained. If those conditions degrade, Cognitive Integrity may decline.

The Correspondence Loop operates at two levels. At the organization level, the mechanism remains available as long as organizational cognition continues and the representation network remains capable of being updated, evaluated, and revised. Availability does not mean that local activity is always executing. Local activity is activated when a specific organizational event affects, or may affect, a region of the representation network. Such local activity may complete, fail, suspend, cascade, conflict, or trigger future activity.

The mechanism also explains characteristic outputs. Immediate outputs may include updated representations, updated semantic relationships, updated causal relationships, detected correspondence degradation, resolved contradictions, reduced degradation, restored local representation-network coherence, rejected or deferred revisions, and updated organizational memory. These outputs should be distinguished from downstream organizational outcomes such as performance, resilience, adaptability, accountability, execution consistency, or effective human-AI collaboration. The Correspondence Loop may influence such outcomes, but they are not part of the mechanism definition.

The Correspondence Loop is distinct from existing organizational mechanisms. It complements cybernetics, organizational learning, distributed cognition, systems dynamics, adaptive systems, PDCA, OODA, knowledge management, sensemaking, and organizational memory. Those theories explain regulation, learning, cognition distribution, dynamic system behavior, adaptation, process improvement, decision cycles, knowledge flows, interpretation, and memory persistence. The Correspondence Loop addresses the residual explanatory gap: existing theories do not explicitly explain the recursive maintenance of the evolving organizational representation network from which Organizational Correspondence emerges.

The mechanism is not a governance process, software workflow, AI orchestration pattern, documentation routine, graph database, or tool architecture. Human-only, hybrid, and AI-native organizations may all exhibit the mechanism if they possess persistent organizational representations, meaningful semantic or causal relationships, recurring organizational cognition, and the capacity to update, evaluate, and revise the representation network.

The boundary conditions follow from this definition. The mechanism applies where persistent organizational representations exist, where semantic or causal relationships among them are possible, where a representation network can be identified, and where organizational cognition recurs over time. It does not apply where there are no persistent representations, no meaningful relationships, no representation network, or no recurring

organizational cognition. This boundary prevents the mechanism from expanding into a general theory of all organizational activity.

The Correspondence Loop can also fail. Trigger failures occur when representation-network events do not activate relevant local activity. Activity failures occur when representation update, relationship update, correspondence evaluation, or correspondence revision fails. Network failures occur when local degradation accumulates into semantic drift, causal traceability collapse, fragmentation, or memory disconnection. These failure modes do not require additional canonical micro-mechanisms. They follow from the four activities and the representation-network substrate on which they operate.

The Correspondence Loop completes the core theoretical architecture of the paper. The representation network provides the substrate. Organizational Correspondence identifies the emergent property at stake. Cognitive Integrity identifies the dynamic organizational condition associated with sustained correspondence. The Correspondence Loop explains the recursive, event-driven mechanism through which organizations maintain, revise, restore, or degrade the representation-network conditions from which correspondence emerges.

8. Integrated Conceptual Model

The preceding sections introduced the components of Correspondence Theory separately: a representation-network ontology, the construct of Organizational Correspondence, the construct of Cognitive Integrity, and the Correspondence Loop as the canonical mechanism. This section integrates those components into a single conceptual model. The model explains how organizations preserve or lose semantic and causal coherence across evolving organizational representations.

