Issues in Information Systems
Volume 27, Issue 3, pp. 716-733, 2026

DOI: https://doi.org/10.48009/3_iis_2026_149

Mapping IS research on Gen-to-Agentic AI adoption: A systematic bibliometric review and research agenda

Leila Gluszak, *Paris Nanterre University, leila.gluszak@outlook.com*

Abstract

As organizations transition from generative AI tools to semi-autonomous agentic systems, IS scholarship lacks consolidated understanding of this paradigm shift's adoption determinants. We ask: How is IS scholarship structuring the organizational transition from generative to agentic AI? Through systematic bibliometric review of 2,789 Scopus publications (2023–2026) combining RCCA, temporal DBCA, and interpretive coding of 205 multi-paradigm documents we identify five intellectual schools, a 492% growth in agentic AI discourse, and four discourse-grounded adoption pathways. We develop a four-block contingency framework (technological, organizational, individual, institutional) highlighting orchestration and governance as under-theorized constructs, and formulate two testable propositions validated through regression and fsQCA designs. Contributions are methodological, conceptual, and programmatic, positioning IS scholarship to lead empirical investigation of the most consequential AI transition currently underway in organizations.

Keywords: generative AI, agentic AI, organizational AI adoption, bibliometric analysis, contingency framework, dynamic capabilities

Introduction

Generative Artificial Intelligence (Gen AI) based on Large Language Models (LLM) entered organizational practice in late 2022, with hundreds of pilots promising productivity gains in knowledge work (Dwivedi et al., 2023). Simultaneously, Agentic AI systems combining generative capabilities with autonomous goal pursuit, planning, and tool use are entering research discourse (Chui et al., 2024). As organizations transition from reactive AI tools to semi-autonomous agents, Information Systems scholarship lacks consolidated understanding of this shift and its adoption determinants.

Existing Information Systems (IS) research on AI adoption builds on Technology-Organization-Environment (TOE) frameworks, Technology Acceptance Models (TAM), and Dynamic Capabilities (DC) perspectives (Tornatzky & Fleischer, 1990; Davis, 1989; Teece et al., 1997). However, most work predates the Gen AI surge or treats AI as undifferentiated. Recent reviews focus primarily on Gen AI (Dwivedi et al., 2023), while literature on agentic AI and multi-paradigm ecosystems remains nascent and fragmented. This raises the central research question: How is IS scholarship structuring the organizational transition from generative to agentic AI?

RQ1: What intellectual structure and temporal evolution characterize the emerging discourse on Gen-to-Agentic AI adoption?
RQ2: What adoption conditions and transition pathways does this discourse associate with the organizational adoption of multi-paradigm AI?

RQ3: What theoretical gaps does this discourse reveal, and what propositions follow for organizational validation?

We address this question through systematic bibliometric analysis (Zupic & Čater, 2015) of 2,789 Scopus publications (2023-2026) using Reference Co-Citation Analysis (RCCA) to identify intellectual schools, temporal Document Bibliographic Coupling Analysis (DBCA) to track thematic evolution, and interpretive coding explicitly discussing both paradigms.

We make three contributions. First, methodological: this bibliometric triangulation documents the Gen-to-Agentic transition's intellectual structure and temporal evolution. Second, conceptual: synthesizing patterns across intellectual schools, temporal trends, and transition documents, we develop contingency framework organizing adoption conditions into four blocks and four ideal-type pathways, highlighting orchestration and governance as under-theorized constructs. Third, programmatic: we formulate two testable propositions with detailed validation designs providing roadmap for organizational research.
We position this work as systematic literature review mapping scholarly discourse, not empirical organizational study. Our contributions lie in documenting field emergence, identifying theoretical gaps, and providing research infrastructure for hypothesis testing. The contingency framework and propositions represent synthesis of literature patterns requiring empirical validation through organizational data.
Section 2 reviews conceptual background. Section 3 details bibliometric methodology. Section 4 presents findings. Section 5 synthesizes results, derives propositions, and discusses implications and limitations. Section 6 concludes.

Conceptual background

To address the research question, we first review theoretical perspectives structuring IS scholarship on organizational AI adoption.

Distinguishing Generative and Agentic AI

Gen AI systems, exemplified by LLMs and other foundation models, produce content in response to human prompts and operate in predominantly reactive mode (Dwivedi et al., 2023). These systems excel at tasks including text generation, summarization, translation, code completion, and content creation. They are evaluated primarily on output quality, coherence, and usefulness for knowledge work augmentation. The interaction paradigm is human-in-the-loop: users initiate each interaction through explicit prompts, directly review outputs, and make all consequential decisions based on AI-generated recommendations.

Agentic AI extends generative capabilities with five properties distinguishing it from reactive prompt-response systems. First, autonomy: agents pursue objectives without constant human intervention. Second, planning: agents decompose high-level goals into executable subtasks. Third, persistent memory: agents maintain state and context across execution steps. Fourth, tool use: agents invoke external APIs, databases, and enterprise systems. Fifth, goal-directed behavior: agents iteratively refine their approach through observation and adaptation (Wang et al., 2024; Masterman et al., 2024). Unlike generative systems responding to single prompts, agentic systems receive directives, plan execution sequences, observe outcomes, and adjust strategies autonomously.

This distinction has organizational consequences: whereas Gen AI augments individual tasks with human review of each output, agentic systems autonomously execute multi-step workflows across enterprise systems.

Theoretical Lenses: TOE Framework

TOE framework (Tornatzky & Fleischer, 1990) posits that technology adoption depends on technological characteristics (complexity, compatibility, relative advantage), organizational factors (size, resources, managerial support), and environmental conditions (regulatory pressures, competitive dynamics). TOE has been extensively applied to AI adoption contexts (Chatterjee et al., 2021; Dwivedi et al., 2023).
TOE's organizational focus omits individual psychological factors a recognized limitation when technologies require behavioral change (Oliveira & Martins, 2011; Baker, 2012). AI systems requiring delegation and trust calibration necessitate individual acceptance beyond organizational decisions (Glikson & Woolley, 2020). This has motivated integrated frameworks combining TOE with individual constructs (Alshamaila et al., 2013; Gangwar et al., 2015).

