Multidisciplinary Team Architecture in AI Centers of Excellence: Integrating Data Science, Business Strategy, and Governance Expertise

Shuokai Wang[1], Bo Zeng[1,*]

[1] School of Economics and Management, Beijing Forestry University, Beijing, China 100091

* Email: zengbo@bjfu.edu.cn (Corresponding Author)

Received 10 June 2025 · Revised 25 July 2025 · Accepted 8 September 2025 · Published 30 December 2025

Abstract

As organizations accelerate artificial intelligence (AI) adoption, the AI Center of Excellence (CoE) has emerged as a pivotal organizational mechanism for centralizing expertise, standardizing practices, and scaling AI initiatives across enterprise functions. A defining characteristic of effective AI CoEs is their reliance on multidisciplinary teams comprising data scientists, machine learning engineers, business strategists, governance specialists, and domain experts. However, empirical evidence documenting the precise composition, integration patterns, and performance implications of such teams remains sparse. This study addresses this gap through a mixed-methods investigation combining a quantitative survey of 186 organizations across twelve industries with semi-structured interviews of 34 senior AI leaders. The findings reveal that higher levels of disciplinary diversity within AI CoEs are positively associated with project success rates ($r = 0.68$, $p < 0.001$) and stakeholder satisfaction ($r = 0.71$, $p < 0.001$). Organizations employing hub-and-spoke or hybrid operating models demonstrate superior cross-functional integration relative to centralized structures. Governance maturity, assessed through a five-level capability framework, emerges as the strongest single predictor of AI deployment scalability. The study further identifies four archetypal team configurations and maps their alignment with organizational AI maturity stages. These results substantiate the claim that the AI CoE is typically composed of multidisciplinary experts, including data scientists, engineers, business strategists, and governance specialists, who collectively guide the organization's AI agenda, and offer practical frameworks for designing effective CoE team architectures.

Keywords: AI Center of Excellence; Multidisciplinary teams; AI governance; Team architecture; Data science; Business strategy; Organizational AI maturity; Cross-functional integration

Multidisciplinary Team Architecture in AI Centers of Excellence: Integrating Data Science, Business Strategy, and Governance Expertise

1. Introduction

The rapid proliferation of artificial intelligence across industries has compelled organizations to rethink how they coordinate, govern, and operationalize AI capabilities at scale. While early AI adoption was often characterized by isolated pilot projects led by individual data science teams, the transition toward enterprise-wide AI deployment has necessitated more structured organizational mechanisms. Among these, the AI Center of Excellence (CoE) has gained prominence as a centralized or semi-centralized unit tasked with establishing standards, sharing best practices, and guiding the strategic direction of AI initiatives across business units (Ransbotham et al., 2019; Davenport & Ronanki, 2018). The CoE model draws on established precedents in technology management, where centers of excellence have long served as vehicles for accumulating specialized knowledge and diffusing it throughout complex organizations (Frost et al., 2002).

A distinguishing feature of AI CoEs, as opposed to traditional IT governance bodies, is their inherent multidisciplinarity. Effective AI deployment demands the convergence of technical expertise in machine learning and data engineering, strategic insight into business value creation, domain-specific knowledge of the organizational context, and governance capabilities spanning ethics, compliance, and risk management (Benbya et al., 2020; Raisch & Krakowski, 2021). This breadth of required expertise means that AI CoEs cannot be staffed exclusively by data scientists or software engineers; rather, they must integrate diverse professional backgrounds and organizational perspectives into a coherent operational unit. As Berente et al. (2021) and Kolbjørnsrud (2024) observe, managing AI and organizing human–AI collaboration require complementary technical and organizational capabilities; trust also influences how people work with AI systems (Glikson & Woolley, 2020).

Despite the growing practitioner consensus on the importance of multidisciplinary AI CoE teams, the academic literature on this topic remains underdeveloped. Most existing studies focus on either the technical dimensions of AI implementation or the strategic management of digital transformation, without specifically examining the team-level architecture of AI CoEs (Mikalef & Gupta, 2021). The composition of AI CoE teams, the mechanisms through which diverse expertise is integrated, and the relationship between team architecture and AI project outcomes have received limited empirical attention. This gap is particularly consequential because the design of CoE teams has direct implications for how organizations balance competing priorities such as innovation speed, regulatory compliance, and stakeholder alignment (Fountaine et al., 2019).

This study addresses the identified research gap by conducting a mixed-methods investigation of multidisciplinary team architecture within AI Centers of Excellence. The research combines a quantitative survey of 186 organizations across twelve industry sectors with semi-structured interviews of 34 senior AI leaders responsible for CoE operations. The study pursues three primary research objectives: first, to map the disciplinary composition and role allocation patterns within operational AI CoEs; second, to evaluate the relationship between team diversity, integration mechanisms, and AI project performance; and third, to assess how governance structures within CoEs influence the scalability and sustainability of AI initiatives. Through this investigation, the study provides empirical substantiation for the proposition that AI CoEs are typically composed of multidisciplinary experts, including data scientists, engineers, business strategists, and governance specialists, who collectively guide the organization's AI agenda, while also identifying actionable design principles for CoE team architecture.

The remainder of this article is organized as follows. Section 2 reviews the relevant literature on AI CoEs, multidisciplinary teams, data science capabilities, business strategy integration, and AI governance. Section 3 presents the research methodology, including the survey design, interview protocol, and analytical

procedures. Section 4 reports the empirical findings, supported by quantitative data analysis and qualitative thematic insights. Section 5 discusses the theoretical and practical implications of the results. Section 6 concludes with a summary of contributions, limitations, and directions for future research.

2. Literature Review

2.1 AI Centers of Excellence: Origins and Organizational Rationale

The concept of a Center of Excellence (CoE) originates in strategic management and knowledge management theory, where it has been conceptualized as an organizational unit that concentrates specialized expertise and disseminates it across the enterprise to enhance performance and innovation (Frost et al., 2002). In the context of information technology, CoEs have been established for enterprise resource planning, business intelligence, cybersecurity, and cloud computing, among other domains (Weill & Ross, 2004). The emergence of AI-specific CoEs reflects the distinctive challenges posed by AI adoption, which requires coordinating technical, strategic, and ethical considerations that span organizational boundaries and traditional functional hierarchies (Pumplun et al., 2019).

AI CoEs serve multiple organizational functions beyond technical capability building. They act as governance hubs that establish policies for responsible AI use, as knowledge brokers that connect data science teams with business stakeholders, and as strategic advisory bodies that help senior leadership prioritize AI investments and evaluate their returns (Jöhnk et al., 2021). The specific configuration of these functions varies across organizations depending on AI maturity, industry context, and organizational culture. Research by Enholm et al. (2022) on organizational AI adoption highlights the importance of coordinating technical and organizational capabilities to realize business value. Similarly, Ashta and Herrmann (2021) discuss opportunities and risks of AI in financial services, illustrating the need to align technological capabilities with domain expertise.