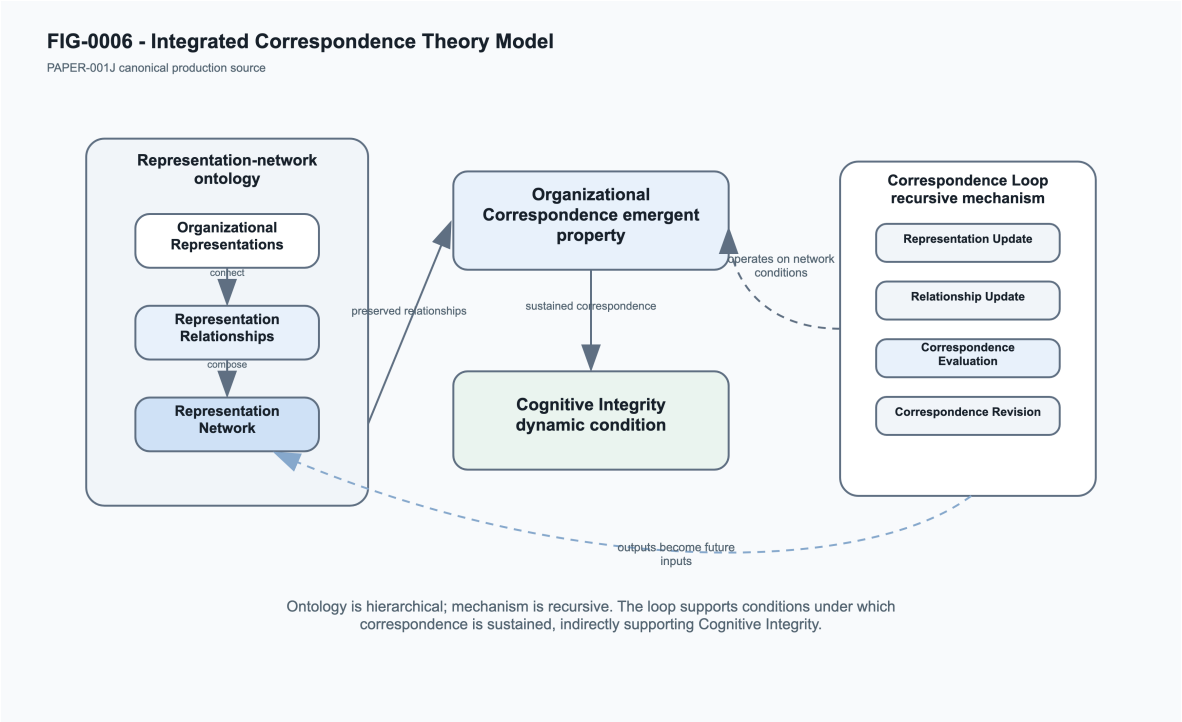

Figure 6 - Integrated Correspondence Theory Model. The integrated model links the representation-network ontology, Organizational Correspondence, Cognitive Integrity, and the Correspondence Loop without collapsing their theoretical roles.

The model begins with organizational representations. These are persistent artifacts, statements, models, records, memories, outputs, or structured expressions through which organizational cognition is represented. They include strategies, decisions, plans, policies, observations, outcomes, governance records, organizational memory, and AI-generated artifacts. Representations matter because they allow organizational cognition to persist beyond a single actor, interaction, or moment.

Representations become theoretically significant through representation relationships. These relationships may be semantic, preserving meaning, interpretation, classification, or implication among representations. They may also be causal, preserving traceability among intent, decision, action, outcome, and learning. The representation network consists of the evolving set of organizational representations and these semantic and causal relationships.

Organizational Correspondence emerges from this network. It is the emergent organizational property by which meaningful semantic and causal relationships among organizational representations are preserved over time. The construct is not a property of isolated artifacts. It depends on the preservation of relationships within the network. A decision, policy, or AI-generated artifact may be internally coherent while still weakening Organizational Correspondence if it is not meaningfully connected to the representations it affects.

Cognitive Integrity is associated with sustained Organizational Correspondence. It is the dynamic organizational condition associated with sustained Organizational Correspondence. This means that Cognitive Integrity depends on the preservation of correspondence conditions across time. It does not refer to individual cognition, organizational performance, or technical sophistication. It names the organization-level condition associated with sustained coherence in represented organizational cognition.

The Correspondence Loop explains how the representation network is maintained, revised, restored, or degraded. The mechanism is recursive and event-driven. It operates on the representation network through four core micro-mechanisms: Representation Update, Relationship Update, Correspondence Evaluation, and Correspondence Revision. These activities preserve or degrade the conditions under which Organizational Correspondence emerges. Cognitive Integrity is therefore sustained indirectly through sustained Organizational Correspondence.

The integrated model should not be read as a simple downstream sequence in which the Correspondence Loop appears after Cognitive Integrity. The ontology is hierarchical, but the mechanism is recursive. The representation network is the substrate. Organizational Correspondence emerges from preserved semantic and causal relationships within that network. Cognitive Integrity is the dynamic condition associated with sustained Organizational Correspondence. The Correspondence Loop recursively operates on the representation network, supporting the conditions under which the constructs are sustained without being ontologically downstream from them.

This architecture preserves the distinctions that make the theory coherent. The representation network is not a construct in the same sense as Organizational Correspondence or Cognitive Integrity; it is the ontological substrate. Organizational Correspondence is not a mechanism; it is an emergent property. Cognitive Integrity is not a capability or process; it is a dynamic organizational condition. The Correspondence Loop is not a construct; it is the mechanism that explains representation-network maintenance.

The integrated model also clarifies the role of AI-native organizations. AI systems may increase the speed, volume, and distribution of organizational representation creation. They may generate new artifacts, revise existing ones, recommend actions, summarize memory, or influence decisions. However, AI does not define the theory. It intensifies the representational problem. The same model applies to human-only, hybrid, and AI-native organizations wherever persistent organizational representations and meaningful relationships among them exist.

The model further clarifies what Correspondence Theory does not explain. It is not a comprehensive theory of performance, strategy, culture, power, leadership, or institutional change. Those phenomena may interact with Organizational Correspondence and Cognitive Integrity, but they are not replaced by them. The conceptual model explains a narrower problem: whether represented organizational cognition preserves meaningful semantic and causal relationships as it evolves. This boundary is important for theoretical precision. It allows the paper to make a distinctive contribution without claiming to account for all sources of organizational success or failure.

The model also establishes the temporal logic of the theory. Organizational representations are not fixed. They are created, revised, retired, reinterpreted, and related through ongoing organizational activity. Correspondence is therefore not a one-time state. It is preserved or degraded over time as the representation network changes. Cognitive Integrity likewise varies as sustained correspondence is maintained, weakened, or restored. The Correspondence Loop explains this dynamic by identifying how organizations update, relate, evaluate, and revise representations and relationships.