Technology Acceptance and Individual-Level Determinants

Technology Acceptance Models, particularly TAM (Davis, 1989) and its extension Unified Theory of Acceptance and Use of Technology (UTAUT) (Venkatesh et al., 2003), emphasize perceived usefulness, perceived ease of use, social influence, and facilitating conditions as determinants of individual adoption. In organizational AI settings, these constructs explain employees' willingness to use AI tools, their trust in AI outputs, and their perceptions of risk. Decades of research demonstrate that even when organizations deploy technologies, individual acceptance mediates actual usage and value realization
(Venkatesh et al., 2003).

However, agentic AI's autonomy fundamentally shifts which acceptance constructs matter most. When AI systems execute multi-step tasks with limited oversight, perceived 'ease of use' becomes less salient than perceived 'controllability' and 'trust in delegation.' Employees must trust not only individual outputs (as with generative AI) but also the process by which autonomous agents pursue goals, select tools, make intermediate decisions, and revise approaches based on environmental feedback (Shavit et al., 2023). This intensified trust requirement creates psychological barriers absent in simpler tool adoption: delegation to autonomous systems requires employees to cede control over process execution while remaining accountable for outcomes as a tension requiring new governance mechanisms (audit trails, override protocols, accountability assignment) to resolve.

Observability emerges as critical mediating construct. When human-in-the-loop interaction allows direct observation of each AI decision, users can calibrate trust through repeated exposure. Human-on-the-loop interaction, by contrast, requires trusting systems whose intermediate steps may not be directly observable until completion. Literature on automation and trust (Lee & See, 2004) suggests that appropriate trust calibration requires transparency into system reasoning, clear communication of uncertainty, and mechanisms for intervention when trust is violated.

Dynamic Capabilities and Orchestration

DC theory (Teece et al., 1997) highlights sensing, seizing, and transforming as core mechanisms through which organizations create value from digital technologies. Sensing involves identifying opportunities and threats through environmental scanning and interpretation. Seizing involves mobilizing resources, making investment decisions, and committing to strategic directions. Transforming involves reconfiguring organizational structures, processes, and asset portfolios to align with new opportunities. This framework has been extensively applied to understand how organizations leverage emerging technologies including cloud computing, big data analytics, and artificial intelligence (Mikalef & Gupta, 2021).

In multi-paradigm AI contexts, however, organizations face new challenge: coordinating heterogeneous AI artifacts with different autonomy levels, failure modes, and governance requirements. Rather than deploying a single monolithic technology (as in classic ERP adoption), organizations now manage ecosystems of AI components: generative models for content creation, agentic systems for workflow automation, specialized models for domain tasks, APIs connecting components, and legacy infrastructure requiring integration. This portfolio-level coordination activity represents novel organizational challenge not fully captured by standard DC formulation.

Integration: Configurational Perspective

TOE contingencies and DC mechanisms suggest configurational view: adoption outcomes depend on combinations of conditions rather than universal best practices or additive effects (Fiss, 2011; Misangyi et al., 2017). Configurational perspectives recognize necessary conditions (required but insufficient alone), equifinality (multiple pathways to same outcome), and causal asymmetry (success conditions differ from failure conditions).

We adopt this logic as interpretive stance for organizing findings into ideal-type pathways and making implicit configurational logic in scholarly discourse explicit.

Methodology

This study employs a systematic bibliometric and interpretive review (Zupic & Čater, 2015; Donthu et al., 2021), combining quantitative citation analysis with qualitative content examination. Our methodological design maps intellectual structures, temporal evolution patterns, and literature-derived organizational pathways. This approach is deliberately anchored at the level of scholarly discourse rather than direct organizational observation. Analytically, our reasoning is abductive: the intellectual schools, themes, pathways, and contingency factors were first allowed to emerge inductively from reference co-citation and bibliographic-coupling patterns, and were then interpreted through established IS adoption theory—the TOE framework (Tornatzky & Fleischer, 1990), TAM/UTAUT (Davis, 1989; Venkatesh et al., 2003), and dynamic capabilities (Teece et al., 1997)—to articulate their theoretical significance. The themes are thus not deduced from theory a priori; rather, theory provides the interpretive lens through which the inductively derived structure is read.

This choice is theoretically grounded. We adopt Swanson and Ramiller's (1997) concept of the organizing vision: the community-level discourse forming around an emerging technology constitutes a legitimate and epistemically valid signal of organizational phenomena in formation. Bibliometric analysis of an emergent corpus is therefore not a methodological limitation. It is a strategically appropriate design for studying fields in statu nascendi, where conference proceedings and emerging journals are precisely where the scientific community works through foundational questions before they crystallize in top-tier outlets.

Data Collection and Corpus Construction

we chose the ARTIREV (ARTificial Intelligence and Literature REView) software which seeks to assist in conducting literature review based on a combination RCCA) and document bibliographic coupling analysis DBCA (Walsh et al., 2021) combining bibliometric techniques with qualitative analysis capabilities.. Literature management and bibliometric analysis were performed using ArtiRev, a specialized tool.
Our systematic literature search targeted peer-reviewed publications addressing organizational implementation of generative and/or agentic AI systems. We queried Scopus databases (January 2023–February 2026) using controlled vocabulary and Boolean operators:

Generative AI: ("generative AI" OR "generative artificial intelligence" OR "GenAI" OR "GPT" OR "large language model" OR "LLM") AND ("organi?ation*" OR "enterprise" OR "business" OR "firm")
Agentic AI: ("agentic AI" OR "AI agent*" OR "autonomous agent*" OR "intelligent agent*") AND ("organi?ation*" OR "enterprise" OR "business" OR "firm")
Inclusion criteria: (1) peer-reviewed articles, conference papers, or book chapters; (2) English language; (3) explicit focus on organizational contexts; (4) published 2023–2026. Exclusion criteria: (1) purely technical papers without organizational focus; (2) dissertations and working papers; (3) duplicate publications.

The initial search yielded 3,145 documents. After deduplication and abstract screening, the final corpus comprised 2,789 documents: 2,610 discussing generative AI, 177 discussing agentic AI, with 205 overlap documents addressing both paradigms simultaneously (7.4% of corpus).