Recent empirical evidence highlights the role of organizational learning in translating AI investments into financial benefits. A global survey by Ransbotham et al. (2020) estimated that organizations which scaled AI by embedding it in business processes had a 39% likelihood of significant financial benefits, increasing to 73% when they also developed organizational learning with AI. However, the mere existence of a CoE is insufficient; its effectiveness depends critically on its composition, governance model, and integration with the broader organization (Fountaine et al., 2019). This observation motivates the present study's focus on the internal team architecture of AI CoEs as a determinant of AI deployment success.

2.2 Multidisciplinary Team Composition in Technology Management

The importance of multidisciplinary teams in technology-driven innovation has been extensively documented across management, engineering, and organizational behavior literatures. Cross-functional teams, which bring together members from different disciplinary backgrounds and organizational functions, have been shown to enhance problem-solving capability, reduce development cycle times, and improve the quality of complex technical outputs (Edmondson & Harvey, 2018). In the specific context of new product development, Bai et al. (2017) demonstrated that organizational structure and cross-functional integration jointly predict team performance, with integration mechanisms serving as critical mediators between structural design and outcomes.

AI projects present particularly compelling requirements for multidisciplinary collaboration because they span the full spectrum from data collection and model development to business deployment and ethical oversight. Davenport and Ronanki (2018) identified three distinct categories of AI capabilities: process automation, cognitive insight, and cognitive engagement, each requiring different combinations of technical, business, and interpersonal expertise. This categorization implies that AI CoEs must be staffed not only with technical specialists capable of building and deploying AI systems, but also with professionals who can

translate business requirements into analytical specifications, evaluate the ethical implications of AI outputs, and manage organizational change associated with AI-driven process transformations (Tambe et al., 2019).

The challenge of effective knowledge integration across disciplines within AI teams has been highlighted by several recent studies. Mikalef and Gupta (2021) found that organizations with stronger integrative capabilities, defined as the ability to combine diverse knowledge sources into actionable insights, achieved better outcomes from their AI investments. Shammar et al. (2021) examined security challenges in integrating blockchain and IoT, illustrating the breadth of expertise needed to manage complex technology deployments. The present study extends this literature by mapping the specific disciplinary composition of AI CoE teams and assessing how different compositions relate to project performance and governance effectiveness. Figure 1 illustrates the AI CoE architecture, its four functional pillars, and their interactions.

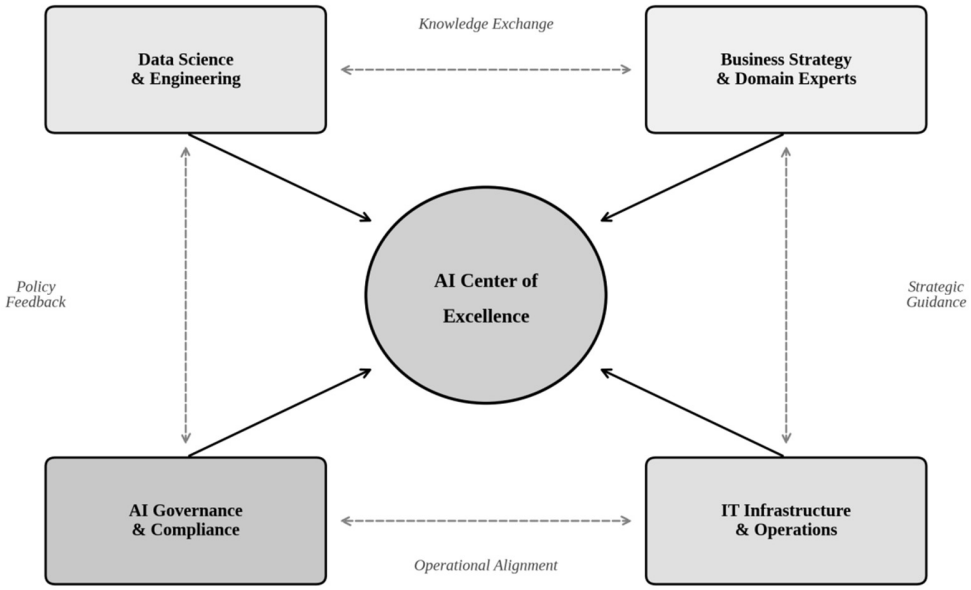


Figure 1. Conceptual Framework of the AI CoE Architecture with Four Functional Pillars and Interactions

2.3 Data Science and Analytics Capabilities

The data science function constitutes the technical core of most AI CoEs, encompassing expertise in statistical modeling, machine learning algorithm development, data engineering, and analytical tool implementation. The maturation of data science as a professional discipline has been accompanied by an expanding scope of required competencies, ranging from classical statistical inference to deep learning, natural language processing, computer vision, and reinforcement learning (Provost & Fawcett, 2013). Barredo Arrieta et al. (2020) review explainable AI models and applications, highlighting the importance of selecting and evaluating models that people can understand and manage.

Within AI CoEs, data scientists typically work alongside machine learning engineers, who focus on the operationalization of models in production environments, and data engineers, who manage the data infrastructure necessary for model training and deployment. This division of labor reflects the recognition that building an effective AI system requires not only algorithmic expertise but also robust data pipelines,

scalable computing infrastructure, and continuous monitoring capabilities (Sculley et al., 2015). The emergence of MLOps (Machine Learning Operations) as a discipline underscores the operational complexity of maintaining AI systems over time and the need for dedicated engineering roles within CoE structures (Kreuzberger et al., 2023). Recent work by Borges et al. (2021) on the strategic use of AI further discusses how analytical capabilities can inform strategic decision-making when integrated with business knowledge.

A critical insight from the data science literature is that technical capability alone does not guarantee organizational impact. Studies have consistently shown that the most common barriers to AI adoption are not technical but organizational, including difficulties in identifying appropriate use cases, integrating AI outputs into existing workflows, and securing stakeholder buy-in for AI-driven decisions (Ransbotham et al., 2019). This finding reinforces the importance of complementing data science expertise with business strategy and governance capabilities within the CoE team structure, a theme that the present study examines empirically.

2.4 Business Strategy Integration and Domain Expertise

The strategic alignment of AI initiatives with organizational goals is widely recognized as a critical success factor for AI adoption. Without clear linkage between AI capabilities and business value creation, organizations risk investing in technically impressive but operationally irrelevant AI projects that fail to generate measurable returns (Westerman et al., 2014). Business strategists within AI CoEs serve as translators between technical teams and organizational leadership, helping to identify high-value use cases, develop business cases for AI investments, and evaluate the competitive implications of AI-driven innovations (Iansiti & Lakhani, 2020).