The next section presents the theoretical propositions that follow from the integrated architecture. These propositions do not introduce new theory. They express expected relationships among the ontology, constructs, and mechanism already developed. The research agenda then explains how future work can investigate those relationships through observable traces while preserving the distinction between theory, operationalization, and empirical evidence.

9. Theoretical Propositions

The preceding sections developed the conceptual architecture of Correspondence Theory. The theory now contains an ontology, two formal constructs, and one canonical mechanism. This section presents the proposition architecture that follows from that theory. The purpose is not to test the theory empirically or define measures. It is to state the expected theoretical relationships that future empirical work can investigate.

Correspondence Theory distinguishes predictions, propositions, hypotheses, measurements, and empirical tests. Predictions describe what the theory expects to occur. Propositions express those expectations as formal theoretical statements. Hypotheses operationalize propositions for empirical testing. Measurements define observable indicators, variables, and metrics. Empirical tests evaluate whether hypotheses are supported. This section remains at the proposition level.

Table 4 - Canonical Proposition Summary

Table 4 - Canonical Proposition Summary

Table 4 - Canonical Proposition Summary

Section 9 - Theoretical Propositions

Proposition	Category	Complete canonical proposition	Antecedent or condition	Theoretical relationship	Expected consequence	Level of analysis
Structural propositions						
P-001	**Structural**	Higher quality representation-network maintenance is positively associated with Organizational Correspondence.	Higher quality representation-network maintenance	Positive association with Organizational Correspondence	Stronger Organizational Correspondence	Representation-network / organization
P-002	**Structural**	Representation-network degradation accumulates when Correspondence Loop activities fail to preserve semantic and causal relationships.	Correspondence Loop activities fail to preserve semantic and causal relationships	Failure permits degradation accumulation	Accumulated representation-network degradation	Representation-network
P-003	**Structural**	Increasing representation complexity without corresponding maintenance capacity reduces Organizational Correspondence.	Increasing representation complexity without corresponding maintenance capacity	Complexity exceeds maintenance capacity	Reduced Organizational Correspondence	Representation-network / organization
Dynamic propositions						
P-004	**Dynamic**	The frequency of local Correspondence Loop mechanism instances increases with the rate of meaningful representation-network change.	Higher rate of meaningful representation-network change	Increased activation frequency of local mechanism instances	More frequent Correspondence Loop activity	Local representation-network region
P-005	**Dynamic**	Higher Mechanism Readiness improves recovery from representation-network degradation.	Higher Mechanism Readiness	Improved recovery capacity	Better recovery from representation-network degradation	Mechanism / organization
P-006	**Dynamic**	Successful Correspondence Revision reduces accumulated correspondence degradation within the affected region of the representation network.	Successful Correspondence Revision	Corrective revision reduces degradation	Reduced accumulated degradation in affected network region	Local representation-network region
Organizational propositions						
P-007	**Organizational**	Higher Organizational Correspondence is positively associated with higher Cognitive Integrity.	Higher Organizational Correspondence	Positive association with Cognitive Integrity	Higher Cognitive Integrity	Organization
P-008	**Organizational**	Representation-network degradation precedes observable declines in Cognitive Integrity.	Representation-network degradation	Temporal precedence before Cognitive Integrity decline	Observable declines in Cognitive Integrity	Organization over time
P-009	**Organizational**	Organizations with similar resources, strategies, and governance may exhibit different levels of Cognitive Integrity because of differences in representation-network maintenance.	Similar resources, strategies, and governance but different representation-network maintenance	Maintenance differences explain variation	Different levels of Cognitive Integrity	Comparative organization
AI-native propositions						
P-010	**AI-native**	AI-native organizations require more frequent Correspondence Loop activation than organizations with lower rates of representation-network change.	AI-native organizations with higher rates of representation-network change	Increased maintenance demand	More frequent Correspondence Loop activation	AI-native organization
P-011	**AI-native**	Increasing AI-generated representation volume without corresponding increases in representation-network maintenance increases correspondence degradation risk.	Increasing AI-generated representation volume without increased maintenance	Volume outpaces maintenance	Increased correspondence degradation risk	AI-native representation network
P-012	**AI-native**	AI systems improve organizational cognition only when representation-network maintenance scales with AI-generated representation change.	AI systems and AI-generated representation change	Improvement depends on scaled maintenance	Organizational cognition improves only under scaled maintenance conditions	AI-native organization

Table 4 summarizes the complete proposition architecture while preserving the canonical wording of each proposition.

Propositions are theoretical statements, not empirical findings.

Table 4 summarizes the complete proposition architecture while preserving the canonical wording of each proposition. Propositions are theoretical statements, not empirical findings.

The canonical proposition catalog organizes propositions into four categories: structural, dynamic, organizational, and AI-native. Structural propositions concern representations, relationships, and representation networks. Dynamic propositions concern Correspondence Loop behavior. Organizational propositions concern Organizational Correspondence and Cognitive Integrity. AI-native propositions concern organizations in which AI-generated representations increase the speed, volume, or distribution of organizational cognition.