Note on 2026 data: Documents indexed for 2026 represent a partial year (January–February 2026). To avoid volume-based bias, all temporal comparisons involving 2026 use normalized ratios (mentions per document) rather than absolute publication counts, ensuring comparability across years.

Transition Document Analysis and Pathway Identification

To identify transition documents addressing both AI paradigms, we extracted corpus metadata from ArtiRev and applied computational filtering using Boolean pattern matching. Documents qualified if titles or abstracts contained terminology from both paradigm categories simultaneously: generative AI indicators (generative|genai|gpt|llm) AND agentic AI indicators (agentic|ai agent|autonomous ai). This identified 205 documents (7,4% of corpus), growing from 1.3% of publications (2023) to 4.5% (2026), signaling increasing scholarly integration of both paradigms.

To validate filtering accuracy, we manually inspected a random 10% subsample (n = 21), confirming 100% precision (all identified documents genuinely discussed both paradigms). Recall was assessed through independent review of 50 randomly selected non-transition documents, identifying zero false negatives, suggesting high recall.

The 205 transition documents underwent systematic thematic content analysis (Krippendorff, 2004). Two researchers independently coded a random sample of 56 documents (27% of the transition corpus) using a structured codebook designed to identify organizational transition pathways. The protocol captured whether authors described the transition as a sequence, a parallel effort, or a direct technological leap. Inter-coder reliability was assessed using Cohen's Kappa ($\kappa = 0.82$), indicating almost perfect agreement. Disagreements were resolved through structured discussion before coding the remainder of the corpus. The resulting percentages are reported to indicate the relative prevalence of these narratives within our specific transition corpus, not as generalized statistical probabilities.

Results

This section presents findings from three complementary analyses addressing our research questions. Reference co-citation analysis reveals intellectual structure (RQ1), temporal document bibliographic coupling analysis tracks thematic evolution (RQ2), and interpretive coding of transition documents identifies relationship framings and contingency factors (RQ3).

Intellectual Structure: Five Schools

RCCA analysis of the full corpus identified five intellectual schools structuring organizational Gen AI and Agentic AI research. The AI and ML Technical school dominates with approximately 74 percent of co-citation weight, anchored in foundational works including Vaswani et al. (2017) on transformer architectures, Brown et al. (2020) on large language models, and Devlin et al. (2018) on BERT. This dominance indicates organizational AI research remains strongly influenced by computer science foundations. Technology Acceptance school draws on Davis (1989), Venkatesh et al. (2003), and Ajzen (1991), informing individual adoption constructs including perceived usefulness and behavioral intention. Resource and DC school anchors in Barney (1991) and Teece et al. (1997), framing AI as strategic resource requiring sensing, seizing, and transforming mechanisms. TOE and Organizational Context school applies Tornatzky and Fleischer (1990) framework to IS adoption contexts. Governance and Ethics represents emerging cluster addressing accountability and regulatory concerns through works including Bender et al. (2021) and Bommasani et al. (2021), signaling shift toward socio-technical concerns as adoption matures.

Temporal Evolution of Key Themes

Temporal DBCA tracks normalized theme frequencies from 2023 to 2026 (Table 1). Three patterns emerge. First, Gen AI reaches saturation at 0.95 in 2026, reflecting baseline discourse normalization. Second, Agentic AI exhibits steepest growth trajectory, increasing 492 percent from 0.12 to 0.71, indicating rapid conceptual penetration. Third, orchestration shows largest relative increase at 340 percent, growing from 0.15 to 0.66, signaling scholarly attention shifting toward coordination of multiple AI artifacts and platform management. Governance discourse doubles from 0.31 to 0.62, reflecting intensifying accountability concerns. Integration themes grow 165 percent while replacement narratives decline 33 percent, indicating scholarly consensus favors complementarity over substitution framings.

Table 1. Temporal Thematic Evolution: Normalized Mentions per Document (2023-2026)

Theme	2023	2024	2025	2026	Change
Generative AI	0.89	0.92	0.94	0.95	+7%
Agentic AI	0.12	0.28	0.52	0.71	+492%
Integration	0.27	0.42	0.58	0.72	+165%
Orchestration	0.15	0.32	0.49	0.66	+340%
Governance	0.31	0.45	0.55	0.62	+100%
Risk/Security	0.44	0.48	0.51	0.53	+20%
Human-AI Collab.	0.38	0.48	0.58	0.66	+73%
Replacement	0.18	0.16	0.14	0.12	-33%

Adoption Pathway Configurations

DBCA of the 205 transition documents identified four distinct clusters characterized by different bibliographic coupling profiles. Coupling weights were computed using shared reference counts normalized by document length. Cluster analysis was performed using the Louvain modularity algorithm (resolution parameter = 0.8; minimum cluster size = 3 documents), yielding a modularity score of Q = 0.61. This score reflects an emerging but structurally coherent multi-paradigm AI discourse at this stage of field emergence and should be interpreted accordingly. The four clusters and their relative sizes are reported in Table 2.
The first cluster, Gen-Only Augmentation (n=62, 30%), groups documents sharing dense bibliographic coupling around productivity augmentation literature, human-in-the-loop workflows, and prompt engineering references. Documents in this cluster cite primarily TAM and TOE foundations combined with output-quality evaluation frameworks, with minimal co-citation of autonomous systems or governance literature.

Table 2. Adoption Pathway Configurations — DBCA Clusters (n=205)

Pathway	n	%	Dominant Bibliographic Profile	Modularity Contribution
Gen-Only Augmentation	62	30%	TAM/TOE + productivity + HCI literature	High internal cohesion
Integrated Portfolios	49	24%	DC + multi-agent + governance architecture	Highest inter-cluster separation
Agentic-Focused Experimentation	41	20%	Risk management + sandboxing + agentic technical	Moderate cohesion
Gen-Persistent Constraint	53	26%	Adoption barriers + capability gaps + governance deficits	Moderate cohesion
Total	205	100%		Q = 0.61

The second cluster, Integrated Portfolios (n=49, 24%), exhibits strong coupling around multi-agent systems architecture, orchestration platform, and AI governance references. Documents in this cluster share dense citation networks combining DC theory with technical coordination literature, indicating an intellectually mature sub-community addressing portfolio-level coordination challenges.