Domain expertise represents another essential dimension of CoE team composition that complements both data science and strategic capabilities. Domain experts bring deep contextual knowledge of the specific industry, function, or process to which AI is being applied, enabling more accurate problem formulation, more relevant feature engineering, and more appropriate evaluation of model outputs (Paschen et al., 2020). In healthcare, for example, clinical domain expertise is essential for ensuring that AI-driven diagnostic tools meet medical standards and align with clinical workflows. Similarly, in financial services, domain expertise in regulatory compliance, market dynamics, and customer behavior is necessary for developing AI applications that are both effective and permissible (Goodell et al., 2021). Vincent (2021) emphasizes that organizational decision-making can integrate human judgment with AI, a principle that applies directly to the design of AI CoE teams (Bansal et al., 2021).

The interaction between business strategy and technical capability within AI CoEs has been modeled through various theoretical frameworks. The dynamic capabilities perspective, as applied by Mikalef and Gupta (2021), suggests that organizations must develop sensing, seizing, and transforming capabilities to extract value from AI investments. Within this framework, the CoE serves as the organizational locus where sensing (identifying opportunities), seizing (allocating resources and building solutions), and transforming (adapting organizational processes) capabilities converge. The multidisciplinary composition of the CoE is thus not merely an operational convenience but a strategic necessity for building the dynamic capabilities required for sustained AI-driven competitive advantage.

2.5 Governance, Ethics, and Regulatory Compliance

AI governance has emerged as a rapidly evolving field encompassing the policies, processes, and organizational structures through which AI systems are developed, deployed, monitored, and retired in accordance with ethical principles, legal requirements, and organizational standards. The increasing regulatory attention to AI, exemplified by the European Union's AI Act, national AI strategies across multiple jurisdictions, and sector-specific guidelines in healthcare and financial services, has elevated governance from a peripheral concern to a central operational requirement for AI CoEs (Mäntymäki et al., 2022).

Governance specialists within AI CoEs fulfill several distinct functions. They develop and maintain AI ethics guidelines and review processes, ensure compliance with applicable regulations and data protection requirements, establish model validation and monitoring frameworks, and coordinate stakeholder engagement around AI transparency and accountability (Jöhnk et al., 2021). The complexity of these functions is compounded by the cross-jurisdictional nature of many organizations' operations, which requires governance frameworks that can accommodate varying regulatory requirements across different geographic markets. Humayun et al. (2024) highlight cybersecurity challenges in IoT and AI deployments, supporting the need for dedicated security expertise within governance frameworks.

The interaction between governance and other CoE functions is bidirectional: governance requirements shape the boundaries within which data science and business strategy operate, while technical and strategic insights inform the development of governance policies that are both protective and enabling. This bidirectional relationship distinguishes effective AI governance from mere compliance, which tends to be reactive and restrictive rather than proactive and enabling (Ashok et al., 2022). Organizations that integrate governance as a core CoE function, rather than treating it as an external oversight mechanism, report higher levels of stakeholder trust, faster regulatory approval processes, and greater willingness among business units to adopt AI solutions (Floridi et al., 2018). The present study examines how governance maturity within AI CoEs relates to project success rates and organizational AI scalability.

3. Research Methodology

3.1 Research Design

This study employs a sequential explanatory mixed-methods design (Creswell & Plano Clark, 2018), combining a quantitative survey phase with a qualitative interview phase. The quantitative phase establishes the distribution of team composition patterns, identifies statistical relationships between team characteristics and performance outcomes, and provides a representative picture of AI CoE structures across industries. The qualitative phase deepens understanding of the mechanisms underlying the quantitative relationships, captures contextual factors that shape CoE design decisions, and surfaces emergent themes not anticipated in the survey instrument. This design is appropriate for the research objectives because it combines the breadth and generalizability of survey data with the depth and explanatory richness of interview data (Teddlie & Tashakkori, 2009).

3.2 Quantitative Phase: Survey Design and Administration

The survey instrument was developed through an iterative process involving literature review, expert consultation, and pilot testing. The instrument comprises five sections: organizational demographics, CoE structure and composition, role allocation and staffing, integration mechanisms, and performance outcomes. Composition and staffing items were developed based on role taxonomies from the literature on AI teams (Tambe et al., 2019; Davenport & Ronanki, 2018) and refined through consultations with eight AI practitioners from diverse industries. Performance was measured through six items capturing project completion rates, time-to-deployment, stakeholder satisfaction, scalability of AI solutions, and model accuracy improvement over time, each rated on a seven-point Likert scale.

The survey was administered electronically to senior AI leaders, including Chief AI Officers, VP of Analytics, Directors of Data Science, and AI CoE Heads, across organizations in twelve industries: financial services, healthcare, manufacturing, technology, retail, telecommunications, energy, pharmaceuticals, insurance, logistics, government, and education. Respondents were identified through professional networks, industry associations, and snowball sampling. A total of 312 invitations were distributed, yielding 198 completed responses (63.5% response rate), of which 186 passed data quality checks and were retained for analysis. The sample included organizations ranging from 500 to over 100,000 employees across North America (39%),

Europe (28%), Asia-Pacific (24%), and other regions (9%). Table 1 summarizes the demographics of the 186 survey respondents.

Table 1. Survey Respondent Demographics (N = 186)

Characteristic	Category	Count (%)
Industry Sector	Financial Services	42 (22.6%)
	Healthcare	28 (15.1%)
	Technology	31 (16.7%)
	Manufacturing	24 (12.9%)
	Others (8 sectors)	61 (32.8%)
Organization Size	500-5,000 employees	47 (25.3%)
	5,001-20,000	58 (31.2%)
	20,001-100,000	52 (28.0%)
	> 100,000	29 (15.6%)
AI CoE Maturity	< 1 year	31 (16.7%)
	1-3 years	78 (41.9%)
	3-5 years	52 (28.0%)
	> 5 years	25 (13.4%)
Geographic Region	North America	72 (38.7%)
	Europe	52 (28.0%)
	Asia-Pacific	45 (24.2%)
	Other regions	17 (9.1%)

3.3 Qualitative Phase: Semi-Structured Interviews

Following the survey analysis, semi-structured interviews were conducted with 34 senior AI leaders from organizations representing different CoE maturity levels, industry sectors, and geographic regions. Interview participants were purposively selected from survey respondents who indicated willingness to participate in follow-up interviews, with selection criteria ensuring diversity across organizational size, industry, and CoE maturity. Interviews lasted between 45 and 75 minutes and were conducted via video conference. The interview protocol explored themes including CoE formation history, team composition rationale, integration challenges and solutions, governance model evolution, and perceived success factors. All interviews were recorded, transcribed, and analyzed using thematic analysis following the six-phase framework of Braun and Clarke (2006).