This categorization is not intended to create additional constructs. It is a manuscript and research-program device for organizing expectations derived from the theory. Each category traces back to canonical elements already introduced. Structural propositions trace to the representation-network ontology. Dynamic propositions trace to the Correspondence Loop. Organizational propositions trace to the dependency between Organizational Correspondence and Cognitive Integrity. AI-native propositions trace to the claim that AI increases the rate and volume of representation-network change.

Structural propositions express the theory's expectation that representation-network maintenance matters for Organizational Correspondence. The first proposition states that higher quality representation-network maintenance is positively associated with Organizational Correspondence. The second states that representation-network degradation accumulates when Correspondence Loop activities fail to preserve semantic and causal relationships. The third states that increasing representation complexity without corresponding maintenance capacity reduces Organizational Correspondence.

Dynamic propositions follow from the mechanism. They concern the activity of the Correspondence Loop rather than the constructs themselves. The fourth proposition states that the frequency of local Correspondence Loop mechanism instances increases with the rate of meaningful representation-network change. The fifth states that higher Mechanism Readiness improves recovery from representation-network degradation. The sixth states that successful Correspondence Revision reduces accumulated correspondence degradation within the affected region of the representation network.

Organizational propositions connect the two formal constructs. The seventh proposition states that higher Organizational Correspondence is positively associated with higher Cognitive Integrity. The eighth states that representation-network degradation precedes observable declines in Cognitive Integrity. The ninth states that organizations with similar resources, strategies, and governance may exhibit different levels of Cognitive Integrity because of differences in representation-network maintenance.

AI-native propositions specify expected implications for organizations in which AI systems generate, modify, or act upon organizational representations at increased speed and scale. The tenth proposition states that AI-native organizations require more frequent Correspondence Loop activation than organizations with lower rates of representation-network change. The eleventh states that increasing AI-generated representation volume without corresponding increases in representation-network maintenance increases correspondence degradation risk. The twelfth states that AI systems improve organizational cognition only when representation-network maintenance scales with AI-generated representation change.

These propositions are theoretical statements rather than empirical findings. They should not be read as validated claims. Their value is architectural: they show how the theory can generate testable expectations without redefining its constructs or mechanism. They also provide a disciplined bridge from conceptual theory to future research. Future hypotheses

should derive from these propositions rather than creating independent theoretical claims.

The proposition architecture also protects the theory from performance reductionism. Correspondence Theory does not claim that high Organizational Correspondence guarantees organizational success or that low Cognitive Integrity explains every failure. Market conditions, regulation, leadership, resources, competition, politics, and external shocks all matter. The propositions focus on relationships among representation-network maintenance, Organizational Correspondence, Cognitive Integrity, and AI-native representation change.

The propositions also indicate the empirical shape of the research program. Structural propositions point researchers toward evidence concerning representations and relationships. Dynamic propositions point toward traces of mechanism activity, including updates, evaluations, and revisions. Organizational propositions point toward inferred relationships between correspondence conditions and Cognitive Integrity. AI-native propositions point toward contexts in which representation-network change is accelerated by machine-generated artifacts. These are research directions, not measurement specifications.

The publication version summarizes these propositions for scholarly presentation; future empirical work should retain traceability to the underlying proposition catalogue without exposing repository identifiers in journal-facing prose. Doing so preserves cumulative theory development. It allows later papers to derive hypotheses, identify observable traces, construct variables, and design empirical studies while remaining traceable to the conceptual model developed here.

10. Research Agenda

Correspondence Theory is developed here as a conceptual theory, but its architecture is designed to support a cumulative research program. The preceding section presented propositions derived from the representation-network ontology, Organizational Correspondence, Cognitive Integrity, and the Correspondence Loop. This section outlines how future research can investigate those propositions without redefining the theory.