The third cluster, Agentic-Focused Experimentation (n=41, 20%), groups documents coupled around sandboxing methodologies, controlled deployment protocols, and risk management frameworks. The bibliographic profile combines agentic AI technical literature with organizational risk and learning references, suggesting a distinct intellectual community addressing constrained autonomous deployment.
The fourth cluster, Gen-Persistent Constraint (n=53, 26%), groups documents coupled around adoption barrier literature, readiness deficit frameworks, and failed implementation references. This cluster exhibits the densest coupling with governance gap and organizational capability literature, reflecting scholarly attention to the conditions preventing agentic deployment despite strategic intent.

Table 3 reports the distribution of the four clusters across the observation window. The pathway configurations derived from these clusters are interpreted theoretically in the following section within a four-block contingency framework.

The percentages reported above (30%, 24%, 20%, 26%) reflect the distribution of documents across DBCA clusters within this specific corpus of 205 transition documents. They are descriptive characterizations of scholarly discourse proportions rather than statistically stable probability estimates. Given the corpus size and moderate modularity score (Q=0. 61), these figures should be interpreted as indicative of relative discourse emphasis rather than precise population parameters.

Relationship Framings in Transition Documents

Coding 205 transition documents describing explicitly both paradigm reveals how literature frames Gen-Agentic relationship (Table 3). Integration and complementarity emerges as dominant framing, appearing in 68 percent of documents describing paradigms as complementary portfolio components where generative models provide reasoning while agentic layers orchestrate execution. Orchestration and platform framing appears in 26 percent, proposing meta-AI infrastructures coordinating heterogeneous components through routing logic, monitoring dashboards, and policy enforcement. Evolution and sequential framing appears in 21 percent, suggesting organizations progress from Gen-only experimentation toward agentic architectures as digital maturity increases. Replacement narratives remain rare at 9 percent, contrasting

sharply with integration dominance. Parallel and coexistence framing appears in 13 percent, describing disconnected systems serving distinct use cases. This distribution indicates scholarly discourse positions Gen and Agentic AI as synergistic rather than substitutive technologies.

Table 3. Relationship Framings in Transition Documents (n=205)

Framing	n	%	Description
Integration/Complementarity	140	68%	Gen reasoning + Agentic execution
Orchestration/Platform	53	26%	Meta-AI coordinates components
Evolution/Sequential	43	21%	Gen pilots → Agentic maturity
Replacement/Substitution	18	9%	Agentic supersedes Gen
Parallel/Coexistence	27	13%	Disconnected systems

Contingency Factors Across Four Blocks

Transition documents emphasize contingency factors organized into four blocks (Table 4). Organizational factors appear most frequently: digital maturity in 77 percent of documents (n=158), governance readiness in 74 percent (n=152), orchestration infrastructure in 46 percent (n=94) despite being emerging construct, and resource availability in 42 percent (n=86). Technological factors include use-case complexity cited in 73 percent of documents (n=150), compatibility and integration requirements in 63 percent (n=129), and observability of AI behavior in 55 percent (n=113). Individual factors appear prominently: AI literacy in 68 percent (n=139), trust in autonomous systems in 62 percent (n=127), and resistance to delegation in 48 percent (n=98). Institutional factors show sector-specific patterns: regulatory pressure appears in 57 percent (n=117) concentrated heavily in finance and healthcare discussions, professional norms in 41 percent (n=84), and industry standards in 32 percent (n=66). This distribution validates four-block contingency structure with all blocks receiving substantial scholarly attention, though organizational factors are most emphasized reflecting their gating role for agentic deployment beyond sandboxes.

Table 4. Contingency Factors Mentioned in Transition Documents (n=205)

Factor Category	Specific Factors	Documents	%
Technological	Use-case complexity	150	73%
	Compatibility/Integration	129	63%
	Observability of AI behavior	113	55%
Organizational	Digital maturity	158	77%
	Governance readiness	152	74%
	Orchestration infrastructure	94	46%
	Resource availability	86	42%
Individual	AI literacy	139	68%
	Trust in autonomous systems	127	62%
	Resistance to delegation	98	48%
Institutional	Regulatory pressure	117	57%
	Professional norms	84	41%
	Industry standards	66	32%

Discussion

Theoretical contributions

Read through the theoretical lenses introduced in Section 2, the four contingency blocks map onto distinct constructs, making explicit how the inductively derived structure relates to established IS adoption theory. The technological block (use-case complexity, compatibility and integration, observability of AI behavior) corresponds to the technology dimension of the TOE framework (Tornatzky & Fleischer, 1990). The organizational block (digital maturity, resource availability) corresponds to the organization dimension of

TOE; two of its factors—orchestration infrastructure and governance readiness—are, however, better read through dynamic capabilities (Teece et al., 1997), with orchestration functioning as a sensing–seizing–transforming coordination capability for heterogeneous AI portfolios and governance readiness as a transforming/reconfiguration prerequisite for delegating autonomy. The individual block (AI literacy, trust in autonomous systems, resistance to delegation) corresponds to TAM/UTAUT acceptance constructs (Davis, 1989; Venkatesh et al., 2003), here reframed for delegation under the human-in-the-loop to human-on-the-loop shift, where controllability and trust in the execution process supersede perceived ease of use. The institutional block (regulatory pressure, professional norms, industry standards) corresponds to the environment dimension of TOE. This mapping clarifies that orchestration and governance readiness are precisely the constructs least accommodated by classical adoption models, and therefore the most theoretically generative for the agentic transition.

This structure systematizes the discourse-derived pathways by revealing underlying theoretical dimensions distinguishing them. It enables systematic hypothesis generation across technological, organizational, individual, and institutional conditions. This multi-dimensional structure supports configurational hypothesis development examining how combinations of conditions rather than individual factors in isolation shape adoption pathways, though empirical validation of such configurational relationships requires organizational data.