3.4 Data Analysis Procedures

Quantitative data were analyzed using descriptive statistics, Pearson correlation analysis, multiple regression modeling, and hierarchical cluster analysis. Descriptive statistics characterized the distribution of team composition and governance patterns. Correlation analysis examined bivariate relationships between team diversity metrics and performance outcomes. Multiple regression models tested the predictive contribution of team composition, integration mechanisms, and governance maturity to AI project success, controlling for organizational size, industry, and CoE age. Cluster analysis identified archetypal team configurations across the sample. Qualitative data were analyzed through thematic analysis, with initial codes generated inductively from interview transcripts and subsequently organized into themes aligned with the research questions. Integration of quantitative and qualitative findings followed a joint display approach, where qualitative themes were mapped onto quantitative patterns to provide explanatory depth (Creswell & Plano Clark, 2018).

4. Findings and Analysis

4.1 Team Composition Patterns

The survey results reveal that AI CoEs are overwhelmingly multidisciplinary in composition, with 94.1% of responding organizations reporting at least three distinct professional roles within their CoE. Data scientists are the most prevalent role, present in 92% of CoEs, followed by machine learning engineers (87%), business analysts (78%), IT and DevOps engineers (73%), domain experts (71%), project managers (68%), governance specialists (64%), and dedicated ethics officers (41%). These findings confirm that the AI CoE is characteristically a multidisciplinary unit, and that data science, engineering, business strategy, and governance expertise are all represented in the majority of operational CoEs. Figure 2 presents the full distribution of role prevalence across the survey sample.

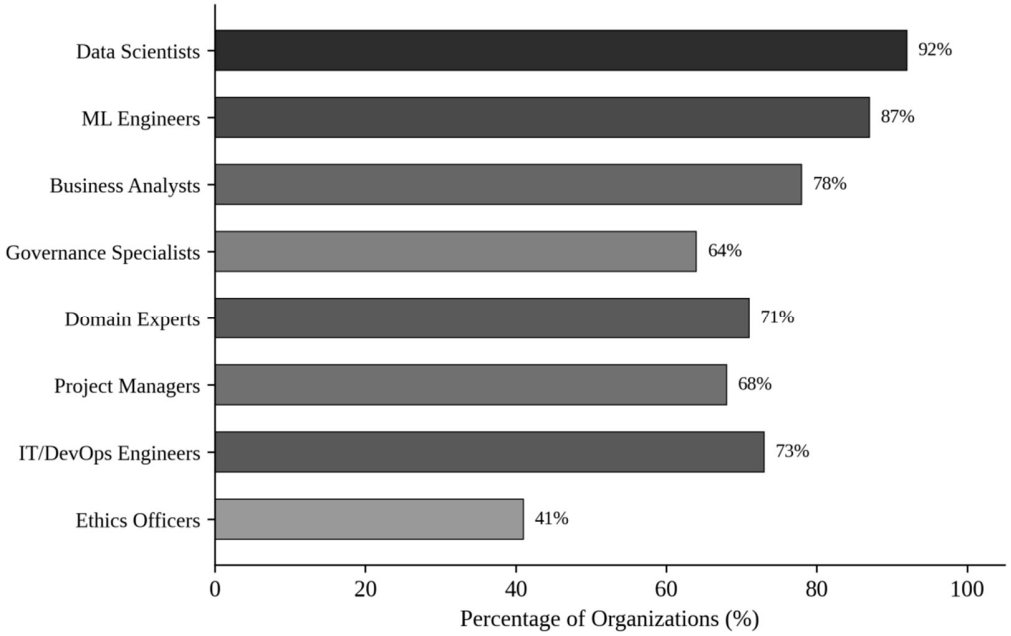


Figure 2. Prevalence of Professional Roles within AI Centers of Excellence across 186 Organizations Surveyed

The median CoE team size is 18 members (IQR: 11-32), with significant variation by organizational size and AI maturity. Organizations with more than 20,000 employees report median CoE sizes of 28 members, compared to 12 for organizations with fewer than 5,000 employees. More mature CoEs (over three years old) tend to have both larger teams and greater disciplinary diversity, suggesting that team composition evolves as organizations accumulate AI experience and recognize the need for complementary expertise. Interview data corroborate this pattern, with multiple participants describing an evolutionary trajectory from initially small, technically focused teams to progressively broader multidisciplinary structures as the CoE's mandate expanded to encompass governance, change management, and strategic advisory functions.

4.2 Role Allocation and Functional Distribution

To understand how staffing resources are distributed across functional domains within AI CoEs, respondents were asked to estimate the percentage of CoE personnel allocated to each of four functional categories: technical and data science, business strategy and domain expertise, governance and compliance, and project management and operations. The results, summarized in Table 2, reveal that technical staff constitute the largest share of CoE personnel on average, but that non-technical functions collectively represent a substantial proportion of CoE capacity.

Table 2. Average Functional Allocation of AI CoE Personnel by AI Maturity Stage

Functional Category	Early (< 1 yr)	Growing (1-3 yr)	Mature (3-5 yr)	Advanced (> 5 yr)
Technical / Data Science	62.4%	51.8%	44.3%	38.7%
Business Strategy / Domain	14.1%	20.6%	24.8%	26.2%
Governance / Compliance	8.3%	13.9%	18.5%	22.4%
Project Mgmt / Operations	15.2%	13.7%	12.4%	12.7%

As shown in Table 2, the proportion of CoE personnel allocated to technical and data science functions decreases from 62.4% in early-stage CoEs to 38.7% in advanced CoEs, while governance and compliance staffing increases from 8.3% to 22.4% over the same maturity trajectory. Business strategy and domain expertise allocation also increases significantly, from 14.1% to 26.2%. These trends indicate that as organizations gain AI experience, they increasingly recognize the importance of non-technical capabilities for sustaining AI value creation. Interview participants attributed this shift to several factors, including regulatory pressure, stakeholder demand for AI explainability, and accumulated experience with AI projects that failed due to poor business alignment rather than technical shortcomings.

4.3 Integration Mechanisms and Operating Models

Effective multidisciplinary teamwork requires not only diverse composition but also mechanisms for integrating diverse perspectives into coherent decision-making processes. The survey examined four categories of integration mechanisms: structural (cross-functional meetings, embedded liaisons), procedural (joint planning sessions, shared evaluation criteria), technological (integrated platforms, shared dashboards), and cultural (shared training programs, rotational assignments). Organizations with high integration scores, defined as those reporting frequent use of at least three mechanism categories, reported significantly higher AI project success rates (mean = 78.3%) compared to organizations with low integration scores (mean = 54.1%), a difference that was statistically significant ($t = 4.87$, $p < 0.001$).

The operating model adopted by the AI CoE also significantly influences integration effectiveness. The survey identified four primary operating models: centralized (all AI expertise concentrated in the CoE, 24.7% of respondents), federated (AI expertise distributed across business units with CoE providing coordination, 18.8%), hub-and-spoke (central CoE with embedded liaisons in business units, 36.0%), and hybrid (dynamic combination of approaches, 20.4%). Figure 3 compares the three integration effectiveness dimensions, namely knowledge sharing, decision speed, and innovation output, across the four operating models.