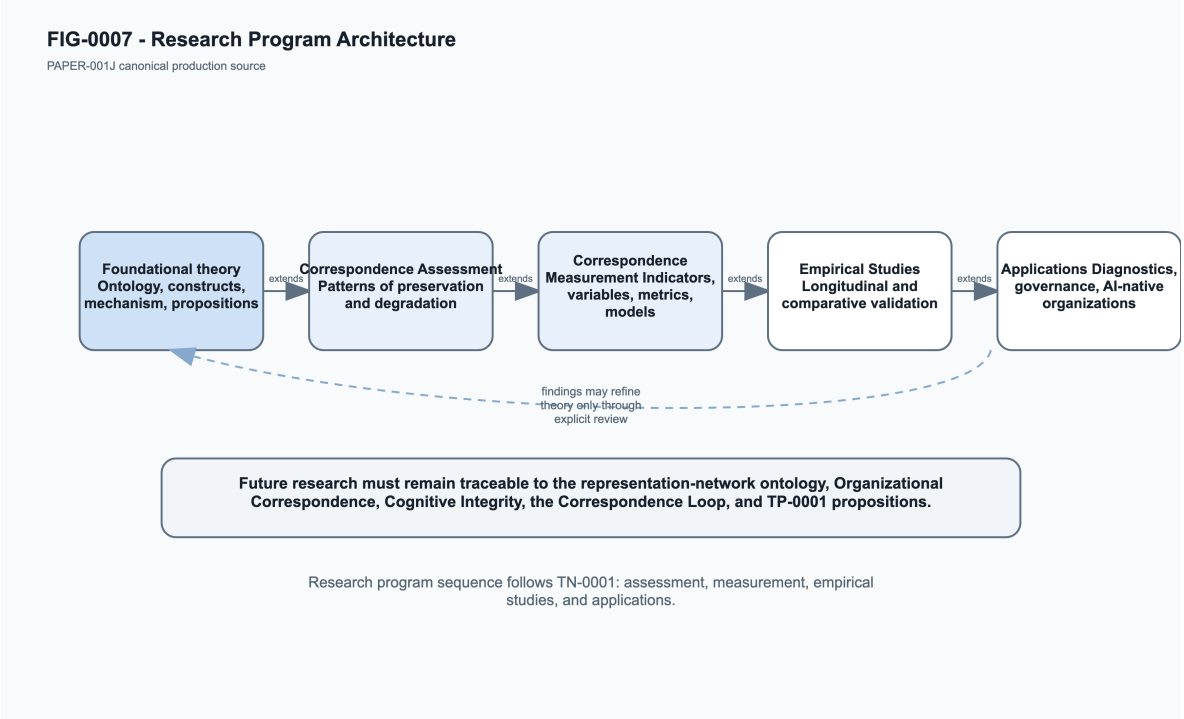

Figure 7 - Research Program Architecture. Paper-001 establishes a cumulative research program from conceptual theory toward assessment, measurement, empirical studies, and applications.

The research program should begin from the distinction between theory and empirical evidence. Theoretical constructs and mechanisms are not directly observed. Researchers observe traces and infer theoretical phenomena. For Correspondence Theory, relevant traces may include representation changes, relationship changes, evaluation records, revision events, contradiction signals, decision traces, governance records, organizational memory updates, and AI reasoning traces. These traces may support inferences about representation-network maintenance, Organizational Correspondence, Cognitive Integrity, and Correspondence Loop activity.

Future diagnostic work should focus on identifying patterns of correspondence preservation and degradation. Such work may examine semantic continuity, causal traceability, representation fragmentation, ambiguity, contradiction, or missing relationships. However, diagnostic frameworks should remain downstream of the theory. They should interpret patterns in the representation network without redefining Organizational Correspondence, Cognitive Integrity, or the Correspondence Loop.

Future measurement work should proceed separately from diagnostics. The research methodology for this program distinguishes observable traces, indicators, variables, metrics, measurement models, validation studies, and empirical evidence. This sequence matters because indicators are not variables, variables are not metrics, and metrics are not theory. Measurement frameworks should define how empirical observations represent theoretical concepts, but they should not alter the canonical definitions.

Table 5 - Future Research Agenda

Table 5 - Future Research Agenda

Table 5 - Future Research Agenda

Section 10 - Research Agenda

Research domain	Illustrative research question	Potential unit or level of analysis	Possible empirical setting	Relevant evidence or observable traces	Intended theoretical contribution
Diagnostics and measurement					
Diagnostic work	What patterns indicate correspondence preservation or degradation in an evolving representation network?	Representation, relationship, network region, organization	Organizations with decision records, policies, plans, memory artifacts, AI-generated outputs, or governance records	Representation changes, relationship changes, contradiction signals, ambiguity signals, missing relationships, fragmentation patterns	Identify patterns of correspondence preservation and degradation without redefining Organizational Correspondence, Cognitive Integrity, or the Correspondence Loop
Measurement work	How can observable traces become indicators, variables, metrics, measurement models, and validation studies without becoming theory?	Indicator, variable, metric, construct-measure relationship	Research settings with structured traces and sufficient evidence quality for measurement development	Observable traces, indicators, variables, metrics, measurement-model evidence, validation records	Build measurement frameworks that represent theoretical concepts without altering canonical definitions
Empirical design					
Empirical validation	Which hypotheses derived from structural, dynamic, organizational, and AI-native propositions receive empirical support?	Hypothesis, proposition category, organization, representation network	Comparative or field studies of organizations with observable representation-network activity	Representation-network maintenance evidence, Correspondence Loop activation traces, revision events, AI-generated representation changes	Test hypotheses derived from the proposition architecture while preserving the distinction between propositions and empirical findings
Longitudinal research	How do correspondence conditions persist, degrade, revise, or recover over time?	Representation network over time, network region, organization	Longitudinal organizational records, decision histories, learning histories, AI-mediated work histories	Time-stamped representation changes, relationship revisions, degradation accumulation, recovery events, learning updates	Explain temporal preservation, degradation, revision, and recovery dynamics
Comparative organizational research	Why might organizations with similar resources, strategies, or governance exhibit different levels of Cognitive Integrity?	Organization, comparable organizational unit, governance context	Human-only, hybrid, and AI-native organizations with comparable strategies or operating conditions	Differences in representation-network maintenance, governance records, decision traces, memory updates, execution outcomes	Explain variation in Cognitive Integrity through differences in representation-network maintenance
AI and intervention research					
AI-native organizational research	Does increased AI-generated representation volume increase maintenance demands and correspondence degradation risk when maintenance capacity does not scale?	AI-native organization, AI-generated representation stream, human-AI representation network	Organizations using AI to generate, summarize, classify, recommend, revise, or act upon representations	AI-generated artifacts, prompts or reasoning traces where available, recommendation records, policy links, outcome links, memory links	Clarify how AI-generated representation change affects Correspondence Loop activation and correspondence degradation risk
Intervention-oriented research	Which practices help organizations detect contradiction, preserve causal traceability, revise obsolete representations, or connect learning back to prior decisions?	Intervention, practice, representation-network region, organization	Organizations testing practices for traceability, revision, contradiction detection, or learning integration	Intervention records, revised representations, repaired relationships, evidence notes, authority or governance review records	Study applications of the theory without adding practices to the theory itself
Evidence quality					
Evidence-quality research	How does evidence quality affect the strength of inferences about Organizational Correspondence, Cognitive Integrity, and Correspondence Loop activity?	Evidence source, trace, inference, study design	Settings with explicit, partial, informal, or retrospective records	Time-stamped evidence, traceability quality, completeness of records, retrospective accounts, missing evidence	Separate theoretical inference from the quality of evidence supporting that inference