Pattern 1: Individual Determinants Beyond TOE Framework

Transition documents consistently emphasize AI literacy, trust, and resistance as adoption mediators appearing in 68 percent of coded documents. This pattern reveals determinants beyond classical TOE framework: Tornatzky and Fleischer (1990) conceptualize adoption as function of technological characteristics, organizational readiness, and environmental pressures, explicitly excluding individual-level psychological factors. Technology Acceptance Model and its extensions address individual acceptance through perceived usefulness and ease of use (Davis, 1989; Venkatesh et al., 2003), but assume users evaluate discrete outputs before acting.

Agentic AI presents different challenge. Delegation requires trusting processes, not just outputs—accepting that systems pursue goals through action sequences not fully observable (Glikson & Woolley, 2020; Lee & See, 2004). Three individual-level constructs appear particularly salient: AI literacy enabling trust calibration (Ng et al., 2021), controllability perceptions determining delegation thresholds (Dietvorst et al., 2015), and resistance from job threat or autonomy concerns (Kellogg et al., 2020). Scholarly discourse frames these factors as critical determinants, suggesting they may constitute microfoundations of organizational sensing, seizing, and transforming capabilities (Teece, 2007). These factors may exhibit properties consistent with necessary conditions or equifinal configurations (Fiss, 2011) hypotheses warranting organizational validation. This discourse pattern motivates empirical investigation testing whether individual factors explain unique variance beyond organizational context or mediate organizational readiness-adoption relationships (Baird & Maruping, 2021).

Pattern 2: Orchestration as Coordination Mechanism

Orchestration exhibits the largest relative growth at 340 percent in normalized mentions from 2023 to 2026 (from 15% to 66% of transition documents, representing absolute growth from ~31 to ~135 documents, emerging as central construct in multi-paradigm AI discourse. This pattern raises theoretical question for DC research: what role does orchestration play in enabling sensing, seizing, and transforming mechanisms when organizations simultaneously deploy heterogeneous AI artifacts with different autonomy levels (Teece, 2007; Eisenhardt & Martin, 2000)?

Our bibliometric evidence documents orchestration discourse emphasizing model registries, routing logic, monitoring dashboards, and policy enforcement technical infrastructure coordinating heterogeneous AI artifacts. Drawing on coordination theory (Malone & Crowston, 1994; Okhuysen & Bechky, 2009), we interpret orchestration as coordination mechanism enabling DC deployment rather than distinct capability itself. This interpretation aligns with coordination theory's focus on interdependence management, but alternative framings treating orchestration as distinct meta-capability (Helfat & Peteraf, 2015) remain theoretically viable. Our bibliometric evidence cannot adjudicate between these interpretations. Orchestration infrastructure supports sensing by revealing which AI artifacts perform effectively for which tasks, supports seizing by enabling systematic deployment of appropriate models, and supports transforming by providing version control and governance integration enabling workflow redesign (Sambamurthy et al., 2003; Helfat et al., 2007). Organizations lacking this coordination infrastructure cannot effectively manage multiple AI paradigms, constraining them to simpler single-paradigm deployments.

Pattern 3: Governance as Necessary Condition

Transition documents consistently position governance readiness as prerequisite for agentic deployment beyond sandboxes rather than continuous enabler, suggesting threshold effect. Governance frameworks addressing autonomous AI systems including observability mechanisms, accountability structures, and scope delimitation have emerged as critical organizational concern (Weill & Ross, 2004; Abraham et al., 2019; Birkstedt et al., 2023).

This discourse pattern challenges dominant theoretical assumptions. Traditional variance-based approaches treat governance as continuous predictor (Venkatraman, 1989), while TOE logic assumes additive effects where organizational, technological, and environmental factors combine linearly. Scholarly discourse suggests different logic: governance may constitute organizational meta-capability whose absence prevents sensing AI risks, seizing agentic opportunities, or transforming workflows regardless of technical expertise or market pressures (Helfat & Peteraf, 2015; Teece, 2007). Transition documents frame organizations lacking governance infrastructure as unable to coordinate multi-paradigm portfolios (Pattern 2) or support individual trust calibration (Pattern 1), constraining them to simpler deployments where governance requirements remain minimal. This discourse pattern aligns with configurational thinking (Fiss, 2011; Dul, 2016): governance may function as necessary condition whose absence prevents agentic adoption regardless of favorable orchestration, maturity, or use-case configurations. Above threshold, other factor combinations determine outcomes. Testing this necessary condition hypothesis requires organizational data enabling set-theoretic analysis rather than variance-based methods.

This discourse-derived hypothesis suggests causal asymmetry (Ragin, 2008; Misangyi et al., 2017): governance absence may predict adoption failure while governance presence does not alone predict success. Empirical validation of this asymmetric relationship requires organizational data enabling set-theoretic analysis rather than symmetric regression models assuming linear additive effects.

These three patterns individual determinants, orchestration infrastructure, and governance readiness emerge from analysis of scholarly discourse rather than organizational data. Our contribution lies in documenting how IS scholarship conceptualizes the Gen-to-Agentic transition and identifying theoretical gaps these discourse patterns reveal. The theoretical interpretations proposed above micro foundations, coordination mechanisms, necessary condition logic represent hypotheses derived from bibliometric analysis requiring validation through organizational research. Where we propose configurational interpretations (necessary

conditions, causal asymmetry, equifinality), these represent discourse-based theoretical hypotheses rather than empirically validated relationships.

Contingency Framework as Theoretical Synthesis

The three patterns above emerge from thematic analysis of discourse content. To systematize adoption approaches documented in scholarly discourse, the DBCA revealed four discourse-grounded pathways characterized by different technological, organizational, individual, and institutional configurations. We develop contingency framework organizing these pathways across theoretical dimensions.