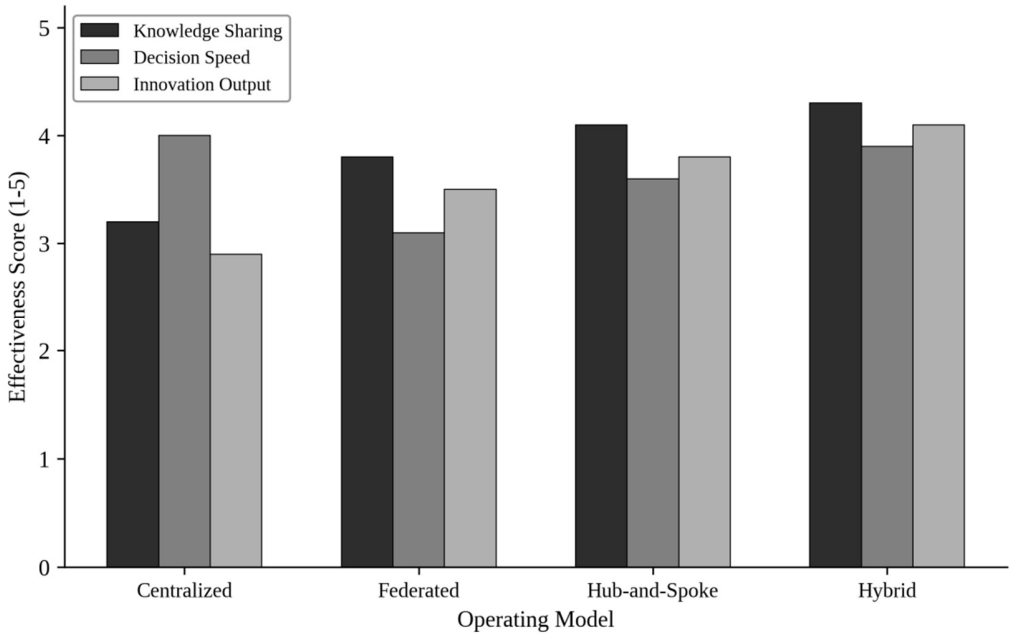


Figure 3. Integration Effectiveness of AI CoE Operating Models across Three Dimensions

The hub-and-spoke and hybrid models consistently outperform centralized and federated models on knowledge sharing and innovation output, while the centralized model excels in decision speed due to concentrated authority. Interview data suggest that the hub-and-spoke model achieves its integration advantages through embedded liaisons who serve as boundary spanners, translating between the central CoE's technical language and business units' operational priorities. One interview participant, a Chief AI Officer at a multinational financial services firm, described this role: the embedded AI leads understand both the data science possibilities and the business constraints, so they can prevent the misalignments that kill AI projects before they reach production.

4.4 Governance Maturity and AI Deployment Scalability

Governance maturity was assessed using a five-level framework adapted from established IT governance maturity models (Weill & Ross, 2004) and calibrated for AI-specific governance requirements. The five levels range from Level 1 (Ad Hoc: no formalized AI governance processes) through Level 3 (Defined: documented governance policies with assigned accountability) to Level 5 (Optimized: continuous improvement of governance processes based on quantitative feedback). The distribution across the sample was: Level 1 (12.4%), Level 2 (23.1%), Level 3 (32.8%), Level 4 (21.5%), and Level 5 (10.2%), indicating that the majority of organizations have at least foundational governance structures but that advanced governance maturity remains relatively uncommon.

The relationship between governance maturity and AI project success rate is illustrated in Figure 4. A strong positive correlation ($r = 0.72$, $p < 0.001$) was observed between governance maturity level and the proportion of AI projects that successfully progressed from pilot to production deployment. Organizations at governance Level 5 reported a mean deployment success rate of 82.6%, compared to 38.4% at Level 1. Multiple regression analysis, controlling for organizational size, industry sector, CoE age, and team size, confirmed that governance maturity was the strongest single predictor of deployment success (standardized beta = 0.41, $p < 0.001$), exceeding the predictive contribution of team size (beta = 0.18, $p < 0.05$) and technical staffing ratio (beta = 0.22, $p < 0.01$).

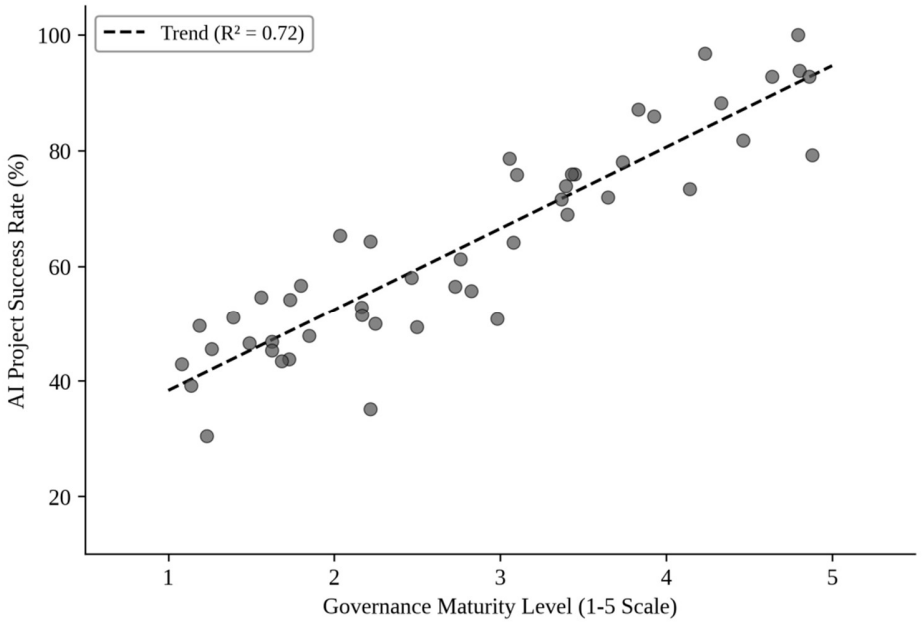


Figure 4. Relationship between Governance Maturity and AI Project Success Rate

This finding has important implications for CoE team design. It suggests that investing in governance capabilities, including dedicating senior personnel to ethics review, compliance monitoring, and stakeholder communication, yields greater returns in terms of deployment scalability than investing exclusively in technical talent. Interview participants reinforced this interpretation, with several senior leaders noting that governance shortcomings, rather than technical limitations, were the primary bottleneck in scaling AI beyond pilot stages. As one VP of Data Science in the healthcare sector stated, we had excellent models sitting in notebooks for months because we lacked the governance infrastructure to approve them for clinical use.

4.5 Archetypal Team Configurations

Hierarchical cluster analysis, using Ward's method with Euclidean distance on the functional allocation variables from Table 2, identified four distinct cluster archetypes across the 186 organizations. These archetypes, presented in Table 3, represent characteristic patterns of CoE team configuration that correspond to different organizational contexts and AI maturity stages.