Table 5 translates the proposition architecture into prospective research domains without presenting them as completed validation.

Research directions are prospective; they do not constitute validated measures or empirical findings.

Table 5 translates the proposition architecture into prospective research domains without presenting them as completed validation. Research directions are prospective; they do not constitute validated measures or empirical findings.

Empirical validation should derive hypotheses from the proposition architecture. Structural propositions invite studies of whether representation-network maintenance is associated with Organizational Correspondence and whether degradation accumulates when semantic and causal relationships are not preserved. Dynamic propositions invite studies of Correspondence Loop activation, recovery from degradation, and revision effectiveness. Organizational propositions invite studies of the relationship between Organizational Correspondence and Cognitive Integrity. AI-native propositions invite studies of whether AI-generated representation change increases maintenance demands.

Longitudinal research is especially important. Correspondence Theory concerns preservation, degradation, revision, and recovery over time. Cross-sectional studies may provide useful snapshots, but they are limited for studying whether correspondence conditions are sustained or degraded. Longitudinal designs can examine how representations change, how relationships are revised, whether degradation accumulates, and whether local revision restores coherence in affected regions of the representation network.

Comparative organizational research is also promising. Organizations with similar strategies, resources, or governance structures may differ in Cognitive Integrity because they differ in representation-network maintenance. Comparative studies could examine human-only, hybrid, and AI-native organizations to understand how rates of representation change affect the frequency and burden of Correspondence Loop activity. Such studies should preserve implementation independence by treating tools and systems as possible sources of traces, not as definitions of the theory.

AI-native organizational research is a particularly important domain. AI systems can generate, summarize, classify, recommend, and revise representations at high speed. Future research should investigate whether increased AI-generated representation volume increases correspondence degradation risk when maintenance capacity does not scale. It should also examine whether AI improves organizational cognition only when its outputs remain connected to intent, decisions, policies, outcomes, and organizational memory.

Intervention-oriented research may eventually examine how organizations improve representation-network maintenance. Such work should remain downstream of the conceptual theory. It may investigate practices that help organizations detect contradiction, preserve causal traceability, revise obsolete representations, or connect learning back to prior decisions. However, these practices should be treated as possible applications of the theory rather than additions to the theory itself.

Evidence quality should be reported separately from theoretical findings. Some organizations may provide explicit, time-stamped, traceable evidence linked to relevant representations and relationships. Others may provide partial records, informal documentation, or retrospective accounts. Weak evidence does not invalidate the theory, but it limits the strength of research inference. Future studies should therefore distinguish between the theoretical condition being inferred and the quality of evidence supporting that inference.

This research agenda positions Paper 001 as the conceptual foundation rather than the empirical endpoint. Future work should develop diagnostic frameworks, measurement frameworks, hypothesis sets, empirical study designs, validation studies, longitudinal datasets, and AI-native organizational case studies. Each should remain traceable to the canonical theory and proposition catalog. The objective is cumulative research: extending,

testing, refining, and, where necessary, revising Correspondence Theory through disciplined empirical work.

The discussion now turns from research program design to the broader implications of the theory. If the future work described here is to remain cumulative, the manuscript must be clear about what Correspondence Theory contributes, what it does not claim, and how it should be interpreted by organizational scholars and practitioners.

11. Discussion

This paper develops Correspondence Theory as a representation-network-centered theory of organizational cognition. Its central claim is that organizations increasingly depend on evolving networks of representations and that the preservation of meaningful semantic and causal relationships among those representations is a distinct theoretical problem.

The paper makes six contributions. First, it introduces a representation-network ontology for studying organizational cognition through persistent representations and their semantic and causal relationships. Second, it defines Organizational Correspondence as the first canonical construct of the theory. Third, it defines Cognitive Integrity as a second construct that captures the dynamic organizational condition associated with sustained Organizational Correspondence. Fourth, it introduces the Correspondence Loop as the recursive, event-driven mechanism of representation-network maintenance. Fifth, it organizes the theory into a proposition architecture for future empirical work. Sixth, it positions Correspondence Theory as a cumulative research program that can support later diagnostics, measurement, and validation without redefining the core theory.