Four-Block Framework: Foundation and Justification

Each block addresses a distinct analytical level with independent causal mechanisms. The Individual block captures cognitive and affective factors: trust calibration, AI literacy, and controllability perceptions that operate at the individual level and mediate organizational capability effects (Venkatesh et al., 2003; Baird and Maruping, 2021); organizations with high governance readiness but low individual trust exhibit resistance patterns not predicted by organizational factors alone (Kellogg et al., 2020; Pattern 1). Individual-level factors explain unique variance beyond organizational context when technologies require behavioral delegation rather than simple tool use (Glikson and Woolley, 2020; Lee and See, 2004). The Technological block addresses intrinsic task properties: use-case complexity, task decomposability, and workflow integration requirements, independent of organizational capabilities (Tornatzky and Fleischer, 1990); complex multi-step workflows require agentic capabilities regardless of organizational readiness, consistent with TOE's proposition that technological characteristics constitute independent adoption determinants (Chatterjee et al., 2021). The Organizational block integrates developed capabilities digital maturity, orchestration infrastructure, and governance readiness that emerge from deliberate capability building over time (Teece et al., 1997; Weill and Ross, 2004; Patterns 2–3); these operate at organizational architecture level, distinct from individual cognitions or external institutional pressures, and their effects on adoption are mediated by sensing, seizing, and transforming mechanisms (Eisenhardt and Martin, 2000; Mikalef and Gupta, 2021). The Institutional block captures field-level forces: regulatory pressures, compliance requirements, and industry norms that constrain or enable adoption independently of internal organizational capabilities (DiMaggio and Powell, 1983); regulated sectors exhibit systematically different adoption patterns than non-regulated sectors, controlling for organizational readiness, and coercive isomorphism shapes technology adoption trajectories in AI-intensive contexts (Birkstedt et al., 2023; Dwivedi et al., 2023).

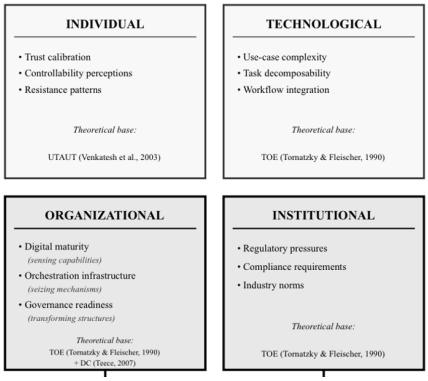


Figure 1. Contingency framework

Each block exhibits discriminant validity: technological conditions distinguish Gen-Only (modular tasks) from Integrated (complex workflows); organizational conditions distinguish Integrated (high maturity, orchestration, established governance) from Gen-Persistent Constraint (capability gaps, emerging

governance); individual conditions distinguish pathways by trust levels and resistance patterns; institutional conditions capture regulatory constraints varying across deployment contexts.

Table 5. Adoption Pathway Configurations

Pathway	Technological Conditions	Organizational Conditions	Individual Conditions	Institutional Conditions	Dominant Discourse Themes	n	%
Gen-Only Augmentation	Modular, decomposable tasks	Moderate digital maturity Emerging governance No orchestration platform	Moderate trust High controllability needs	Low regulatory pressure	Productivity tools Human-in-loop No agency granted	62	30%
Integrated Portfolios	Complex, multi-step workflows	High digital maturity Established governance Orchestration infrastructure	High trust Calibrated controllability	Moderate regulatory pressure	Portfolio coordination Multi-agent systems Governance architectures	49	24%
Agentic-Focused Experimentation	High complexity tasks Bounded problem scope	Moderate maturity Strong risk management Sandbox constraints	High risk tolerance Learning orientation	Variable regulatory context	Constrained learning Rapid iteration Explicit scope limits	41	20%
Gen-Persistent Constraint	Articulated complex goals	Low governance readiness Limited orchestration Capability gaps	Low trust High resistance Autonomy concerns	Variable regulatory context	Adoption barriers Failed pilots Readiness deficits	53	26%

Practical Implications

Scholarly discourse patterns suggest four practice-relevant considerations for organizations navigating the Gen-to-Agentic transition. These insights derive from bibliometric analysis of peer-reviewed literature rather than systematic observation of organizational deployments. Empirical validation with organizational data is required before prescriptive application, as discourse patterns may not fully reflect organizational practice.

First, governance readiness assessment. Discourse patterns position governance infrastructure as potential prerequisite for agentic deployment beyond sandboxes (Pattern 3). Organizations may benefit from assessing governance readiness using the four-block framework (Table 3) before initiating agentic pilots. Unlike traditional IT adoption where technical capability often drives deployment decisions, agentic AI discourse emphasizes governance foundations including established audit trails, accountability assignment mechanisms, and policy enforcement infrastructure. Organizations lacking these mechanisms might consider investing in governance frameworks before granting decision-making autonomy to AI systems. This governance-first sequencing reflects agentic AI's unique challenge of delegating consequential actions to stochastically determined systems, though the optimal sequencing remains an empirical question requiring organizational validation.

Second, orchestration infrastructure for portfolio coordination. Organizations simultaneously deploying generative models and autonomous agents face coordination challenges absent in single-paradigm

deployments. Scholarly discourse emphasizes orchestration mechanisms (Pattern 2) addressing which AI artifact handles which task, how systems route requests across heterogeneous components, and how organizations monitor performance across paradigms. Literature patterns suggest implementing model registries, routing logic, and monitoring dashboards early in adoption trajectories may reduce transition friction when scaling from Gen-only to integrated portfolios. This differs from traditional enterprise architecture where system boundaries remain relatively stable; AI portfolio coordination may require dynamic mechanisms adapting to evolving model capabilities.

Third, delegation calibration training. Agentic AI introduces human-on-the-loop paradigms where employees monitor autonomous execution rather than directly controlling outputs. Discourse patterns suggest training programs might shift from prompt engineering skills toward delegation calibration competencies. Organizations may benefit from tracking individual-level metrics including trust calibration appropriateness, override frequency patterns, and delegation threshold effectiveness alongside system-level performance indicators. Unlike traditional technology acceptance focusing on perceived ease of use, agentic adoption discourse emphasizes developing appropriate skepticism: trusting outputs sufficiently to delegate while maintaining vigilant monitoring capacity.

Fourth, graduated autonomy progression. Literature patterns suggest staged deployment approaches progressing from sandbox environments through human-approval gates to monitored autonomy rather than binary on/off activation. Traditional staged rollouts test technical stability; agentic staged rollouts may test organizational governance capacity to detect and respond to autonomous system failures before consequences cascade. Organizations might consider tracking audit trail coverage completeness, incident detection latency, and escalation protocol effectiveness as leading indicators of governance readiness for expanded autonomy scope.