Table 3. Four Archetypal AI CoE Team Configurations Identified through Cluster Analysis

Attribute	Tech-Centric (n=52)	Balanced (n=67)	Governance-Led (n=38)
Technical staff %	65.2%	45.8%	35.1%
Business staff %	14.8%	24.3%	22.7%
Governance staff %	7.6%	17.2%	28.9%
Median team size	14	22	26
Avg. project success	56.3%	74.8%	81.2%
Dominant industry	Tech / Startups	Cross-sector	Finance / Health
Typical AI maturity	Early	Growing-Mature	Mature-Advanced

Note: A fourth cluster (Strategy-Driven, n=29) with 31.4% business staff, 42.1% technical, and 15.3% governance was also identified, predominantly in retail and consumer sectors.

The Tech-Centric archetype, characterized by heavily technical staffing and smaller teams, is most prevalent among early-stage CoEs in technology companies and startups, where rapid model development takes priority over governance infrastructure. The Balanced archetype, the most common configuration (36.0% of

the sample), distributes resources more evenly across functions and is associated with growing-to-mature CoEs in diverse industry settings. The Governance-Led archetype, predominant in regulated industries such as financial services and healthcare, allocates the highest proportion of staff to governance and compliance functions and achieves the highest average project success rate. These patterns suggest that CoE team architecture should be calibrated to the organization's regulatory environment, AI maturity, and strategic priorities rather than following a universal template.

4.6 Correlation Analysis: Team Diversity and Performance Outcomes

To quantify the relationship between multidisciplinary team composition and AI project outcomes, a disciplinary diversity index (DDI) was computed for each organization using the Blau heterogeneity index applied to the distribution of personnel across seven role categories. The DDI ranges from 0 (all personnel in a single role) to 1 (equal distribution across all roles). The sample mean DDI was 0.74 (SD = 0.12), indicating generally high diversity. Pearson correlation analysis revealed statistically significant positive associations between DDI and multiple performance outcomes: project success rate ($r = 0.68$, $p < 0.001$), stakeholder satisfaction ($r = 0.71$, $p < 0.001$), model reusability ($r = 0.54$, $p < 0.001$), and time-to-production ($r = -0.47$, $p < 0.001$, negative indicating faster deployment with higher diversity).

Multiple regression analysis confirmed that DDI remains a significant predictor of project success after controlling for organizational size, industry sector, CoE age, and total AI investment. The full model explained 61.3% of the variance in project success rate (adjusted R-squared = 0.589, $F(7, 178) = 40.26$, $p < 0.001$). The standardized coefficients indicate that governance maturity (beta = 0.41), DDI (beta = 0.32), and integration mechanism usage (beta = 0.24) are the three strongest predictors, collectively accounting for the majority of explained variance. These results provide robust empirical support for the claim that multidisciplinary team composition is a significant contributor to AI CoE effectiveness, while also highlighting that team diversity must be accompanied by governance maturity and effective integration mechanisms to realize its full potential.

5. Discussion

5.1 Theoretical Implications

The findings of this study contribute to the intersection of organizational design theory, technology management, and AI governance literature. First, the empirical documentation of team composition patterns across 186 AI CoEs provides the most comprehensive mapping to date of the professional roles and functional distributions within these organizational units. The results confirm and extend the practitioner-oriented descriptions of AI CoE staffing found in industry reports (Fountaine et al., 2019; Ransbotham et al., 2020) by demonstrating statistically significant relationships between team composition characteristics and measurable performance outcomes. The strong positive correlation between disciplinary diversity and project success ($r = 0.68$) aligns with the broader multidisciplinary team literature (Edmondson & Harvey, 2018) while specifying the contextual conditions under which diversity yields performance benefits in AI settings.

Second, the identification of governance maturity as the strongest single predictor of AI deployment scalability extends the AI governance literature by establishing an empirical link between governance capability and operational outcomes. Previous studies have primarily addressed AI governance at the policy or ethical framework level (Floridi et al., 2018), with limited attention to the organizational capability dimension. The present findings suggest that governance should be conceptualized not merely as a compliance function but as a core operational capability that enables the scaling of AI from pilot to production. This perspective is consistent with the evolving understanding of AI as a sociotechnical system whose deployment requires coordinated attention to technical, organizational, and institutional factors (Papagiannidis et al., 2023). The emerging significance of responsible AI practices also connects to broader discussions of ethics-based and internal algorithmic auditing (Mökander et al., 2021; Raji et al., 2020).

Third, the four archetypal team configurations identified through cluster analysis offer a typological framework that connects CoE team design to organizational context. The finding that governance-led configurations achieve the highest project success rates in regulated industries, while tech-centric configurations predominate in early-stage CoEs, suggests that optimal team architecture is contingent on environmental and developmental factors. This contingency perspective echoes foundational organizational design theory (Galbraith, 1974) and extends it to the specific domain of AI organizational structures. The implication is that AI CoE design guidelines should be adaptive rather than prescriptive, offering different team composition recommendations for different organizational contexts and maturity stages.

5.2 Practical Implications

For organizational leaders responsible for establishing or evolving AI CoEs, the findings offer several actionable insights. First, the consistent relationship between disciplinary diversity and project success underscores the importance of resisting the temptation to staff CoEs exclusively with technical specialists. While data scientists and machine learning engineers form the essential technical core, the inclusion of business strategists, governance specialists, and domain experts is associated with meaningfully better outcomes. Organizations in the early stages of CoE development may begin with a technical focus but should plan for deliberate diversification of the team composition as the CoE matures, as reflected in the maturity-dependent staffing patterns documented in Table 2.

Second, the strong predictive contribution of governance maturity to deployment scalability suggests that organizations should invest in governance capabilities early in the CoE lifecycle, rather than treating governance as an afterthought to be addressed once technical capabilities are established. Practical governance investments include appointing a dedicated governance lead within the CoE, establishing model review and approval processes, developing AI risk assessment frameworks, and implementing monitoring systems for deployed models. The discussion by Gauthier and Brender (2021) of blockchain technology and auditing points to emerging technological tools that may support AI governance infrastructure within CoEs.

Third, the hub-and-spoke operating model emerges as a particularly effective configuration for balancing centralized expertise with business unit responsiveness. Organizations considering CoE structural design should evaluate the feasibility of embedded liaison roles that bridge the central CoE and operational units, as these boundary-spanning positions were consistently identified as enablers of cross-functional integration in both the survey and interview data. The work by Saad et al. (2020) on 6G technologies further illustrates how emerging technological capabilities may reshape the communication and coordination infrastructure available to distributed CoE teams. Additionally, advances in quantum algorithms reviewed by Bharti et al. (2022) suggest that CoE teams may need to monitor frontier developments that could reshape organizational analytics capabilities.