The first implication is theoretical. Correspondence Theory shifts attention from isolated organizational artifacts, individual cognition, or local decision processes to the representation network through which organizational cognition becomes persistent and revisable. This does not replace existing theories of learning, sensemaking, memory, feedback, or distributed cognition. It complements them by identifying a residual problem: how meaningful relationships among organizational representations are preserved or degraded over time. The theory therefore provides a new way to study coherence in organizational cognition without reducing that coherence to communication, alignment, knowledge storage, or performance.

The second theoretical implication concerns construct architecture. Organizational Correspondence and Cognitive Integrity are related but distinct. Organizational Correspondence identifies the emergent property by which meaningful semantic and causal relationships among organizational representations are preserved over time. Cognitive Integrity identifies the dynamic organizational condition associated with sustained Organizational Correspondence. Separating these constructs allows the theory to distinguish the property of correspondence from the organizational condition associated with its sustained preservation.

The third theoretical implication concerns mechanism. The Correspondence Loop explains how the representation network is recursively maintained through Representation Update, Relationship Update, Correspondence Evaluation, and Correspondence Revision. This mechanism is not an execution cycle, decision loop, governance process, or software workflow. It is a recursive, event-driven network maintenance mechanism. This distinction matters because organizations may make decisions, learn, adapt, and improve processes while still losing coherence among the representations those activities generate.

The paper also has managerial implications, although it does not prescribe a management method. Managers increasingly operate in organizations where representations proliferate across documents, systems, AI outputs, decisions, plans, policies, and memory. The theory suggests that the central challenge is not only producing better representations, but preserving relationships among them. Strategic intent must remain connected to decisions. Decisions must remain connected to actions. Actions must remain connected to outcomes. Outcomes must remain connected to learning. Learning must revise the representations it affects. When these relationships degrade, organizations may appear active and information-rich while becoming cognitively fragmented.

The logistics example illustrates the implication. A leadership team may believe it has adopted a reliability strategy while operational teams act on an AI-generated routing recommendation still tied to short-term efficiency. The managerial issue is not simply that a decision was poorly communicated. It is that the representations linking strategy, recommendation, action, outcome, and learning were not maintained. Restoring coherence requires reconnecting these representations so that future organizational cognition reflects the revised relationship.

For AI-native organizations, the implications are especially salient. AI systems can increase the speed and volume of representation creation. They may also make representation change less visible because outputs can be generated, summarized, recombined, and acted upon rapidly. Correspondence Theory suggests that AI-native cognition requires attention to whether AI-generated representations remain connected to organizational intent, decision history, causal assumptions, governance commitments, outcomes, and memory. AI can support organizational cognition, but only if the representation network remains coherent enough for Organizational Correspondence to be sustained.

The theory also clarifies what this paper does not claim. Correspondence Theory is not a general theory of organizational performance. High Organizational Correspondence does not guarantee success, and low Cognitive Integrity does not explain every failure. Organizations may fail because of market shifts, poor strategy, resource constraints, regulation, competition, politics, leadership failures, or external shocks. The theory is narrower: it explains coherence and degradation in represented organizational cognition. This boundary protects the theory from overclaiming while preserving its distinctive contribution.

Several limitations follow. First, the paper is conceptual and does not empirically validate the theory. The propositions remain theoretical statements requiring future research. Second, the theory depends on trace-based inference. Researchers rarely observe complete representation networks, so empirical work must account for partial observability and evidence quality. Third, the paper introduces the core theory but does not develop full diagnostic or measurement frameworks. Those are necessary next steps but should remain downstream of the canonical theory. Fourth, the literature integration offered here will require deeper citation work in a full journal manuscript.

Future research should therefore proceed along several paths. Conceptual work can refine the relationship between Correspondence Theory and adjacent literatures. Diagnostic work can identify patterns of correspondence degradation. Measurement work can operationalize variables and validation models. Empirical work can test propositions using longitudinal and comparative designs. AI-native organizational studies can examine how machine-generated representations alter the burden of representation-network maintenance. Across these paths, future work should preserve the theory's central architecture rather than redefining its

foundational constructs or mechanism.

The broader contribution of the paper is to make organizational coherence theoretically visible at the level of representations and relationships. As organizations become more distributed, AI-mediated, and representationally dense, this level of analysis becomes increasingly important. Correspondence Theory offers a vocabulary and architecture for studying whether organizational cognition remains meaningfully and causally connected as it changes.

12. Conclusion

Organizations increasingly think, decide, remember, and act through evolving networks of representations. Strategies, decisions, policies, work products, outcomes, governance records, organizational memory, and AI-generated artifacts all contribute to the represented structure of organizational cognition. As these representations multiply and change, the central challenge is not simply whether organizations possess information, knowledge, or memory. It is whether meaningful semantic and causal relationships among representations remain intact over time.

This paper introduced Correspondence Theory to address that challenge. It developed a representation-network ontology, defined Organizational Correspondence as the emergent property by which meaningful semantic and causal relationships among organizational representations are preserved over time, defined Cognitive Integrity as the dynamic organizational condition associated with sustained Organizational Correspondence, and introduced the Correspondence Loop as the recursive, event-driven mechanism through which the representation network is maintained, revised, or degraded.