These practice-oriented insights require empirical validation. The extent to which discourse patterns reflect organizational reality, the causal mechanisms underlying suggested relationships, and the comparative effectiveness of alternative approaches all constitute open empirical questions addressed in the future research agenda (Section 5.4).

Limitations and research boundaries

Our bibliometric design has four inherent limitations bounding scope and interpretation.
First, discourse–reality gap and publication bias. We analyze scholarly discourse, not organizational behavior; the four adoption pathways represent intellectual patterns within the literature (2023–2026), not validated descriptions of organizational deployments. Because academic incentives favor novelty and successful adoptions over failed initiatives, the discourse likely undercounts governance failures and organizations that stall before agentic deployment (field studies suggest 40–60 percent of AI initiatives fail to progress beyond pilots). Empirical validation with organizational data is required before prescriptive application.

Second, pathway distribution and clustering uncertainty. The pathway percentages (Gen-Only 30%, Integrated 24%, Agentic-Focused 20%, Gen-Persistent Constraint 26%) reflect the corpus distribution, not organizational base rates, and carry inherent clustering uncertainty given the corpus size and moderate modularity ($Q = 0.61$). Successful deployments may be underrepresented due to competitive sensitivity and experimental implementations overrepresented due to novelty bias; sensitivity to clustering parameters cannot be fully excluded. Future work on larger multi-paradigm corpora should assess whether these configurations replicate.

Third, corpus boundaries. Our Scopus focus reflects Anglophone IS-centric scholarship, excluding practitioner reports, non-English academic work, and grey literature. The 2,789-document corpus provides comprehensive coverage of indexed academic discourse but incomplete coverage of organizational practice.

Fourth, no causal testing. Our design identifies associations within scholarly discourse rather than causal relationships in organizations. We cannot establish temporal precedence or causal mechanisms. The propositions in the following section translate patterns into empirically testable predictions requiring longitudinal organizational data.

These limitations are inherent to bibliometric literature reviews. This research establishes field infrastructure through intellectual structure revelation, thematic evolution quantification, and testable proposition derivation, enabling transition from discourse synthesis to organizational hypothesis testing.

Future Research Agenda

Grounded in DC and IS Governance theory, we derive two propositions translating discourse patterns into testable predictions (Table 5).

Proposition 1: Orchestration Infrastructure and Portfolio Integration. Pattern 2 positions orchestration infrastructure as coordination mechanism enabling sensing, seizing, and transforming in multi-paradigm AI contexts (Teece, 2007). This predicts organizations with established orchestration infrastructure will exhibit Integrated Portfolio adoption more frequently than Gen-Only patterns, controlling for use-case complexity and regulatory context. Multinomial logistic regression with organizational survey data (n=300+) would test this relationship.

Proposition 2: Governance Readiness as Necessary Condition. Pattern 3 positions governance readiness as threshold enabling agentic deployment (Weill & Ross, 2004). Governance infrastructure provides observability, enforceability, and accountability prerequisites for autonomous agent delegation. This predicts governance readiness appears as necessary condition present in all Integrated and Agentic-Focused cases but absent in Gen-Only and Gen-Persistent Constraint cases. Fuzzy-set QCA necessary condition analysis (n=40-60) would test consistency ≥0.90 across cases.

Pattern 1 Integration. Pattern 1 (individual determinants) is operationalized through the Individual block but not formalized as standalone proposition pending empirical validation. Individual factors: trust calibration, AI literacy, resistance patterns, may function as moderators or mediators of organizational readiness effects on adoption pathways. Testing these multilevel relationships requires nested data structures (individuals within organizations) and represents a priority direction for future research.

Conclusion

This bibliometric synthesis documents a rapid intellectual shift in organizational AI research: from technical foundations toward integration, governance, and orchestration concerns. Scholarly discourse suggests adoption pathways are associated with configurations of organizational capabilities (digital maturity, governance readiness), technical attributes (use-case complexity, task decomposability), and individual readiness (trust calibration, AI literacy). Orchestration emerges as coordination mechanism enabling organizations to manage multi-paradigm AI portfolios. Our structured contingency framework synthesizes bibliometric and coded evidence into typified configurations linking readiness factors to adoption pathways. The framework specifies which combinations of conditions favor Gen-Only, Integrated, Agentic-Focused,

or Gen-Persistent Constraint pathways. Two testable propositions provide foundation for empirical validation, linking theoretical mechanisms to falsifiable predictions.
This foundation is critical as organizations navigate the Gen-to-Agentic transition during a pre-paradigmatic period where scholarly frameworks are emerging contemporaneously with deployment decisions. The framework offers practitioners a discourse-grounded readiness assessment tool for diagnosing organizational configurations and identifying capability gaps, though these remain theory-building constructs pending organizational validation.

We acknowledge significant limitations Anglophone corpus bias, discourse-reality gap, publication bias favoring successful deployments, and C3 structural immaturity. Section 5.4 proposes empirical validation through multinomial regression testing orchestration-integration links (n=300+), fuzzy-set QCA testing governance as necessary condition (n=40-60), and complementary longitudinal case studies (n=8-12) tracing micromechanisms over time.

We call for organizational researchers to test these propositions with longitudinal data establishing temporal precedence, and for practitioners to use the framework as diagnostic instrument requiring context-specific calibration rather than prescriptive blueprint.

Acknowledgements

The author used Claude (Anthropic), a generative AI assistant, during the preparation of this manuscript for language refinement, editing, and formatting assistance. The author reviewed and verified all content and takes full responsibility for the originality, validity, and integrity of the work, in accordance with the IIS AI Use Policy (https://iacis.org/iis/AI_Policy.php).

References

Abraham, R., Schneider, J., & vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. International Journal of Information Management, 49, 424–438. https://doi.org/10.1016/j.ijinfomgt.2019.07.008

Ajzen, I. (1991). The theory of planned behavior. Organizational Behavior and Human Decision Processes, 50(2), 179-211. https://doi.org/10.1016/0749-5978(91)90020-T

Baird, A., & Maruping, L. M. (2021). The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. MIS Quarterly, 45(1), 315-341.

Baker, J. (2012). The technology-organization-environment framework. In Y. K. Dwivedi, M. R. Wade, & S. L. Schneberger (Eds.), Information systems theory (pp. 231-245). Springer.