5.3 Limitations and Future Research

This study is subject to several limitations that suggest directions for future research. The cross-sectional survey design captures associations between team characteristics and outcomes but does not establish causal relationships; longitudinal studies tracking CoE evolution and performance over time would strengthen causal inference. The reliance on self-reported performance measures introduces potential common method bias, which could be mitigated in future studies through the incorporation of objective performance metrics such as model deployment counts, revenue attribution, and audit compliance records. The sample, while diverse across industries and geographies, may be subject to selection bias if organizations with more successful CoEs were more willing to participate. Future research should also examine the internal dynamics of CoE teams, including leadership styles, conflict resolution processes, and knowledge transfer mechanisms, to better understand how diverse expertise is translated into effective collaborative outcomes. As Zhu et al. (2024) discuss regarding Web 3.0 applications, the continued evolution of digital infrastructure will likely create new demands on CoE team architecture that warrant ongoing investigation. Similarly, advances in

quantum machine learning discussed by Huang et al. (2021) and Cerezo et al. (2022) may ultimately necessitate new specialist roles within CoEs as quantum machine learning matures into a deployable enterprise capability.

6. Conclusion

This study has provided comprehensive empirical evidence documenting the multidisciplinary team architecture of AI Centers of Excellence across 186 organizations in twelve industries. The findings substantiate the claim that the AI CoE is typically composed of multidisciplinary experts, including data scientists, engineers, business strategists, and governance specialists, who collectively guide the organization's AI agenda. Specifically, the study demonstrates that data scientists are present in 92% of surveyed CoEs, machine learning engineers in 87%, business analysts in 78%, and governance specialists in 64%, confirming the inherently cross-disciplinary nature of these organizational units.

Beyond documenting composition patterns, the study reveals that higher disciplinary diversity within CoEs is significantly associated with improved AI project outcomes, including higher success rates, greater stakeholder satisfaction, and faster time-to-production. Governance maturity emerges as the strongest single predictor of AI deployment scalability, suggesting that organizations should prioritize governance capability development alongside technical talent acquisition. The four archetypal team configurations identified through cluster analysis offer a practical typology for calibrating CoE design to organizational context, with governance-led configurations achieving the highest performance in regulated industries and balanced configurations demonstrating broad applicability across sectors.

These findings carry implications for both organizational practice and academic understanding of AI adoption. For practitioners, the results provide evidence-based guidance for CoE team design, staffing prioritization, and governance investment. For researchers, the study advances the theoretical understanding of how organizational design principles apply to the distinctive challenges of AI governance and deployment, and identifies productive avenues for future investigation, including longitudinal studies of CoE evolution, examination of internal team dynamics, and assessment of emerging technology impacts on CoE requirements. As organizations continue to expand their AI ambitions, the architecture of their AI Centers of Excellence will remain a critical determinant of whether those ambitions translate into sustainable, responsible, and scalable value creation.

Acknowledgement

The authors sincerely thank the three anonymous reviewers for their constructive and insightful comments, which have significantly improved the theoretical rigor of the organizational framework, the clarity of the mixed-methods analysis, and the practical value of the research conclusions. We also extend our gratitude to the School of Economics and Management at Beijing Forestry University for providing the research environment, administrative support, and academic resources necessary for this work. Special thanks are due to all participating organizations and senior AI leaders who took time to complete the survey and participate in the interviews, without whose generous input this empirical investigation would not have been possible. We also acknowledge the editorial team of the journal for their professional and efficient handling of the submission and peer review process.

Author Contributions

Conceptualization, S.K.W. and B.Z.; Methodology, S.K.W. and B.Z.; Formal Analysis, S.K.W. and B.Z.; Investigation, S.K.W. and B.Z.; Data Curation, S.K.W. and B.Z.; Software, S.K.W.; Validation, S.K.W. and B.Z.; Resources, S.K.W. and B.Z.; Visualization, S.K.W.; Writing – Original Draft, S.K.W.; Writing – Review & Editing, S.K.W. and B.Z.; Supervision, B.Z.; Project Administration, S.K.W. and B.Z. We have read and approved the final manuscript and take responsibility for the work.

Funding

This research did not receive any financial support from external funding institutions.

Conflicts of Interest

The authors declare that there are no financial or professional conflicts of interest that could affect the objectivity of the research design, data analysis, or conclusion presentation.

Ethical Approval and Consent

This study was conducted in accordance with ethical research guidelines for social science investigations. All survey and interview participants provided informed consent prior to participation, and all response data were anonymized and aggregated for analysis to protect participant confidentiality. No identifiable personal or organizational sensitive information was retained in the final dataset. The study did not involve human biological subjects, animal experiments, or clinical trials

Data Availability Statement

The core statistical indicators, model coefficients, and performance metrics supporting the findings of this study are all presented in the tables and figures within the main text. The raw survey and interview dataset contains organization-specific information and is subject to confidentiality agreements with participants, so it cannot be publicly disclosed in its original form. The analysis code for statistical modeling and thematic coding can be requested from the corresponding author on reasonable request, subject to compliance with data use requirements.

Declaration of Generative AI Use

Generative artificial intelligence tools were used to assist with language polishing of the manuscript text and the initial generation of schematic diagrams, including the conceptual framework diagram, operating model comparison chart, and archetypal configuration table. All AI-assisted content has been carefully reviewed and verified by the authors, who take full responsibility for the final content of the manuscript.

References

Ashok, M., Madan, R., Joha, A., & Sivarajah, U. (2022). Ethical framework for artificial intelligence and digital technologies. International Journal of Information Management, 62, 102433. https://doi.org/10.1016/j.ijinfomgt.2021.102433

Ashta, A., & Herrmann, H. (2021). Artificial intelligence and fintech: An overview of opportunities and risks for banking, investments, and microfinance. Strategic Change, 30(3), 211–222. https://doi.org/10.1002/jsc.2404

Bai, W., Feng, Y., Yue, Y., & Feng, L. (2017). Organizational structure, cross-functional integration and performance of new product development team. Procedia Engineering, 174, 621–629. https://doi.org/10.1016/j.proeng.2017.01.198

Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M. T., & Weld, D. S. (2021). Does the whole exceed its parts? The effect of AI explanations on complementary team performance. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Article 81, pp. 1–16). https://doi.org/10.1145/3411764.3445717

Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Benbya, H., Davenport, T. H., & Pachidi, S. (2020). Special issue editorial: Artificial intelligence in organizations: Current state and future opportunities. MIS Quarterly Executive, 19(4), ix–xxi. https://aisel.aisnet.org/misqe/vol19/iss4/4

Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. MIS Quarterly, 45(3), 1433–1450. https://doi.org/10.25300/MISQ/2021/16274

Bharti, K., Cervera-Lierta, A., Kyaw, T. H., Haug, T., Alperin-Lea, S., Anand, A., ... & Aspuru-Guzik, A. (2022). Noisy intermediate-scale quantum algorithms. Reviews of Modern Physics, 94(1), 015004. https://doi.org/10.1103/RevModPhys.94.015004

Borges, A. F. S., Laurindo, F. J. B., Spínola, M. M., Gonçalves, R. F., & Mattos, C. A. (2021). The strategic use of artificial intelligence in the digital era: Systematic literature review and future research directions. International Journal of Information Management, 57, 102225. https://doi.org/10.1016/j.ijinfomgt.2020.102225

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative research in psychology, 3(2), 77-101. https://doi.org/10.1191/1478088706qp063oa

Cerezo, M., Verdon, G., Huang, H.-Y., Cincio, L., & Coles, P. J. (2022). Challenges and opportunities in quantum machine learning. Nature Computational Science, 2, 567–576. https://doi.org/10.1038/s43588-022-00311-3

Creswell, J. W., & Plano Clark, V. L. (2018). Designing and conducting mixed methods research (3rd ed.). SAGE Publications.

Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. Harvard Business Review, 96(1), 108–116.

Edmondson, A. C., & Harvey, J.-F. (2018). Cross-boundary teaming for innovation: Integrating research on teams and knowledge in organizations. Human Resource Management Review, 28(4), 347–360. https://doi.org/10.1016/j.hrmr.2017.03.002

Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. Information Systems Frontiers, 24, 1709–1734. https://doi.org/10.1007/s10796-021-10186-w

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2018). AI4People: An ethical framework for a good AI society. Minds and Machines, 28(4), 689–707. https://doi.org/10.1007/s11023-018-9482-5

Fountaine, T., McCarthy, B., & Saleh, T. (2019). Building the AI-powered organization. Harvard Business Review, 97(4), 62–73.

Frost, T. S., Birkinshaw, J. M., & Ensign, P. C. (2002). Centers of excellence in multinational corporations. Strategic Management Journal, 23(11), 997–1018. https://doi.org/10.1002/smj.273

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

Gauthier, M. P., & Brender, N. (2021). How do the current auditing standards fit the emergent use of blockchain? Managerial Auditing Journal, 36(3), 365–385. https://doi.org/10.1108/MAJ-12-2019-2513

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

Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. Journal of Behavioral and Experimental Finance, 32, 100577. https://doi.org/10.1016/j.jbef.2021.100577

Huang, H.-Y., Broughton, M., Mohseni, M., Babbush, R., Boixo, S., Neven, H., & McClean, J. R. (2021). Power of data in quantum machine learning. Nature Communications, 12, 2631. https://doi.org/10.1038/s41467-021-22539-9

Humayun, M., Tariq, N., Alfayad, M., Zakwan, M., Alwakid, G., & Assiri, M. (2024). Securing the Internet of Things in artificial intelligence era: A comprehensive survey. IEEE Access, 12, 25469–25490. https://doi.org/10.1109/ACCESS.2024.3365634

Iansiti, M., & Lakhani, K. R. (2020). Competing in the age of AI: Strategy and leadership when algorithms and networks run the world. Harvard Business Press.

Jöhnk, J., Weißert, M., & Wyrtki, K. (2021). Ready or not, AI comes: An interview study of organizational AI readiness factors. Business & Information Systems Engineering, 63(1), 5–20. https://doi.org/10.1007/s12599-020-00676-7

Kolbjørnsrud, V. (2024). Designing the intelligent organization: Six principles for human-AI collaboration. California Management Review, 66(2), 44–64. https://doi.org/10.1177/00081256231211020

Kreuzberger, D., Kühl, N., & Hirschl, S. (2023). Machine learning operations (MLOps): Overview, definition, and architecture. IEEE Access, 11, 31866–31879. https://doi.org/10.1109/ACCESS.2023.3262138

Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. AI and Ethics, 2(4), 603–609. https://doi.org/10.1007/s43681-022-00143-x

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

Mökander, J., Morley, J., Taddeo, M., & Floridi, L. (2021). Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. Science and Engineering Ethics, 27, 44. https://doi.org/10.1007/s11948-021-00319-4

Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2023). Toward AI governance: Identifying best practices and potential barriers and outcomes. Information Systems Frontiers, 25, 123–141. https://doi.org/10.1007/s10796-022-10251-y

Paschen, J., Wilson, M., & Ferreira, J. J. (2020). Collaborative intelligence: How human and artificial intelligence create value along the B2B sales funnel. Business Horizons, 63(3), 403–414. https://doi.org/10.1016/j.bushor.2020.01.003

Provost, F., & Fawcett, T. (2013). Data science and its relationship to big data and data-driven decision making. Big Data, 1(1), 51–59. https://doi.org/10.1089/big.2013.1508

Pumplun, L., Tauchert, C., & Heidt, M. (2019). A new organizational chassis for artificial intelligence: Exploring organizational readiness factors. In Proceedings of the 27th European Conference on Information Systems (ECIS). https://aisel.aisnet.org/ecis2019_rp/106/

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

Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (pp. 33–44). https://doi.org/10.1145/3351095.3372873

Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., & Kiron, D. (2019). Winning with AI. MIT Sloan management review.

Ransbotham, S., Khodabandeh, S., Kiron, D., Candelon, F., Chu, M., & LaFountain, B. (2020). Expanding AI's impact with organizational learning.

Saad, W., Bennis, M., & Chen, M. (2020). A vision of 6G wireless systems: Applications, trends, technologies, and open research problems. IEEE Network, 34(3), 134–142. https://doi.org/10.1109/MNET.001.1900287

Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... & Dennison, D. (2015). Hidden technical debt in machine learning systems. Advances in Neural Information Processing Systems, 28, 2503–2511.

Shammar, E. A., Zahary, A. T., & Al-Shargabi, A. A. (2021). A survey of IoT and blockchain integration: Security perspective. IEEE Access, 9, 156114–156150. https://doi.org/10.1109/ACCESS.2021.3129697

Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. California Management Review, 61(4), 15–42. https://doi.org/10.1177/0008125619867910

Teddlie, C., & Tashakkori, A. (2009). Foundations of mixed methods research: Integrating quantitative and qualitative approaches in the social and behavioral sciences. SAGE Publications.

Vincent, V. U. (2021). Integrating intuition and artificial intelligence in organizational decision-making. Business Horizons, 64(4), 425–438. https://doi.org/10.1016/j.bushor.2021.02.008

Weill, P., & Ross, J. W. (2004). IT governance: How top performers manage IT decision rights for superior results. Harvard Business Press.

Westerman, G., Bonnet, D., & McAfee, A. (2014). Leading digital: Turning technology into business transformation. Harvard Business Press.

Zhu, J., Li, F., & Chen, J. (2024). A survey of blockchain, artificial intelligence, and edge computing for Web 3.0. Computer Science Review, 54, 100667. https://doi.org/10.1016/j.cosrev.2024.100667