The paper also positioned the theory as a cumulative research program. Its proposition architecture identifies expected relationships among representation-network maintenance, Organizational Correspondence, Cognitive Integrity, and AI-native organizational cognition. Its research methodology preserves the distinction between theory, observable traces, indicators, variables, measurement, and empirical evidence. Future research can therefore test, refine, and extend the theory without redefining its foundations.

The theory is intentionally bounded. It does not claim to explain every source of organizational performance, failure, adaptation, or strategic advantage. Instead, it isolates a specific problem that becomes increasingly consequential as organizations rely on distributed human and machine cognition: the preservation of correspondence across evolving representations. This boundary allows the theory to complement existing accounts of learning, feedback, memory, sensemaking, and distributed cognition while contributing a distinct explanatory architecture.

The core contribution is a new conceptual architecture for studying coherence in organizational cognition. Existing theories explain learning, feedback, memory, adaptation, sensemaking, and distributed cognition. Correspondence Theory complements these accounts by explaining the preservation and degradation of semantic and causal relationships across evolving organizational representation networks. AI-native organizations require mechanisms for preserving correspondence across evolving representation networks.

References

Alavi, M. and Leidner, D.E., 2001. Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues. MIS Quarterly, 25(1), pp.107-136. Available at: https://doi.org/10.2307/3250961.

Argyris, C. and Schon, D.A., 1978. Organizational Learning: A Theory of Action Perspective. Reading, MA: Addison-Wesley.

Ashby, W.R., 1956. An Introduction to Cybernetics. London: Chapman & Hall.

von Bertalanffy, L., 1968. General System Theory: Foundations, Development, Applications. New York: George Braziller.

Daft, R.L. and Lengel, R.H., 1986. Organizational information requirements, media richness and structural design. Management Science, 32(5), pp.554-571. Available at: https://doi.org/10.1287/mnsc.32.5.554.

Faraj, S., Pachidi, S. and Sayegh, K., 2018. Working and organizing in the age of the learning algorithm. Information and Organization, 28(1), pp.62-70. Available at: https://doi.org/10.1016/j.infoandorg.2018.02.005.

Galbraith, J.R., 1974. Organization design: An information processing view. Interfaces, 4(3), pp.28-36. Available at: https://doi.org/10.1287/inte.4.3.28.

Henderson, J.C. and Venkatraman, N., 1993. Strategic alignment: Leveraging information technology for transforming organizations. IBM Systems Journal, 32(1), pp.4-16.

Hutchins, E., 1995. Cognition in the Wild. Cambridge, MA: MIT Press.

Kellogg, K.C., Valentine, M.A. and Christin, A., 2020. Algorithms at work: The new contested terrain of control. Academy of Management Annals, 14(1), pp.366-410. Available at: https://doi.org/10.5465/annals.2018.0174.

March, J.G., 1991. Exploration and exploitation in organizational learning. Organization Science, 2(1), pp.71-87. Available at: https://doi.org/10.1287/orsc.2.1.71.

March, J.G. and Simon, H.A., 1958. Organizations. New York: Wiley.

Nonaka, I., 1994. A dynamic theory of organizational knowledge creation. Organization Science, 5(1), pp.14-37. Available at: https://doi.org/10.1287/orsc.5.1.14.

Orlikowski, W.J., 2007. Sociomaterial practices: Exploring technology at work. Organization Studies, 28(9), pp.1435-1448. Available at: https://doi.org/10.1177/0170840607081138.

Raisch, S. and Krakowski, S., 2021. Artificial intelligence and management: The automation-augmentation paradox. Academy of Management Review, 46(1), pp.192-210. Available at: https://doi.org/10.5465/amr.2018.0072.

Ross, J.W., Weill, P. and Robertson, D.C., 2006. Enterprise Architecture as Strategy: Creating a Foundation for Business Execution. Boston, MA: Harvard Business School Press.

Star, S.L. and Griesemer, J.R., 1989. Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. Social Studies of Science, 19(3), pp.387-420. Available at: https://doi.org/10.1177/030631289019003001.

Teece, D.J., Pisano, G. and Shuen, A., 1997. Dynamic capabilities and strategic management. Strategic Management Journal, 18(7), pp.509-533. Available at: https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z.

Venkatraman, N., 1989. The concept of fit in strategy research: Toward verbal and statistical correspondence. Academy of Management Review, 14(3), pp.423-444. Available at: https://doi.org/10.5465/amr.1989.4279078.

Walsh, J.P. and Ungson, G.R., 1991. Organizational memory. Academy of Management Review, 16(1), pp.57-91. Available at: https://doi.org/10.5465/amr.1991.4278992.

Wegner, D.M., 1987. Transactive memory: A contemporary analysis of the group mind. In B. Mullen and G.R. Goethals, eds. Theories of Group Behavior. New York: Springer, pp.185-208. Available at: https://doi.org/10.1007/978-1-4612-4634-3_9.

Weick, K.E., 1995. Sensemaking in Organizations. Thousand Oaks, CA: SAGE.