Barney, J. (1991). Firm resources and sustained competitive advantage. Journal of Management, 17(1), 99-120. https://doi.org/10.1177/014920639101700108

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT). https://doi.org/10.1145/3442188.3445922

Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. Internet Research, 33(7), 133-167. https://doi.org/10.1108/INTR-01-2022-0042

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33 (NeurIPS 2020), 1877-1901.

Chatterjee, S., Rana, N. P., Dwivedi, Y. K., & Baabdullah, A. M. (2021). Understanding AI adoption in manufacturing and production firms using an integrated TAM-TOE model. Technological Forecasting and Social Change, 170, 120880. https://doi.org/10.1016/j.techfore.2021.120880

Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zemmel, R. (2024). The economic potential of generative AI: The next productivity frontier. McKinsey Global Institute.

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Quarterly, 13(3), 319-340. https://doi.org/10.2307/249008

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 4171-4186.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. Journal of Experimental Psychology: General, 144(1), 114-126. https://doi.org/10.1037/xge0000033

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. American Sociological Review, 48(2), 147-160. https://doi.org/10.2307/2095101

Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. Journal of Business Research, 133, 285-296. https://doi.org/10.1016/j.jbusres.2021.04.070

Dul, J. (2016). Necessary condition analysis (NCA): Logic and methodology of "necessary but not sufficient" causality. Organizational Research Methods, 19(1), 10-52. https://doi.org/10.1177/1094428115584005

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., ... & Wright, R. (2023). Opinion Paper: 'So what if ChatGPT wrote it?' Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. International Journal of Information Management, 71, 102642. https://doi.org/10.1016/j.ijinfomgt.2023.102642

Eisenhardt, K. M., & Martin, J. A. (2000). Dynamic capabilities: What are they? Strategic Management Journal, 21(10-11), 1105-1121.

Fiss, P. C. (2011). Building better causal theories: A fuzzy set approach to typologies in organization research. Academy of Management Journal, 54(2), 393-420. https://doi.org/10.5465/amj.2011.60263120

Gangwar, H., Date, H., & Ramaswamy, R. (2015). Understanding determinants of cloud computing adoption using an integrated TAM-TOE model. Journal of Enterprise Information Management, 28(1), 107-130.

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. Academy of Management Annals, 14(2), 627-660. https://doi.org/10.5465/annals.2018.0057

Helfat, C. E., Finkelstein, S., Mitchell, W., Peteraf, M. A., Singh, H., Teece, D. J., & Winter, S. G. (2007). Dynamic capabilities: Understanding strategic change in organizations. Blackwell.

Helfat, C. E., & Peteraf, M. A. (2015). Managerial cognitive capabilities and the microfoundations of dynamic capabilities. Strategic Management Journal, 36(6), 831-850. https://doi.org/10.1002/smj.2247

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

Krippendorff, K. (2004). Content analysis: An introduction to its methodology (2nd ed.). Sage.

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46(1), 50-80. https://doi.org/10.1518/hfes.46.1.50_30392

Malone, T. W., & Crowston, K. (1994). The interdisciplinary study of coordination. ACM Computing Surveys, 26(1), 87-119. https://doi.org/10.1145/174666.174668

Masterman, T., Besen, S., Sawtell, M., & Chao, A. (2024). The landscape of emerging AI agent architectures for reasoning, planning, and tool calling: A survey. arXiv preprint arXiv:2404.11584.

Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. Information & Management, 58(3), 103434. https://doi.org/10.1016/j.im.2021.103434

Misangyi, V. F., Greckhamer, T., Furnari, S., Fiss, P. C., Crilly, D., & Aguilera, R. (2017). Embracing causal complexity: The emergence of a neo-configurational perspective. Journal of Management, 43(1), 255-282. https://doi.org/10.1177/0149206316679252

Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. Computers and Education: Artificial Intelligence, 2, 100041. https://doi.org/10.1016/j.caeai.2021.100041

Okhuysen, G. A., & Bechky, B. A. (2009). Coordination in organizations: An integrative perspective. Academy of Management Annals, 3(1), 463-502. https://doi.org/10.5465/19416520903047533

Oliveira, T., & Martins, M. F. (2011). Literature review of information technology adoption models at firm level. Electronic Journal of Information Systems Evaluation, 14(1), 110-121.

Ragin, C. C. (2008). Redesigning Social Inquiry: Fuzzy Sets and Beyond. University of Chicago Press. https://doi.org/10.7208/chicago/9780226702797.001.0001

Sambamurthy, V., Bharadwaj, A., & Grover, V. (2003). Shaping agility through digital options: Reconceptualizing the role of information technology in contemporary firms. MIS Quarterly, 27(2), 237-263. https://doi.org/10.2307/30036530

Shavit, Y., Agarwal, S., Brundage, M., Adler, S., O'Keefe, C., Campbell, R., ... & Robinson, D. G. (2023). Practices for governing agentic AI systems. OpenAI Research Paper.

Swanson, E. B., & Ramiller, N. C. (1997). The organizing vision in information systems innovation. Organization Science, 8(5), 458-474. https://doi.org/10.1287/orsc.8.5.458

Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. Strategic Management Journal, 28(13), 1319-1350. https://doi.org/10.1002/smj.640

Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. Strategic Management Journal, 18(7), 509-533.

Tornatzky, L. G., & Fleischer, M. (1990). The Processes of Technological Innovation. Lexington Books.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30 (NIPS 2017), 5998-6008.

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. MIS Quarterly, 27(3), 425-478. https://doi.org/10.2307/30036540

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

Walsh, I., & Rowe, F. (2022). BIBGT: Combining bibliometrics and theory-driven coding for systematic literature reviews. Communications of the Association for Information Systems, 50, Article 25. https://doi.org/10.17705/1CAIS.05025

Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., ... & Wen, J. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science, 18(6), 186345. https://doi.org/10.1007/s11704-024-40231-1

Weill, P., & Ross, J. W. (2004). IT Governance: How Top Performers Manage IT Decision Rights for Superior Results. Harvard Business School Press.

Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. Organizational Research Methods, 18(3), 429-472. https://doi.org/10.1177/1094428114562629